Annals of Mathematics, **161** (2005), 1–59

The Tits alternative for $\mathrm{Out}(F_n)$ II: A Kolchin type theorem

By Mladen Bestvina, Mark Feighn, and Michael Handel*

Abstract

This is the second of two papers in which we prove the Tits alternative for $\mathrm{Out}(F_n)$.

Contents

*The authors gratefully acknowledge the support of the National Science Foundation.

1. Introduction and outline

Recent years have seen a development of the theory for $\mathrm{Out}(F_n)$, the outer automorphism group of the free group F_n of rank n, that is modeled on Nielsen-Thurston theory for surface homeomorphisms. As mapping classes have either exponential or linear growth rates, so free group outer automorphisms have either exponential or polynomial growth rates. (The degree of the polynomial can be any integer between 1 and $n-1$; see [BH92].) In [BFH00], we considered individual automorphisms with primary emphasis on those with exponential growth rates. In this paper, we focus on subgroups of $\mathrm{Out}(F_n)$ all of whose elements have polynomial growth rates.

To remove certain technicalities arising from finite order phenomena, we restrict our attention to those outer automorphisms of polynomial growth whose induced automorphism of $H_1(F_n;\mathbb{Z}) \cong \mathbb{Z}^n$ is unipotent. We say that such an outer automorphism is *unipotent*. The subset of unipotent outer automorphisms of F_n is denoted $\mathrm{UPG}(F_n)$ (or just UPG). A subgroup of $\mathrm{Out}(F_n)$ is *unipotent* if each element is unipotent. We prove (Proposition 3.5) that any polynomially growing outer automorphism that acts trivially in $\mathbb{Z}/3\mathbb{Z}$-homology is unipotent. Thus every subgroup of polynomially growing outer automorphisms has a finite index unipotent subgroup.

The archetype for the main theorem of this paper comes from linear groups. A linear map is unipotent if and only if it has a basis with respect to which it is upper triangular with 1's on the diagonal. A celebrated theorem of Kolchin [Ser92] states that for any group of unipotent linear maps there is a basis with respect to which all elements of the group are upper triangular with 1's on the diagonal.

There is an analogous result for mapping class groups. We say that a mapping class is unipotent if it has linear growth and if the induced linear map on first homology is unipotent. The Thurston classification theorem implies that a mapping class is unipotent if and only if it is represented by a composition of Dehn twists in disjoint simple closed curves. Moreover, if a pair of unipotent mapping classes belongs to a unipotent subgroup, then their twisting curves cannot have transverse intersections (see for example [BLM83]). Thus every unipotent mapping class subgroup has a characteristic set of disjoint simple closed curves and each element of the subgroup is a composition of Dehn twists along these curves.

Our main theorem is the analogue of Kolchin's theorem for $\mathrm{Out}(F_n)$. Fix once-and-for-all a wedge Rose_n of n circles and permanently identify its fundamental group with F_n. A marked graph (of rank n) is a graph equipped with a homotopy equivalence from Rose_n; see [CV86]. A homotopy equivalence $f : G \to G$ on a marked graph G induces an outer automorphism of the fundamental group of G and therefore an element $\mathcal{O}$ of $\mathrm{Out}(F_n)$; we say that $f : G \to G$ is a *representative* of $\mathcal{O}$.

Suppose that G is a marked graph and that $\emptyset = G_0 \subsetneq G_1 \subsetneq \cdots \subsetneq G_K = G$ is a filtration of G where G_i is obtained from G_{i-1} by adding a single edge E_i. A homotopy equivalence $f : G \to G$ is *upper triangular with respect to the filtration* if each $f(E_i) = v_i E_i u_i$ (as edge paths) where u_i and v_i are closed paths in G_{i-1}. If the choice of filtration is clear then we simply say that $f : G \to G$ is *upper triangular*. We refer to the u_i's and v_i's as *suffixes* and *prefixes* respectively.

An outer automorphism is unipotent if and only if it has a representative that is upper triangular with respect to some filtered marked graph G (see Section 3).

For any filtered marked graph G, let $\mathcal{Q}$ be the set of upper triangular homotopy equivalences of G up to homotopy relative to the vertices of G. By Lemma 6.1, $\mathcal{Q}$ is a group under the operation induced by composition. There is a natural map from $\mathcal{Q}$ to $\mathrm{UPG}(F_n)$. We say that a unipotent subgroup of $\mathrm{Out}(F_n)$ is *filtered* if it lifts to a subgroup of $\mathcal{Q}$ for some filtered marked graph.

We denote the conjugacy class of a free factor F^i by $[[F^i]]$. If $F^1 * F^2 * \cdots * F^k$ is a free factor, then we say that the collection $\mathcal{F} = \{[[F^1]], [[F^2]], \ldots, [[F^k]]\}$ is a *free factor system*. There is a natural action of $\mathrm{Out}(F_n)$ on free factor systems and we say that $\mathcal{F}$ is $\mathcal{H}$-invariant if each element of the subgroup $\mathcal{H}$ fixes $\mathcal{F}$. A (not necessarily connected) subgraph K of a marked real graph determines a free factor system $\mathcal{F}(K)$. A partial order on free factor systems is defined in subsection 2.8.

We can now state our main theorem.

THEOREM 1.1 (Kolchin theorem for $\mathrm{Out}(F_n)$). *Every finitely generated unipotent subgroup $\mathcal{H}$ of $\mathrm{Out}(F_n)$ is filtered. For any $\mathcal{H}$-invariant free factor system $\mathcal{F}$, the marked filtered graph G can be chosen so that $\mathcal{F}(G_r) = \mathcal{F}$ for some filtration element G_r. The number of edges of G can be taken to be bounded by $\frac{3n}{2} - 1$ for $n > 1$.*

It is an interesting question whether or not the requirement that $\mathcal{H}$ be finitely generated is necessary or just an artifact of our proof.

Question. Is every unipotent subgroup of $\mathrm{Out}(F_n)$ contained in a finitely generated unipotent subgroup?

Remark 1.2. In contrast to unipotent mapping class subgroups which are all finitely generated and abelian, unipotent subgroups of $\mathrm{Out}(F_n)$ can be quite large. For example, if G is a wedge of n circles, then a filtration on G corresponds to an ordered basis $\{e_1, \dots, e_n\}$ of F_n and elements of $\mathcal{Q}$ correspond to automorphisms of the form $e_i \mapsto a_i e_i b_i$ with $a_i, b_i \in \langle e_1, \dots, e_{i-1}\rangle$. When $n > 2$, the image of $\mathcal{Q}$ in $\mathrm{UPG}(F_n)$ contains a product of nonabelian free groups.

This is the second of two papers in which we establish the Tits alternative for $\mathrm{Out}(F_n)$.

THEOREM (The Tits alternative for $\mathrm{Out}(F_n)$). *Let $\mathcal{H}$ be any subgroup of* $\mathrm{Out}(F_n)$. *Then either $\mathcal{H}$ is virtually solvable, or contains a nonabelian free group.*

For a proof of a special (generic) case, see [BFH97a]. The following corollary of Theorem 1.1 gives another special case of the Tits alternative for $\mathrm{Out}(F_n)$. The corollary is then used to prove the full Tits alternative.

COROLLARY 1.3. *Every unipotent subgroup $\mathcal{H}$ of* $\mathrm{Out}(F_n)$ *either contains a nonabelian free group or is solvable.*

Proof. We first prove that if $\mathcal{Q}$ is defined as above with respect to a marked filtered graph G, then every subgroup $\mathcal{Z}$ of $\mathcal{Q}$ either contains a nonabelian free group or is solvable.

Let $i \geq 0$ be the largest parameter value for which every element of $\mathcal{Z}$ restricts to the identity on G_{i-1}. If $i = K + 1$, then $\mathcal{Z}$ is the trivial group and we are done. Suppose then that $i \leq K$. By construction, each element of $\mathcal{Z}$ satisfies $E_i \mapsto v_i E_i u_i$ where v_i and u_i are paths (that depend on the element of $\mathcal{Z}$) in G_{i-1} and are therefore fixed by every element of $\mathcal{Z}$. The suffix map $\mathcal{S} : \mathcal{Z} \to F_n$, which assigns the suffix u_i to the element of $\mathcal{Z}$, is therefore a homomorphism. The prefix map $\mathcal{P} : \mathcal{Z} \to F_n$, which assigns the inverse of v_i to the element of $\mathcal{Z}$, is also a homomorphism.

If the image of $\mathcal{P} \times \mathcal{S} : \mathcal{Z} \to F_n \times F_n$ contains a nonabelian free group, then so does $\mathcal{Z}$ and we are done. If the image of $\mathcal{P} \times \mathcal{S}$ is abelian then, since $\mathcal{Z}$ is an abelian extension of the kernel of $\mathcal{P} \times \mathcal{S}$, it suffices to show that the kernel of $\mathcal{P} \times \mathcal{S}$ is either solvable or contains a nonabelian free group. Upward induction on i now completes the proof. In fact, this argument shows that $\mathcal{Z}$ is polycyclic and that the length of the derived series is bounded by $\frac{3n}{2} - 1$ for $n > 1$.

For $\mathcal{H}$ finitely generated the corollary now follows from Theorem 1.1. When $\mathcal{H}$ is not finitely generated, it can be represented as the increasing union of finitely generated subgroups. If one of these subgroups contains a nonabelian free group, then so does $\mathcal{H}$, and if not then $\mathcal{H}$ is solvable with the length of the derived series bounded by $\frac{3n}{2} - 1$. ∎

Proof of the Tits alternative for $\mathrm{Out}(F_n)$. Theorem 7.0.1 of [BFH00] asserts that if $\mathcal{H}$ does not contain a nonabelian free group then there is a finite index subgroup $\mathcal{H}_0$ of $\mathcal{H}$ and an exact sequence

$$1 \to \mathcal{H}_1 \to \mathcal{H}_0 \to \mathcal{A} \to 1$$

with $\mathcal{A}$ a finitely generated free abelian group and with $\mathcal{H}_1$ a unipotent subgroup of $\mathrm{Out}(F_n)$. Since $\mathcal{H}_1$ does not contain a nonabelian free group, by Corollary 1.3, $\mathcal{H}_1$ is solvable. Thus, $\mathcal{H}_0$ is solvable and $\mathcal{H}$ is virtually solvable. □

In [BFH04] we strengthen the Tits alternative for $\mathrm{Out}(F_n)$ further by proving:

THEOREM (Solvable implies abelian). *A solvable subgroup of* $\mathrm{Out}(F_n)$ *has a finitely generated free abelian subgroup of index at most* 3^{5n^2}.

Emina Alibegović [Ali02] has since provided an alternate shorter proof. The rank of an abelian subgroup of $\mathrm{Out}(F_n)$ is $\leq 2n-3$ for $n > 1$ [CV86].

We reformulate Theorem 1.1 in terms of trees, and it is in this form that we prove the theorem. There is a natural right action of the automorphism group of F_n on the set of simplicial F_n-trees produced by twisting the action. See Section 2 for details. If we identify trees that are equivariantly isomorphic then this action descends to give an action of $\mathrm{Out}(F_n)$. A simplicial F_n-tree is nontrivial if there is no global fixed point. If T is a simplicial real F_n-tree with trivial edge stabilizers, then the set of conjugacy classes of nontrivial vertex stabilizers of T is a free factor system denoted $\mathcal{F}(T)$. The reformulation is as follows.

THEOREM 5.1. *For every finitely generated unipotent subgroup* $\mathcal{H}$ *of* $\mathrm{Out}(F_n)$ *there is a nontrivial simplicial* F_n*-tree* T *with all edge stabilizers trivial that is fixed by all elements of* $\mathcal{H}$. *Furthermore, there is such a tree with exactly one orbit of edges and if* $\mathcal{F}$ *is any maximal proper* $\mathcal{H}$*-invariant free factor system then* T *may be chosen so that* $\mathcal{F}(T) = \mathcal{F}$.

Such a tree can be obtained from the marked filtered graph produced by Theorem 1.1 by taking the universal cover and then collapsing all edges except for the lifts of the highest edge E_K. For a proof of the reverse implication, namely that Theorem 5.1 implies Theorem 1.1, see Section 6.

Along the way we obtain a result that is of interest in its own right. The necessary background material on trees may be found in Section 2, but also we give a quick review here. Simplicial F_n-trees may be endowed with metrics by equivariantly assigning lengths to edges. Given a simplicial real F_n-tree T and an element $a \in F_n$, the number $\ell_T(a)$ is defined to be the infimum of the distances that a translates elements of T. It is through these length functions that the space of simplicial real F_n-trees is topologized. Again there is a natural right action of $\mathrm{Out}(F_n)$. We will work in the $\mathrm{Out}(F_n)$-subspace $\mathcal{T}$ consisting of those nontrivial simplicial real trees that are limits of free actions.

THEOREM 1.4 *Suppose $T \in \mathcal{T}$ and $\mathcal{O} \in \mathrm{UPG}(F_n)$. There is an integer $d = d(\mathcal{O}, T) \geq 0$ such that the sequence $\{T\mathcal{O}^k/k^d\}$ converges to a tree $T\mathcal{O}^\infty \in \mathcal{T}$.*

This is proved in Section 4 as Theorem 4.22, which also contains an explicit description of the limit tree in the case that $d(\mathcal{O}, T) \geq 1$.

Section 5 is the heart of the proof of Theorem 5.1. For notational simplicity, let us assume that $\mathcal{H}$ is generated by two elements, $\mathcal{O}_1$ and $\mathcal{O}_2$. Given $T \in \mathcal{T}$, let Elliptic(T) be the subset of F_n consisting of elements fixing a point of T. Elements of Elliptic(T) are *elliptic*. Choose $T_0 \in \mathcal{T}$ such that T_0 has trivial edge stabilizers and such that Elliptic(T_0) is $\mathcal{H}$-invariant and maximal, i.e. such that if $T \in \mathcal{T}$ has trivial edge stabilizers, if Elliptic(T) is $\mathcal{H}$-invariant, and if Elliptic(T_0) $\subset$ Elliptic(T), then Elliptic(T_0) = Elliptic(T).

We prove that T_0 satisfies the conclusions of Theorem 5.1 but not by a direct analysis of T_0. Rather, we consider the "*bouncing sequence*" $\{T_0, T_1, T_2, \cdots\}$ in $\mathcal{T}$ defined inductively by $T_{i+1} = T_i\mathcal{O}_{i+1}^\infty$ where the subscripts of the outer automorphisms are taken mod 2. We establish properties of T_i for large i and then use these to prove that T_0 is the desired tree.

The key arguments in Section 5 are Proposition 5.5, Proposition 5.7, and Proposition 5.13. They focus not on discovering "ping-pong" dynamics ($\mathcal{H}$ may well contain a nonabelian free group), but rather on constructing an element in $\mathcal{H}$ of exponential growth. The connection to the bouncing sequence is as follows. Properties of the tree $T_k = T_0\mathcal{O}_1^\infty\mathcal{O}_2^\infty \ldots \mathcal{O}_{k-1}^\infty$ are reflected in the dynamics of the 'approximating' outer automorphism $\mathcal{O}(k) = \mathcal{O}_1^{N_1}\mathcal{O}_2^{N_2} \ldots \mathcal{O}_{k-1}^{N_{k-1}}$ where $N_1 \gg N_2 \gg \cdots \gg N_{k-1} \gg 1$. We verify properties of T_k by proving that if the property did not hold, then $\mathcal{O}(k)$ would have exponential growth.

After the breakthrough of E. Rips and the subsequent successful applications of the theory by Z. Sela and others, it became clear that trees were the right tool for proving Theorem 1.1. Surprisingly, under the assumption that $\mathcal{H}$ is finitely generated (which is the case that we are concerned with in this paper and which suffices for proving that the Tits alternative holds), we only work with simplicial real trees and the full scale $\mathbb{R}$-tree theory is never used. However, its existence gave us a firm belief that the project would succeed, and, indeed, the first proof we found of the Tits alternative used this theory. In a sense, our proof can be viewed as a development of the program, started by Culler-Vogtmann [CV86], to use spaces of trees to understand $\mathrm{Out}(F_n)$ in much the same way that Teichmüller space and its compactification were used by Thurston and others to understand mapping class groups.

2. F_n-trees

In this section, we collect the facts about real F_n-trees that we will need. This paper will only use these facts for simplicial real trees, but we sometimes record more general results for anticipated later use. Much of the material in this section can be found in [Ser80], [SW79], [CM87], or [AB87].

2.1. *Real trees.* An arc in a topological space is a subspace homeomorphic to a compact interval in $\mathbb{R}$. A point is a *degenerate* arc. A *real tree* is a metric space with the property that any two points may be joined by a unique arc, and further, this arc is isometric to an interval in $\mathbb{R}$ (see for example [AB87] or [CM87]). The arc joining points x and y in a real tree is denoted by $[x, y]$. A *branch point* of a real tree T is a point $x \in T$ whose complement has other than 2 components. A real tree is *simplicial* if it is equipped with a discrete subspace (the set of *vertices*) containing all branch points such that the *edges* (closures of the components of the complement of the set of vertices) are compact. If the subspace of branch points of a real tree T is discrete, then it admits a (nonunique) structure as a simplicial real tree. The simplicial real trees appearing in this paper will come with natural maps to compact graphs and the vertex sets of the trees will be the preimages of the vertex sets of the graphs.

For a real tree T, a map $\sigma : J \to T$ with domain an interval J is a *path in* T if it is an embedding or if J is compact and the image is a single point; in the latter case we say that σ is a *trivial* path.

If the domain J of a path σ is compact, define the *inverse of* σ, denoted $\overline{\sigma}$ or σ^{-1}, to be $\sigma \circ \rho$ where $\rho : J \to J$ is a reflection.

We will not distinguish paths in T that differ only by an orientation-preserving change of parametrization. Hence, every map $\sigma : J \to T$ with J compact is properly homotopic rel endpoints to a unique path $[\sigma]$ called its *tightening.*

If $\sigma : J \to T$ is a map from the compact interval J to the simplicial real tree T and the endpoints of J are mapped to vertices, then the image of $[\sigma]$, if nondegenerate, has a natural decomposition as a concatenation $E_1 \cdots E_k$ where each E_i, $1 \le i \le k$, is a directed edge of T. The sequence $E_1 \cdots E_k$ is called *the edge path associated to* σ. We will identify $[\sigma]$ with its associated edge path. This notation extends naturally if the domain of the path is a ray or the entire line and σ is an embedding. A path *crosses* an edge of T if the edge appears in the associated edge path. A path is *contained in* a subtree if it crosses only edges of the subtree. A *ray* in T is a path $[0, \infty) \to T$ that is an embedding.

2.2. *Real F_n-trees.* By F_n denote a fixed copy of the free group with basis $\{e_1, \ldots, e_n\}$. A real F_n-tree is a real tree equipped with an action of F_n by isometries. It is *minimal* if it has no proper F_n-invariant subtrees. If H is a subgroup of F_n then $\mathrm{Fix}_T(H)$ denotes the subset of T consisting of points that are fixed by each element of H. If $a \in F_n$, then $\mathrm{Fix}_T(a) := \mathrm{Fix}_T(\langle a \rangle)$. If $X \subset T$, then $\mathrm{Stab}_T(X)$ is the subgroup of F_n consisting of elements that leave X invariant. If $x \in T$, then $\mathrm{Stab}_T(x) := \mathrm{Stab}_T(\{x\})$. The symbol '$[[\cdot]]$' denotes 'conjugacy class'. Define

$$\mathrm{Point}(T) := \{[[\mathrm{Stab}_T(x)]] \mid x \in T,\ \mathrm{Stab}_T(x) \neq \langle 1 \rangle\}$$

and

$$\mathrm{Arc}(T) := \{[[\mathrm{Stab}_T(\sigma)]] \\ \mid \sigma \text{ is a nondegenerate arc in } T,\ \mathrm{Stab}_T(\sigma) \neq \langle 1 \rangle\}.$$

The *length function* of a real F_n-tree T assigns to $a \in F_n$ the number

$$\ell_T(a) := inf_{x \in T}\{d_T(x, ax)\}.$$

Length is constant on conjugacy classes, so we also write $\ell_T([[a]])$ for $\ell_T(a)$. If $\ell_T(a)$ is positive, then a (or $[[a]]$) is *hyperbolic in* T, otherwise a is *elliptic*. If a is hyperbolic in T, then $\{x \in T \mid d_T(x, ax) = \ell_T(a)\}$ is isometric to $\mathbb{R}$. This set is called the *axis of* a and is denoted $\mathrm{Axis}_T(a)$. The restriction of a to its axis is translation by $\ell_T(a)$. If a is elliptic in T then a fixes a point of T. Thus, an element of F_n is in $\mathrm{Elliptic}(T)$ if it is trivial or if its conjugacy class is in $\mathrm{Point}(T)$. A subgroup of F_n is *elliptic* if all elements are elliptic.

A real F_n-tree T is *trivial* if $\mathrm{Fix}_T(F_n) \neq \emptyset$. In particular, a minimal tree is trivial if and only if it is a point. We will need the following special case of a result of Serre.

Theorem 2.1 ([Ser80]). *Suppose that T is a real F_n-tree where $F_n = \langle a_1, \ldots, a_k \rangle$. Suppose that $a_i a_j$ is elliptic in T for $1 \leq i, j \leq k$. Then T is trivial.*

2.3. *Very small trees.* We will only need to consider a restricted class of real trees.

A real F_n-tree T is *very small* [CL95] if

(1) T is nontrivial,

(2) T is minimal.

(3) The subgroup of F_n of elements pointwise fixing a nondegenerate arc of T is either trivial or maximal cyclic, and

(4) for each $1 \neq a \in F_n$, $\mathrm{Fix}_T(a)$ is either empty or an arc.

It follows from (3) that if T is very small and $x, y \in T$, then each element of $\mathrm{Stab}_T([x, y])$ fixes $[x, y]$ pointwise. In particular, if T is simplicial, then no element of F_n inverts an edge.

We will need:

THEOREM 2.2 ([CM87], [AB87]). *Let Q be a finitely generated group, and let T be a minimal nontrivial Q-tree. Then the axes of hyperbolic elements of Q cover T.*

In the case of simplicial trees, the following theorem is established by an easy Euler characteristic argument. The generalization to $\mathbb{R}$-trees due to Gaboriau and Levitt uses more sophisticated techniques.

THEOREM 2.3 ([GL95]). *Let T be a very small F_n-tree. There is a bound depending only on n to the number of conjugacy classes of point and arc stabilizers. The rank of a point stabilizer is no more than n with equality if and only if T/F_n is a wedge of circles and each edge of T has infinite cyclic stabilizer.*

2.4. *Spaces of real F_n-trees.* Let $\mathbb{R}_+$ denote the ray $[0, \infty)$ and let $\mathcal{C}$ denote the set of conjugacy classes of elements in F_n. The space $\mathcal{T}_{all}$ of non-trivial minimal real F_n-trees is given the smallest topology such that the map $\theta : \mathcal{T}_{all} \to \mathbb{R}_+^{\mathcal{C}}$, given by $\theta(T) = (\ell_T(a))_{[[a]]\in\mathcal{C}}$ is continuous.

Let $\mathcal{T}_{CV}$ denote the subspace of $\mathcal{T}_{all}$ consisting of free simplicial actions. The closure of $\mathcal{T}_{CV}$ in $\mathcal{T}_{all}$ is denoted $\mathcal{T}_{VS}$. The subspace of simplicial trees in $\mathcal{T}_{VS}$ is denoted $\mathcal{T}$. The map θ is injective when restricted to $\mathcal{T}_{VS}$; see [CM87]. In other words, if $S, T \in \mathcal{T}_{VS}$ satisfy $\theta(S) = \theta(T)$, then S and T are equivariantly isometric. In this paper, we only need to work in $\mathcal{T}$ although some results are presented in greater generality.

The automorphism group $\mathrm{Aut}(F_n)$ acts naturally on $\mathcal{T}_{all}$ on the right by twisting the action; i.e., if the action on $T \in \mathcal{T}_{all}$ is given by $(a, t) \mapsto a \cdot t$ and if $\Phi \in \mathrm{Aut}(F_n)$ then the action on $T\Phi$ is given by $(a, t) \mapsto \Phi(a) \cdot t$. In terms of the length functions, the action is given by $\ell_{T\Phi}(a) = \ell_T(\Phi(a))$ for $\Phi \in \mathrm{Aut}(F_n)$, $T \in \mathcal{T}_{all}$, and $a \in F_n$. The subgroup $\mathrm{Inner}(F_n)$ of inner automorphisms acts trivially, and we have an action of $\mathrm{Out}(F_n) = \mathrm{Aut}(F_n)/\mathrm{Inner}(F_n)$. The spaces $\mathcal{T}_{CV}$, $\mathcal{T}_{VS}$, and $\mathcal{T}$ are all $\mathrm{Out}(F_n)$-invariant.

To summarize, for $\mathcal{O} \in \mathrm{Out}(F_n)$ and $T \in \mathcal{T}_{VS}$, the following are equivalent.

- $\mathcal{O}$ fixes T.
- $\ell_T(\mathcal{O}([[\gamma]])) = \ell_T([[\gamma]])$ for all $\gamma \in F_n$.
- For any $\Phi \in \mathrm{Aut}(F_n)$ representing $\mathcal{O}$, there is a Φ-equivariant isometry $f_\Phi : T \to T$.

2.5. *Bounded cancellation constants.* We will often need to compare the length of the same element of F_n in different real F_n-trees. This is facilitated by the existence of bounded cancellation constants.

Definition 2.4. Let S and T be real F_n-trees. The *bounded cancellation constant* of an F_n-map $f : S \to T$, denoted $\mathrm{BCC}(f)$, is the least upper bound of numbers B with the property that there exist points $x, y, z \in S$ with $y \in [x, z]$ so that the distance between $f(y)$ and $[f(x), f(z)]$ is B.

Cooper [Coo87] showed that if S and T are in $\mathcal{T}_{CV}$ and if f is PL, then $\mathrm{BCC}(f)$ is finite. For a map $f : X \to Y$ between metric spaces we denote by $\mathrm{Lip}(f)$ the Lipschitz constant of f; i.e.,

$$\mathrm{Lip}(f) := \sup\{d_Y(f(x_1), f(x_2))/d_X(x_1, x_2) \mid (x_1, x_2) \in X \times X, x_1 \neq x_2\}.$$

The map f is *Lipschitz* if $\mathrm{Lip}(f) < \infty$. The following generalization of Cooper's result is an immediate consequence of Lemma 3.1 of [BFH97a].

PROPOSITION 2.5. *Suppose that* $S \in \mathcal{T}_{CV}$, $T \in \mathcal{T}_{VS}$, *and* $f : S \to T$ *is a Lipschitz* F_n*-map. Then,* $\mathrm{BCC}(f) < \infty$.

2.6. *Real graphs.* In [BFH00], marked graphs were used. Here we will need graphs with a metric structure.

A *real graph* is a locally finite graph (one-dimensional CW-complex) whose universal cover has the structure of a simplicial real tree with covering transformations acting by isometries. A locally finite graph with specified edge lengths determines a real graph. Occasionally, it is convenient to view a locally finite graph as a real graph. To do this, we will specify edge lengths. If no lengths are mentioned, then they are assumed to be 1.

Let G be a real graph with universal covering $p : \Gamma \to G$. A map $\sigma : J \to G$ with domain an interval J is a *path* if $\sigma = p \circ \tilde{\sigma}$ where $\tilde{\sigma}$ is a path in Γ. The terminology for paths in trees transfers directly over to real graphs; cf. [BFH00, p. 525].

A *closed path* in G is a path whose initial and terminal endpoints coincide. A *circuit* is an immersion from the circle S^1 to G; homotopic circuits are not distinguished. Any homotopically nontrivial map $\sigma : S^1 \to G$ is homotopic to a unique circuit $[[\sigma]]$. Circuits are identified with cyclically ordered edge paths which we call *associated edge circuits.* A circuit *crosses* an edge if the edge appears in the circuit's associated edge circuit. A circuit is *contained in* a subgraph if it crosses only edges of the subgraph. We make standard identifications between based closed paths and elements of the fundamental group and between circuits and conjugacy classes in the fundamental group.

A *marked real graph* is a real graph G together with a homotopy equivalence $\mu : \mathrm{Rose}_n \to G$. The universal cover of a marked real graph has a structure of a real free F_n-tree that is well-defined up to equivariant isometry. A real F_n-tree T admits an F_n-equivariant map $\tilde{\mu} : \widetilde{\mathrm{Rose}_n} \to T$. This map is well-defined up to equivariant homotopy. If the action is free, then the quotient $\mu : \mathrm{Rose}_n \to G$ is a marking.

A *core graph* is a finite graph with no vertices of valence 1 or 0. Any connected graph with finitely generated fundamental group has a unique maximal core subgraph, called its *core*. The core of a forest is empty.

2.7. *Models and normal forms for simplicial F_n-trees.* References for this section are [SW79] and [Ser80]. A map $h : Y \to Z$ with Y and Z CW-complexes is *cellular* if, for all k, the k-skeleton of Y maps into the k-skeleton of Z, i.e. $h(Y^{(k)}) \subset Z^{(k)}$. Given CW-complexes Y, Z_0, and Z_1 and cellular maps $g_i : Y \to Z_i$, the *double mapping cylinder* $D(g, h)$ of g and h is the quotient $(Y \times [0, 1]) \sqcup (Z_0 \sqcup Z_1)/\sim$ where $\sim$ is the equivalence relation generated by $(y, 0) \sim g_0(y)$ and $(y, 1) \sim g_1(y)$. The double mapping cylinder is naturally a CW-complex with a map to $[0, 1]$. In the case where $Z_0 = Z_1$, we modify the definition of $D(g, h)$ so that corresponding points of Z_0 and Z_1 are also identified. In this case, $D(g, h)$ has a natural map to S^1.

Let Rose_n denote a fixed wedge of n oriented circles with a fixed identification of $\pi_1(\mathrm{Rose}_n, *)$ with F_n such that the i^{th} circle corresponds to e_i. Also fix a compatible identification of F_n with the covering transformations of the universal cover $\widetilde{\mathrm{Rose}_n}$ of Rose_n.

Let T be a simplicial real F_n-tree and let $\overline{T}$ denote the real graph T/F_n. A *graph of spaces over* $\overline{T}$ is a CW-complex X with a cellular map $q : X \to \overline{T}$ such that:

- For each vertex x of $\overline{T}$, $q^{-1}(x)$ is a subcomplex of X.
- For each edge e of $\overline{T}$ with endpoints v and w (possibly equal), there is a CW-complex X_e, a pair of cellular maps $g : X_e \to q^{-1}(v)$ and $h : X_e \to q^{-1}(w)$, and isomorphisms $D(g, h) \to q^{-1}(e)$ and S^1 or $[0, 1] \to e$ such that the following diagram commutes.

$$\begin{array}{ccc} D(g,h) & \longrightarrow & q^{-1}(e) \\ \downarrow & & \downarrow \\ S^1 \text{or } [0,1] & \longrightarrow & e \end{array}$$

A *vertical subspace of* X is a subcomplex of the form $q^{-1}(x)$ for some vertex $x \in \overline{T}$. An edge of a vertical subspace is *vertical*. Other edges of X are *horizontal*. An edge path consisting of vertical edges is *vertical*.

Example 2.6. The quotient $\widetilde{\mathrm{Rose}_n} \times_{F_n} T$ of $\widetilde{\mathrm{Rose}_n} \times T$ by the diagonal action of F_n with the map $Q : \widetilde{\mathrm{Rose}_n} \times_{F_n} T \to \overline{T}$ induced by projection onto the second coordinate is naturally a graph of spaces over $\overline{T}$. It is an Eilenberg-MacLane space. Its fundamental group is naturally identified with F_n by the map to Rose_n induced by projection onto the first coordinate. If x is a vertex of $\overline{T}$, then $Q^{-1}(x)$ is a full subcomplex of $\widetilde{\mathrm{Rose}_n} \times_{F_n} T$ isomorphic to $\widetilde{\mathrm{Rose}}_n/\mathrm{Stab}_T(\tilde{x})$ where $\tilde{x}$ is a lift of x to T. In particular, $Q^{-1}(x)$ is a graph homotopy equivalent to the wedge Rose_x of n_x circles where n_x is the rank of $\mathrm{Stab}_T(\tilde{x})$. Similarly, if x is a point in the interior of an edge e of $\overline{T}$, then the preimage of x is isomorphic to $\widetilde{\mathrm{Rose}}_n/\mathrm{Stab}_T(\tilde{e})$ where $\tilde{e}$ is a lift of e to T. In particular, $Q^{-1}(x)$ is a graph homotopy equivalent to the wedge Rose_e of n_e circles where n_e is the rank of $\mathrm{Stab}_T(\tilde{e})$.

Let $\mathcal{M}$ be the set of midpoints of edges of $\overline{T}$. A *model for* T is a graph of spaces X over $\overline{T}$ with a homotopy equivalence $\widetilde{\mathrm{Rose}}_n \times_{F_n} T \to X$ such that the following diagram commutes up to a homotopy supported over the complement of $\mathcal{M}$:

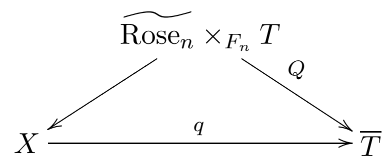

and such that the induced map $Q^{-1}(\mathcal{M}) \to q^{-1}(\mathcal{M})$ is a homotopy equivalence. The homotopy equivalence $X \leftarrow \widetilde{\mathrm{Rose}}_n \times_{F_n} T \to \mathrm{Rose}_n$ identifies conjugacy classes in $\pi_1(X)$ with conjugacy classes in F_n and is called *the induced marking.*

The trees in this paper will all be minimal with finitely generated vertex and edge stabilizers. (In fact, edge stabilizers will be cyclic.) Until Section 5.4, we will make the following additional requirements of our models.

- If x is a vertex of $\overline{T}$, then the vertical subspace $q^{-1}(x)$ is a subcomplex of X isomorphic to Rose_x.
- If e is an edge of $\overline{T}$, then X_e is isomorphic to Rose_e.

Models satisfying these properties are constructed in [SW79].

Example 2.7. Pictured below is an example of a model X together with the induced marking $\mu : \mathrm{Rose}_2 \to X$ and the quotient map $q : X \to \overline{T}$. The μ-image of the edge 'e_1' is the edge 'a' and the μ-image of the edge 'e_2' is the only horizontal edge 't'. The vertical space is a wedge of two circles. This model satisfies all of the above properties.

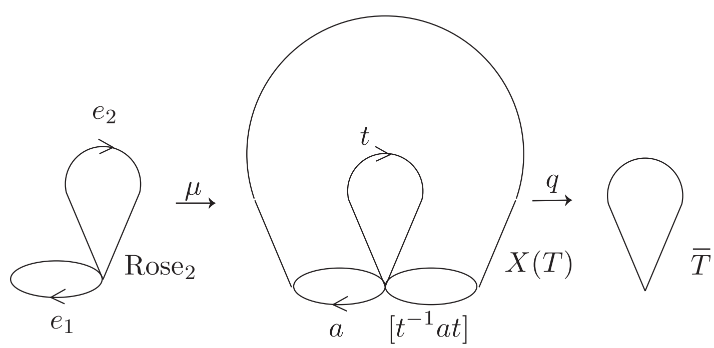

Any path in X whose endpoints are vertices is homotopic rel endpoints to an edge path in $X^{(1)}$ of the form

$$\nu_0 H_1 \nu_1 H_2 \nu_2 \cdots H_m \nu_m$$

where ν_i is a (possibly trivial) vertical edge path and H_i is a horizontal edge of X. The *length* of the path is the sum of the lengths $q(H_i)$. Such an edge path is *in normal form* unless for some i we have that $H_i \nu_i H_{i+1}$ is homotopic rel endpoints into a vertical subspace. If σ is a path in X whose endpoints are vertices, then $[\sigma]$ is an edge path homotopic rel endpoints to σ that is in normal form.

If the displayed edge path is not in normal form and if i is as above, then the path is homotopic to the path obtained by replacing $\nu_{i-1} H_i \nu_i H_{i+1} \nu_{i+1}$ by a path in a vertical subspace that is homotopic rel endpoints. We call this process *erasing a pair of horizontal edges*. Any edge path in X may be put into normal form by iteratively erasing pairs horizontal edges (see [SW79]). Two paths in normal form that are homotopic rel endpoints have the same length.

In an analogous fashion, circuits in X have lengths and normal forms. If σ is a circuit in X, then $[[\sigma]]$ is a circuit freely homotopic to σ that is in normal form. Note that $length_X([[\sigma]]) = \ell_T([[a]])$ where $[[a]]$ is the conjugacy class of F_n represented by the image of σ under the induced marking of X.

If σ_1 and σ_2 are paths in X with the same initial points, then the *overlap length of* σ_1 *and* σ_2 is defined to be

$$\frac{1}{2} \cdot \big(length_X([\sigma_1]) + length_X([\sigma_2]) - length_X([\overline{\sigma}_1 \sigma_2]) \big).$$

Remark 2.8. Suppose that σ_1, σ_2 and σ_3 are paths in X with endpoints at vertices, that the terminal endpoint of σ_1 is the initial endpoint of σ_2 and that the terminal endpoint of σ_2 is the initial endpoint of σ_3. Let D be the overlap length of $\bar{\sigma}_1$ and σ_2, let D' be the overlap length of $\bar{\sigma}_2$ and σ_3 and assume that $length_X([\sigma_2]) > D + D'$. In the proof of Proposition 4.21 we use the fact, immediate from the definitions, that the following quantities are realized as lengths of edge paths in T.

- $length_X([\sigma_2])$.
- $length_X([\sigma_2]) - (D + D')$.

Example 2.9. Let X be as in Example 2.7. Then the overlap length of t and at is the same as the length in X of t even though the maximal common initial segment of t and at is degenerate. Of course, $at = t[t^{-1}at]$ are both normal forms and the maximal common initial segment of t and $t[t^{-1}at]$ is t.

2.8. *Free factor systems.* Here we review definitions and background for free factor systems as treated in [BFH00].

We reserve the notation F^i for free factors of F_n. If $F^1 * F^2 * \cdots * F^k$ is a free factor and each F^i is nontrivial (and so has positive rank), then we say that the collection $\mathcal{F} = \{[[F^1]], [[F^2]], \dots, [[F^k]]\}$ is a *nontrivial free factor system.* We refer to $\emptyset$ as *the trivial free factor system.* A free factor system $\mathcal{F}$ is *proper* if it is not $\{[[F_n]]\}$.

We write $[[F^1]] \sqsubset [[F^2]]$ if F^1 is conjugate to a free factor of F^2 and write $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ if for each $[[F^i]] \in \mathcal{F}_1$ there exists $[[F^j]] \in \mathcal{F}_2$ such that $[[F^i]] \sqsubset [[F^j]]$. We say that $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ is *proper* if $\mathcal{F}_1 \neq \mathcal{F}_2$. The next lemma follows immediately from Lemma 2.6.3 of [BFH00].

Lemma 2.10. *There is a bound, depending only on n, to the length of a chain $\mathcal{F}_1 \sqsubset \mathcal{F}_2 \sqsubset \cdots \sqsubset \mathcal{F}_N$ of proper $\sqsubset$'s.*

We say that a subset X of F_n is *carried by* the free factor system $\mathcal{F}$ if $X \subset F^i$ for some $[[F^i]] \in \mathcal{F}$. A collection $\mathcal{X}$ of subsets is *carried by* $\mathcal{F}$ if each $X \in \mathcal{X}$ is carried by some element of $\mathcal{F}$.

Let ∂F_n denote the boundary of F_n. Let $\mathcal{R}_n$ (for rays) denote the quotient of ∂F_n by the action of F_n. The natural action of $\mathrm{Aut}(F_n)$ on ∂F_n descends to an action of $\mathrm{Out}(F_n)$ on $\mathcal{R}_n$. If G is a marked real graph, then $\mathcal{R}_n$ is naturally identified with the set of rays in G where two rays are equivalent if their associated edge paths have a common tail. In [BFH00], a parallel treatment was given using lines instead of rays. The reader is referred there for details.

A free factor F^i of F_n gives rise to a subset $\mathcal{R}^i$ of $\mathcal{R}_n$. In terms of a tree $T \in \mathcal{T}_{CV}$, a ray represents an element of $\mathcal{R}^i$ if it can be F_n-translated so that its image is eventually in the minimal F^i-subtree of T. A ray $R \in \mathcal{R}_n$ is *carried by* F^i if $R \in \mathcal{R}^i$. It is *carried by* the free factor system $\mathcal{F}$ if it is carried by F^i for some $[[F^i]] \in \mathcal{F}$. A subset of $\mathcal{R}_n$ is *carried by* $\mathcal{F}$ if each element of the subset is carried by some element of $\mathcal{F}$.

The proof of the following lemma is completely analogous to the proof of Corollary 2.6.5 of [BFH00].

Lemma 2.11. *Let $\mathcal{X}$ be a collection of subsets of F_n and let $\mathcal{R}$ be a subset of $\mathcal{R}_n$. Then there is a unique minimal (with respect to $\sqsubset$) free factor system $\mathcal{F}$ that carries both $\mathcal{X}$ and $\mathcal{R}$.*

A (not necessarily connected) subgraph K of a marked real graph determines a free factor system $\mathcal{F}(K)$ as in [BFH00, Ex. 2.6.2]. If T is a simplicial real F_n-tree with trivial edge stabilizers, then the set of conjugacy classes of nontrivial vertex stabilizers of T is a free factor system denoted $\mathcal{F}(T)$.

3. Unipotent polynomially growing outer automorphisms

In this section we bring outer automorphisms into the picture. We will consider a class of outer automorphisms that is analogous to the class of unipotent matrices. First we review the linear algebra of unipotent matrices.

3.1. *Unipotent linear maps.* The results in this section are standard. We include proofs for the reader's convenience. Throughout this section, R denotes either $\mathbb{Z}$ or $\mathbb{C}$, and V denotes a free R-module of finite rank.

Proposition 3.1. *Let $f : V \to V$ be an R-module endomorphism. The following conditions are equivalent*:

(1) *V has a basis with respect to which f is upper triangular with* 1*'s on the diagonal.*

(2) $(\mathrm{Id} - f)^{\mathrm{rank}(V)} = 0$.

(3) $(\mathrm{Id} - f)^n = 0$ *for some* $n > 0$.

Proof. It is clear that (1) implies (2) and that (2) implies (3). To see that (3) implies (1), assume that $(\mathrm{Id} - f)^n = 0$. We may assume that $W := \mathrm{Im}(\mathrm{Id} - f)^{n-1} \neq 0$. The restriction of $\mathrm{Id} - f$ to the submodule W is 0, and hence each $0 \neq v \in W$ is fixed by f. After perhaps replacing v by a root in the case $R = \mathbb{Z}$, we may assume that v is an f-fixed basis element of V. The proof now concludes by induction on $\mathrm{rank}(V)$ using the fact that the induced homomorphism $f' : V/\langle v\rangle \to V/\langle v\rangle$ also satisfies $(\mathrm{Id} - f')^n = 0$. □

An endomorphism f satisfying any of the equivalent conditions of Proposition 3.1 is said to be *unipotent.*

Corollary 3.2. *Let $f : V \to V$ be an R-module endomorphism, and let W be an f-invariant submodule of V which is a direct summand of V. Then f is unipotent if and only if both the restriction of f to W and the induced endomorphism on V/W are unipotent.*

Proof. The proof is evident if we use Proposition 3.1(1) in the "if" direction and Proposition 3.1(2) in the "only if" direction. □

Corollary 3.3. *Let $f : V \to V$ be unipotent. If $x \in V$ is f-periodic, i.e. if $f^m(x) = x$ for some $m > 0$, then x is f-fixed, i.e. $f(x) = x$.*

Proof. First assume that $R = \mathbb{C}$. We may assume that

$$V = \mathrm{span}(x, f(x), \cdots, f^{m-1}(x)).$$

Let $e_1, e_2, \ldots, e_m$ be the standard basis for $\mathbb{C}^m$. There is a surjective linear map $\pi : \mathbb{C}^m \to V$ given by $\pi(e_i) = f^{i-1}(x)$, and f lifts to the linear map

$\overline{f} : \mathbb{C}^m \to \mathbb{C}^m$, $\overline{f}(e_i) = e_{i+1 \mod m}$. For $\lambda \in \mathbb{C}$, the generalized λ-eigenspace is defined to be

$$\{x \in \mathbb{C}^m | (\lambda I - \overline{f})^m(x) = 0\}.$$

Since f is unipotent, the linear map π must map the generalized 1-eigenspace onto V (and all other generalized eigenspaces to 0). The characteristic polynomial $\lambda^m - 1$ of $\overline{f}$ has m distinct roots. In particular, the generalized 1-eigenspace is one-dimensional (and equals the 1-eigenspace of $\overline{f}$). It follows that $\dim(V) \leq 1$ and $f(x) = x$.

If $R = \mathbb{Z}$, just tensor with $\mathbb{C}$. □

Corollary 3.4. *Let $f : V \to V$ be unipotent. If W is a direct summand which is periodic (i.e. $f^m(W) = W$ for some $m > 0$), then W is invariant (i.e. $f(W) = W$).*

Proof. The restriction of f^m to W is unipotent, so there is a basis element $x \in W$ fixed by f^m. By Corollary 3.3, $f(x) = x$. The proof concludes by induction on $\mathrm{rank}(W)$. □

Proposition 3.5. *Let $A \in \mathrm{GL}_n(\mathbb{Z})$ have all eigenvalues on the unit circle (i.e. A grows polynomially). If the image of A in $\mathrm{GL}_n(\mathbb{Z}/3\mathbb{Z})$ is trivial, then A is unipotent.*

Proof. We first argue that some power A^N of A is unipotent, i.e. that all eigenvalues of A are roots of unity. Choose N so that all eigenvalues of A^N are close to 1. Then $\mathrm{tr}(A^N)$ is an integer close to n, and thus all eigenvalues of A^N are equal to 1.

Let $f = f_1^{n_1} \cdots f_m^{n_m}$ be the minimal polynomial for A factored into irreducibles in $\mathbb{Z}[x]$. Let $A_i = f_i^{n_i}(A)$ and $K_i = \mathrm{Ker}(A_i)$. First note that each $K_i \neq 0$. For example, $\mathrm{Im}(A_2 A_3 \cdots A_m) \subset K_1$ but $A_2 A_3 \cdots A_m \neq 0$ since f is minimal. If A is not unipotent, then some f_i, say f_1, is not $x - 1$. Since all roots of f are roots of unity, f_1 is the minimal polynomial for a nontrivial root of unity and so it divides $1 + x + x^2 + \cdots + x^{r-1}$ for some $r > 1$. The matrix $I + A + A^2 + \cdots + A^{r-1}$ has nontrivial kernel (since its n_1^{st} power vanishes on K_1). A nonzero integral vector v in this kernel satisfies $A^r(v) = v$ and $A(v) \neq v$. Then $\mathrm{Fix}(A^r)$ is a nontrivial direct summand of $\mathbb{Z}^n$, the restriction of A to this summand is nontrivial and periodic, and the induced endomorphism of $\mathrm{Fix}(A^r) \otimes \mathbb{Z}/3\mathbb{Z}$ is the identity. This contradicts the standard fact that the kernel of $\mathrm{GL}_k(\mathbb{Z}) \to \mathrm{GL}_k(\mathbb{Z}/3\mathbb{Z})$ is torsion-free. □

3.2. *Topological representatives.* A homotopy equivalence $f : G \to G$ of a marked real graph induces an outer automorphism $\mathcal{O}$ of F_n via the fixed identification of F_n with the fundamental group of Rose_n. If f maps vertices to vertices and if the restriction of f to each edge of G is an immersion, then we say that f is a *topological representative of* $\mathcal{O}$.

A *filtration* for a topological representative $f : G \to G$ is an increasing sequence of f-invariant subgraphs $\emptyset = G_0 \subsetneq G_1 \subsetneq \cdots \subsetneq G_K = G$. The closure of $G_r \setminus G_{r-1}$ is called the r^{th} *stratum*.

If the path $\sigma = \sigma_1\sigma_2$ is the concatenation of paths σ_1 and σ_2, then σ *splits*, denoted $\sigma = \sigma_1 \cdot \sigma_2$, if $[f^i(\sigma)] = [f^i(\sigma_1)][f^i(\sigma_2)]$ for all integers $i \geq 0$; see [BFH00, pp. 553–554]. In this paper, as in [BFH00], it is critically important to understand the behavior of paths under iteration by f. If a path splits, the behavior of the path is determined by the behavior of the subpaths.

3.3. *Relative train tracks and automorphisms of polynomial growth.* The techniques of this paper depend on being able to find good representatives for outer automorphisms of polynomial growth.

Definition 3.6. An outer automorphism $\mathcal{O} \in \mathrm{Out}(F_n)$ has *polynomial growth* if, given $a \in F_n$, there is a polynomial $P \in \mathbb{R}[x]$ such that the (reduced) word length of $\mathcal{O}^i([[a]])$ is bounded above by $P(i)$. The set of outer automorphisms having polynomial growth is denoted $\mathrm{PG}(F_n)$ (or just PG).

It follows from [BH92] that the definition of polynomial growth given above agrees with the definition on page 564 of [BFH00]. We start by recalling the topological representatives for automorphisms having polynomial growth that were found in [BH92].

Theorem 3.7 ([BH92]). *An automorphism $\mathcal{O} \in \mathrm{PG}(F_n)$ has a topological representative $f : G \to G$ with a filtration $\emptyset = G_0 \subsetneq G_1 \subsetneq \cdots \subsetneq G_K = G$ such that*

(1) *for every edge $E \in \overline{G_i \setminus G_{i-1}}$, the edge path $f(E)$ crosses exactly one edge in $\overline{G_i \setminus G_{i-1}}$ and it crosses that edge exactly once.*

(2) *If $\mathcal{F}$ is an $\mathcal{O}$-invariant free factor system, it can be arranged that $\mathcal{F} = \mathcal{F}(G_r)$ for some r. If $\mathcal{O}$ is the identity on each conjugacy class in $\mathcal{F}$, it can be arranged that $f = Id$ on G_r.*

Definition 3.8. A topological representative as in Theorem 3.7 is called a *relative train track* (RTT) *representative* for $\mathcal{O}$.

3.4 *Unipotent representatives and* UPG *automorphisms.*

Definition 3.9. An outer automorphism is *unipotent* if it has polynomial growth and its action on $H_1(F_n;\mathbb{Z})$ is unipotent. The set of unipotent automorphisms is denoted by $\mathrm{UPG}(F_n)$ (or just UPG).

We now recall a special case of an improvement of RTT representatives from [BFH00].

Definition 3.10. Let $f : G \to G$ be an RTT representative. A nontrivial path τ in G is a *periodic Nielsen path* if, for some $m \geq 0$, $[f^m(\tau)] = [\tau]$. If $m = 1$ then τ is a *Nielsen path*. An *exceptional path* in G is a path of the form $E_i\tau^m\overline{E}_j$ where $\overline{G_i \setminus G_{i-1}}$ is the single edge E_i, $\overline{G_j \setminus G_{j-1}}$ is the single edge E_j, τ is a Nielsen path, $f(E_i) = E_i\tau^p$, and $f(E_j) = E_j\tau^q$ for some $m \in \mathbb{Z}$, $p, q > 0$.

Theorem 3.11 ([BFH00, Th. 5.1.8]). *Suppose that $\mathcal{O} \in$ UPG(F_n) and that $\mathcal{F}$ is an $\mathcal{O}$-invariant free factor system. Then there is an* RTT *representative $f : G \to G$ and a filtration $\emptyset = G_0 \subsetneq G_1 \subsetneq \cdots \subsetneq G_K = G$ representing $\mathcal{O}$ with the following properties*:

(1) *$\mathcal{F} = \mathcal{F}(G_r)$ for some filtration element G_r.*

(2) *Each $\overline{G_i \setminus G_{i-1}}$ is a single edge E_i satisfying $f(E_i) = E_i \cdot u_i$ for some closed path u_i with edges in G_{i-1}.*

(3) *Every vertex of G is fixed by f.*

(4) *Every periodic Nielsen path has period one.*

(5) *If σ is any path with endpoints at vertices, then there exists $M = M(\sigma)$ so that for each $m \geq M$, $[f^m(\sigma)]$ splits into subpaths that are either single edges or are exceptional.*

(6) *$M(\sigma)$ is a bounded multiple of the edge length of σ.*

Remark 3.12. Another useful condition is

(7) If E_i and E_j are distinct edges of G with nontrivial suffixes u_i and u_j, then $u_i \neq u_j$.

This property is part of the construction of $f : G \to G$ from [BFH00, Th. 5.1.8]. There is an operation called sliding that is used for nonexponentially growing strata. Condition 1 of [BFH00, Prop. 5.4.3] implies Item (7). Alternatively, starting with $f : G \to G$ satisfying (1–6), f may be enhanced to also satisfy (7) by replacing E_j with $E_j\bar{E}_i$.

Definition 3.13. An RTT representative f satisfying Items (1–7) above is a *unipotent representative* or a UR. The based closed paths u_i are *suffixes* of f.

Remark 3.14. Note that Item (2) can be restated as

$$[f^k(E_i)] = E_i \cdot u_i \cdot [f(u_i)] \cdot \cdots \cdot [f^{k-1}(u_i)]$$

for all $k > 0$. Since exceptional paths do not have nontrivial splittings, the splitting of $[f^k(E_i)]$ guaranteed by Item (5) restricts to a splitting of u_i into single edges and exceptional paths. The immersed infinite ray

$$R_i = E_i u_i [f(u_i)] \cdots [f^{k-1}(u_i)] \cdots$$

is the *eigenray* associated to E_i. Lifts of R_i to the universal cover of G are also called eigenrays. The subpaths $[f^m(u_i)]$ of R_i are sometimes referred to as *blocks*.

For example, the map $f : G \to G$ on the wedge of two circles with edges a and b given by $f(a) = a$, $f(b) = ba$ is a UR. For $\omega = ba^{-10}bab^{-1}$ we may take $M(\omega) = 10$ in Item (5), since $[f^{10}(\omega)] = b \cdot (bab^{-1})$ is a splitting into an edge and an exceptional (Nielsen) path. The map given by $a \mapsto a$, $b \mapsto ba$, $c \mapsto cba^{-1}$ on the wedge of three circles is not a UR since $\omega = cba^{-1}$ does not eventually split as in Item (5). Replacing ba^{-1} by b' yields a UR of the same outer automorphism.

Definition 3.15. Let $f : G \to G$ be a UR with filtration $\emptyset = G_0 \subsetneq G_1 \subsetneq \cdots \subsetneq G_K = G$. The *highest edge* (*or stratum*) *of* G is $E_K = G_K \setminus G_{K-1}$. The height of a path σ in G, denoted *height*(σ), is the smallest m such that the path crosses only edges in G_m. If σ is a path of height m, then a *highest edge in* σ is an occurrence of E_m or $\overline{E}_m$ in σ. By [BFH00, Lemma 4.1.4], the path σ naturally splits at the initial endpoints of its highest edges; this is called the *highest edge splitting of* σ.

Many arguments in this paper are inductions on height.

Proposition 3.16. *If* $\mathcal{O} \in \mathrm{UPG}(F_n)$, *then all* $\mathcal{O}$*-periodic conjugacy classes are fixed.*

Proof. Let $f : G \to G$ be a UR for $\mathcal{O}$ and let σ be a circuit in G representing an $\mathcal{O}$-periodic conjugacy class. Consider the highest edge splitting of σ. Each of the resulting subpaths is an f-periodic Nielsen path. Theorem 3.11(4) now implies that each subpath is f-fixed, and thus σ is f-fixed. □

We will also need the following more technical results.

Lemma 3.17 ([BFH00, Lemma 5.7.9]). *Suppose that* $f : G \to G$ *is a* UR. *There is a constant* C *so that if* ω *is a closed path that is not a Nielsen path,* $\sigma = \alpha\omega^k\beta$ *is a path, and* $k > 0$, *then at most* C *copies of* $[f^m(\omega)]$ *are canceled when* $[f^m(\alpha)][f^m(\omega^k)][f^m(\beta)]$ *is tightened to* $[f^m(\sigma)]$.

The following proposition is the analogue of the fact in linear algebra that if A is a unipotent matrix and v a nonzero vector, then projectively the sequence $\{A^k(v)\}$ converges to an eigenspace of A.

PROPOSITION 3.18. *Let $f : G \to G$ be a* UR *with edge E_i and suffix u_i. If $[f(u_i)] \neq u_i$, R^* is an initial segment of R_i, and σ is a path in G that crosses E_i or its inverse, then there is an N such that, for all $k > N$, $[f^k(\sigma)]$ contains R^* or its inverse as a subpath.*

Proof. We argue by induction on *height*(σ). If *height*$(\sigma) = i$, consider the splitting of $[f^M(\sigma)]$ into edges and exceptional paths (see Theorem 3.11(5)). There is a 1-1 correspondence between occurrences of E_i in σ and in $[f^M(\sigma)]$. Since $[f(u_i)] \neq u_i$, E_i does not occur in an exceptional path, and hence one of the subpaths in the splitting is E_i or $\overline{E}_i$. Eventually, the iterates contain R^* or its inverse.

Now assume *height*$(\sigma) = j > i$. Again consider the splitting of $[f^M(\sigma)]$ into edges and exceptional paths. We first claim that an exceptional path $E_s\tau^k\overline{E}_t$ cannot cross E_i or $\overline{E}_i$. Indeed, suppose that this exceptional path does cross E_i or $\overline{E}_i$. It must be then that τ crosses E_i or its inverse because the edges E_s and E_t have fixed suffixes and so are distinct from E_i and $\overline{E}_i$. But, τ cannot cross E_i of $\overline{E}_i$ for otherwise, since *height*$(\tau) < j$, it follows from the induction hypothesis that high iterates of τ (which equal τ) would have to contain arbitrarily long segments of R_i. This contradiction establishes the claim.

If the edge E_i or its inverse occurs in the splitting, we are done. Also, if there is an edge E_l in the splitting whose eigenray R_l crosses E_i, then high iterates of σ contain large segments of R_l, which in turn contain large iterates of u_l, and these eventually contain R^* by induction.

It remains to exclude the possibility that, for all large m, $[f^m(\sigma)]$ crosses only edges whose iterates do not cross E_i. Let G' be the f-invariant subgraph of G consisting of edges whose f-iterates do not cross E_i or $\overline{E}_i$. Since the f-image of an edge crosses that same edge, each component of G' is f-invariant. It follows that the restriction of f to the component G'_0 of G' that contains $[f^m(\sigma)]$, for large m, is a homotopy equivalence, see for example [BFH00, Lemma 6.0.6]. Therefore, σ is homotopic rel endpoints into G'_0. Thus, E_i is an edge in G'_0, a contradiction. □

4. The dynamics of unipotent automorphisms

4.1. *Poloynomial sequences.* Suppose $T \in \mathcal{T}$ and $\mathcal{O} \in \mathrm{UPG}(F_n)$. Our goal in this section is to show that there is a natural number $d = d(\mathcal{O}, T)$ such that the sequence $\{T\mathcal{O}^k/k^d\}_{k=0}^{\infty}$ converges to a tree $T\mathcal{O}^{\infty} \in \mathcal{T}$. This is the content of Theorem 4.22.

Theorem 4.22 will be proved by showing that if $f : G \to G$ is a UR for $\mathcal{O}$, if $h : G \to X$ is a homotopy equivalence from G to a model for T taking vertices to vertices, and if σ is a path in G with endpoints at vertices, then there is a polynomial P such that, for large k, the length of $[h(f^k(\sigma))]$ equals $P(k)$. Theorem 3.11 completely describes the $[f^k(\sigma)]$'s. To measure the length of $[h(f^k(\sigma))]$, we must first transfer $f^k(\sigma)$ to X via h and then put this path into normal form. The main work is in understanding the cancellation that occurs when $[h(f^k(\sigma))]$ is put into normal form. *All paths will be assumed to have endpoints that are vertices.*

The key properties of a sequence of paths $\{[f^k(\sigma)]\}_k$ are captured in the following definition.

Definition 4.1. Let G be a real graph. A sequence of paths in G is *polynomial* if it can be obtained from constant sequences of paths by finitely many operations of the following four basic types.

(1) (*re-indexing and truncation*): The sequence of paths $\{A_k\}_{k=k_0}^{\infty}$ is obtained from the sequence of paths $\{B_k\}_{k=k_1}^{\infty}$ by *re-indexing and truncation* if there is an integer $k' \geq k_1 - k_0$ such that $A_k = B_{k+k'}$.

(2) (*inversion*): The sequence of paths $\{A_k\}_{k=k_0}^{\infty}$ is obtained from the sequence of paths $\{B_k\}_{k=k_0}^{\infty}$ by *inversion* if A_k is the inverse of B_k.

(3) (*concatenation*): The sequence of paths $\{A_k\}_{k=k_0}^{\infty}$ is obtained from the sequences of paths $\{B_k\}_{k=k_0}^{\infty}$ and $\{C_k\}_{k=k_0}^{\infty}$ by *concatenation* if $A_k = B_kC_k$. (As the notation implies, no cancellation occurs in B_kC_k.)

(4) (*integration*): The sequence of paths $\{A_k\}_{k=k_0}^{\infty}$ is obtained from the sequence of paths $\{B_k\}_{k=k_0}^{\infty}$ by *integration* if

$$A_k = B_{k_0}B_{k_0+1}\cdots B_k.$$

(Again no cancellation occurs.)

For example, in a wedge of three circles with edges A, B, and C, the sequences $\{AB^kC\}$ and $\{ABAB^2AB^3\cdots AB^k\}$ are polynomial.

A sequence *eventually* has a property if it may be truncated and re-indexed so that the resulting sequence has the property. The elements of a sequence *eventually* have a property if only finitely many elements do not have the property.

LEMMA 4.2. *Let $f : G \to G$ be a* UR *of a unipotent automorphism. Let σ be a path in G. Then the sequence $\{[f^k(\sigma)]\}_{k=0}^{\infty}$ is eventually polynomial.*

Proof. We use induction on the height of σ. If the height is 1, the sequence is constant. For the induction step, replace σ by the iterate $[f^M(\sigma)]$ from Theorem 3.11 so that it splits into subpaths which are either single edges or exceptional paths. It suffices to prove the statement for these subpaths. The statement is clear for the exceptional subpaths. For a single edge E with

$f(E) = E \cdot u$, the sequence $\{[f^k(E)]\}$ is the concatenation of the constant sequence $\{E\}$ with the integral of the sequence $\{[f^k(u)]\}$ by Theorem 3.11(2). $\square$

The *complexity* of a polynomial sequence $\{A_k\}$, denoted

$$\text{complexity}(\{A_k\}),$$

is the minimal number of basic operations needed to make $\{A_k\}$. The complexity of a constant sequence is 0.

To measure the lengths of polynomial sequences of paths, we have polynomial sequences of numbers.

Definition 4.3. A sequence of nonnegative real numbers is *polynomial* if it can be obtained from constant sequences of nonnegative real numbers by finitely many operations of the following three basic types.

(1) (*re-indexing and truncation*): The sequence $\{p_k\}_{k=k_0}^{\infty}$ is obtained from the sequence $\{q_k\}_{k=k_1}^{\infty}$ by *re-indexing and truncation* if there is an integer $k' \geq k_1 - k_0$ such that $p_k = q_{k+k'}$.

(2) (*concatenation*): The sequence $\{p_k\}_{k=k_0}^{\infty}$ is obtained from the sequences $\{q_k\}_{k=k_0}^{\infty}$ and $\{r_k\}_{k=k_0}^{\infty}$ by *concatenation* if $p_k = q_k + r_k$.

(3) (*integration*): The sequence $\{p_k\}_{k=k_0}^{\infty}$ is obtained from the sequence $\{q_k\}_{k=k_0}^{\infty}$ by *integration* if

$$p_k = q_{k_0} + q_{k_0+1} + \cdots + q_k.$$

The following lemma is immediate from the definitions.

LEMMA 4.4. *Let $\{A_k\}_{k=k_0}^{\infty}$ be a polynomial sequence of paths in G. Then the sequence of nonnegative real numbers $\{\text{length}_G(A_k)\}_{k=k_0}^{\infty}$ is polynomial.*

The *complexity* of a polynomial sequence $\{p_k\}$, denoted

$$\text{complexity}(\{p_k\}),$$

is the minimal number of basic operations needed to make $\{p_k\}$. The complexity of a constant sequence is 0.

LEMMA 4.5. (1) *If 0 is an element of a polynomial sequence of nonnegative real numbers, then the sequence is constantly 0.*

(2) *Unless a polynomial sequence of nonnegative real numbers is constant, it is increasing.*

(3a) *If $\{p_k\}$ is a polynomial sequence of real numbers, then there is a polynomial $P \in \mathbb{R}[x]$ such that $P(k) = p_k$.*

(3b) *If $\{p_k\}$ is not constant, if $\{m_{j,k}\}_k$ are the positive constant sequences used in the integration operations in a particular construction of $\{p_k\}$, if $m = \min_j\{m_{j,k}\}$, and if P has degree d, then the leading coefficient of P is bounded below by $m/d!$.*

(4) *If $\{p_k\}$ is a polynomial sequence of nonnegative real numbers and if $c \in \mathbb{R}$ is eventually not greater than p_k, then the sequence $\{p_k - c\}$ is eventually polynomial.*

Proof. In each case, the proof is by induction on complexity.

(1) The statement is true for constant sequences. If q_k and r_k are never 0, then the same is true for $q_k + r_k$ and $q_{k_0} + \cdots + q_k$.

(2) The sum of constant sequences is constant. The sum of increasing and constant sequences is increasing as long as at least one of the sequences is increasing.

(3) The proof is an induction on the complexity of $\{p_k\}$. If $\{p_k\}$ is constant then $P(k) = p_k$ for a constant polynomial P. Suppose $\{q_k\}$ and $\{r_k\}$ are polynomial sequences, that Q and R are polynomials with $Q(k) = q_k$ and $R(k) = r_k$, that the leading coefficients of Q and R are respectively Q_0 and R_0, that $\deg(Q) \geq \deg(R)$, and that $\deg(Q) \geq 1$.

If $\{p_k\}$ is obtained from $\{q_k\}$ by re-indexing and truncation, then there is a polynomial P with the same degree and leading coefficient as Q so that $P(k) = p_k$. If $p_k = q_k + r_k$ then $P = Q + R$. The leading coefficient of P is Q_0 if $\deg(Q) > \deg(R)$ and is $Q_0 + R_0$ otherwise.

Finally, suppose that $\{p_k\}$ is obtained from $\{r_k\}$ by integration. We will need the fact that $\sum_{i=0}^{k} i^d$ is a polynomial of degree $d{+}1$ with leading coefficient $1/(d+1)$. Using the quoted fact, there is a polynomial P such that $\deg(P) = \deg(R) + 1$, the leading coefficient of P is $R_0/\deg(P)$, and $P(k) = p_k$. Item (3) follows easily.

(4) The statement is clear for constant sequences and if $\{p_k\}$ is obtained by re-indexing and truncating a sequence where the lemma holds. If $\{p_k\}$ is the sum of sequences $\{q_k\}$ and $\{r_k\}$ for which the statement holds and where $\{q_k\}$ is not constant, then $\{q_k - c\}$ is eventually polynomial and hence so is $\{p_k - c\} = \{q_k - c\} + \{r_k\}$. Finally, suppose $\{p_k\}_{k=k_0}^{\infty}$ is obtained from the nonzero sequence $\{q_k\}$ by integration. Suppose that $q_{k_0} + q_{k_0+1} + \cdots + q_{k_1} > c$. Then $p_k - c$ is eventually $(q_{k_0} + \cdots + q_{k_1} - c) + (q_{k_1+1} + q_{k_1+2} + \cdots + q_k)$. Thus, after re-indexing and truncating, we see that $\{p_k - c\}$ is the sum of a constant sequence and the integral of a polynomial sequence. In particular, $\{p_k - c\}$ is eventually polynomial. □

We now record some general properties of polynomial sequences of paths that will be needed.

LEMMA 4.6 (Stability in G). *If $\{A_k\}$ is a polynomial sequence of paths in G, then either $\{A_k\}$ is constant or, for all N, the initial and terminal paths of A_k of length N are eventually constant.*

Proof. The proof is by induction on the complexity of $\{A_k\}$. The conclusion is true if $\{A_k\}$ is constant. The re-indexing and truncation step and the inversion step follow immediately from definitions. Suppose that $\{B_k\}$ and $\{C_k\}$ are polynomial sequences in G for which the conclusion holds.

If $\{A_k\}$ is obtained from $\{B_k\}$ and $\{C_k\}$ by concatenation, then $\{A_k\}$ is constant if and only if $\{B_k\}$ and $\{C_k\}$ are constant. Suppose that $\{B_k = B\}$ is constant, but that $\{C_k\}$ is not. If C is eventually the initial path of length N of C_k, then eventually the initial path of length N of A_k is the initial path of length N of BC. Eventually, the terminal path of length N of A_k is the terminal path of length N of C_k. The case where $\{B_k\}$ is not constant, but $\{C_k\}$ is constant is symmetric. Finally, if neither $\{B_k\}$ nor $\{C_k\}$ is constant, then eventually the initial (respectively terminal) path of length N of A_k equals the initial path of length N of B_k (respectively C_k).

Suppose that $\{A_k\}$ is obtained from $\{B_k\}$ by integration. By definition, the initial path of length N of A_k is eventually constant. If $\{B_k = B\}$ is constant, then eventually the terminal path of length N of A_k is the terminal path of length N of a concatenation of B's. If $\{B_k\}$ is not constant, then eventually the terminal path of length N of A_k is the terminal path of length N of B_k. □

LEMMA 4.7. *If $\{B_k\}$ is a polynomial sequence of paths in G and if $\{A\}$ and $\{C\}$ are constant sequences such that eventually the terminal endpoint of A is the initial endpoint of B_k and the terminal endpoint of B_k is the initial endpoint of C, then the sequence $\{[AB_kC]\}$ is eventually polynomial.*

Proof. The proof is by induction on the complexity of $\{B_k\}$. The statement is true if $\{B_k\}$ is constant. The re-indexing and truncation step and the inversion step follow immediately from definitions.

Suppose $\{B_k\} = \{B'_k B''_k\}$ is the concatenation of $\{B'_k\}$ and $\{B''_k\}$ and the lemma holds for $\{B'_k\}$ and $\{B''_k\}$. If $\{B'_k\}$ is constant, then

$$[AB_kC] = [AB'_k B''_k C] = [[AB'_k][B''_k C]]$$

and we are done by hypothesis. The case that $\{B''_k\}$ is constant is similar. If $\{B'_k\}$ and $\{B''_k\}$ are not constant then eventually $[AB'_k B''_k C] = [AB'_k][B''_k C]$ and again we are done by hypothesis.

Finally, suppose $\{B_k\}$ is obtained from $\{B'_k\}_{k=k_0}^{\infty}$ by integration and that the lemma holds for $\{B'_k\}$. Choose N so that

$$N \cdot \mathit{length}_G(B_{k_0}) > \max\{\mathit{length}_G(A), \mathit{length}_G(C)\}.$$

Then, eventually

$$\begin{aligned}[AB_kC] &= [AB'_{k_0}B'_{k_0+1}\cdots B'_kC] \\ &= [AB'_{k_0}\cdots B'_{k_0+N}][B'_{k_0+N+1}\cdots B'_{k-N}][B'_{k-N+1}\cdots B'_kC].\end{aligned}$$

Since we have already verified the concatenation step, all three terms give polynomial sequences. □

Definition 4.8. A subgroup H of a group J is *primitive* if, for all $a \in J$ and all $i \neq 0$, $a^i \in H$ implies that $a \in H$. An element a of J is *primitive* if $\langle a \rangle$ is a primitive subgroup of J.

LEMMA 4.9. *Let $G' \to G$ be an immersion of finite graphs such that $\mathrm{Im}[\pi_1(G') \to \pi_1(G)]$ is a primitive finitely generated subgroup of $\pi_1(G)$. Let $\{A_k\}$ be a polynomial sequence of paths in G. Assume that for infinitely many values of k the path A_k lifts to G' starting at a given vertex $x \in G'$. Then the same is true for all large k. Furthermore, the lifts form* (*after truncation*) *a polynomial sequence in G'* (*so that in particular — see Lemma* 4.6 *— the terminal endpoint of these lifts is constant*).

The lemma fails if the primitivity assumption is dropped; e.g. take G to be the circle and G' the double cover.

Proof. We proceed by induction on the complexity of $\{A_k\}$.

Suppose first that the last operation is inversion. For infinitely many k the other endpoint of the lift of A_k starting at x is a point $y \in G'$ (there are finitely many preimages of the common terminal endpoint of the A_k's in G). Applying the statement of the lemma to $\{\overline{A}_k\}$ we learn that for all large k there is a lift A'_k of A_k that terminates at y. For infinitely many k, A'_k starts at x, and $\{A'_k\}$ forms a polynomial sequence. Therefore, for all large k, A'_k starts at x.

Suppose next that the last operation is concatenation: $A_k = B_kC_k$. Then B_k lifts to G' starting at x for infinitely many k and thus for all large k, and the lifts B'_k form a polynomial sequence. Let y be the common terminal endpoint of the B'_k. Similarly, for all large k the path C_k lifts to a path C'_k starting at y, and these paths form a polynomial sequence. Thus $A'_k = B'_kC'_k$ is a polynomial sequence starting at x and projecting to A_k.

Finally, suppose that the last operation is integration:

$$A_k = B_1B_2\cdots B_k.$$

Since A_k is a subpath of A_l for all $l \geq k$, it follows from our assumptions that each A_k lifts to a path A'_k starting at $x_0 = x$. Infinitely many of these end at the same point y_1. Thus for infinitely many k the path B_k lifts starting at y_1. It follows that, for all sufficiently large k, B_k lifts to a path starting at y_1

and ending at a point y_2. Repeating this procedure, we produce a sequence $\{y_1, y_2, \cdots\}$ such that B_k eventually lifts to a path starting at y_j and ending at y_{j+1}. Choose $i < j$ such that $y_i = y_j$. For large k there are lifts of B_k that connect y_i to y_{i+1}, y_{i+1} to y_{i+2},..., y_{j-1} to y_j. By the primitivity assumption we must have $y_i = y_{i+1} = \cdots = y_j$. Therefore the sequence $\{y_1, y_2, \cdots\}$ is eventually constant, i.e. $y_k = y$ for all large k. Thus for large k the path B_k lifts to B'_k beginning and ending at y. The claim now follows. □

LEMMA 4.10. *If $\{A_k\}$ is a polynomial sequence of closed paths in G based at say x, if Z is a primitive cyclic subgroup of $\pi_1(G, x)$, and if, for infinitely many k, A_k represents an element of Z, then eventually A_k represents an element of Z.*

Proof. Let $(G', x') \to (G, x)$ be an immersion such that $\mathrm{Im}[\pi_1(G', x') \to \pi_1(G, x)] = Z$. By Lemma 4.9, eventually A_k lifts to a closed path in G' based at x'. □

Throughout the rest of this section, $p : \Gamma \to G$ denotes the universal covering of the marked real graph G, T is a tree in in $\mathcal{T}$, X is a model for T, and $h : G \to X$ a cellular homotopy equivalence such that the image of each edge of G is in normal form. By subdividing if necessary, we may assume that $h^{-1}(X^{(0)}) = G^{(0)}$. In particular, the h-image of each edge of G is either horizontal or vertical. Edges of the former type are *h-horizontal*; edges of the latter type are *h-vertical*. Recall (Section 2.7) that X comes with a map $q : X \to \overline{T} = T/F_n$. By Proposition 2.5, $\mathrm{BCC}(\widetilde{qh} : \Gamma \to T) < \infty$ where $\widetilde{qh}$ is a lift of qh. By definition, if AB is a concatenation of paths in G, then the overlap length of $[h(\overline{A})]$ and $[h(B)]$ is less than $\mathrm{BCC}(\widetilde{qh})$.

The main technical result of this section is Proposition 4.21 which states that if $\{A_k\}$ is a polynomial sequence of closed paths in G, then $\{\ell_T([[A_k]])\}$ is eventually a polynomial sequence of real numbers.

Example 4.11. Suppose that $G = \mathrm{Rose}_2$ with X and f as in Example 2.7. Suppose also that the edge 't' has length 1. The sequence $\{e_1 e_2^{-1} e_1 e_2^k\}_{k=1}^\infty$ is polynomial. Yet,

$$\{\ell_T([[e_1 e_2^{-1} e_1 e_2^k]]\}_{k=1}^\infty = \{0, 1, 2, 3 \cdots\}$$

is eventually polynomial, but is not a polynomial sequence of nonnegative real numbers (see Lemma 4.5(1)).

Definition 4.12. A polynomial sequence $\{A_k\}$ of paths in G is *elliptic* (with respect to h) if $[h(A_k)] = \nu_k$ with ν_k vertical. A sequence $\{\nu_k\}$ of vertical edge paths in X is *elliptic* if for some elliptic sequence $\{A_k\}$ of edge paths in G, $[h(A_k)] = \nu_k$.

LEMMA 4.13. *The image of $G^{(0)}$ under h is $X^{(0)}$. If $x, y \in G^{(0)}$ and if P is a path in X between $h(x)$ and $h(y)$, then there is a unique path Q in G between x and y such that $[hQ] = [P]$.*

Proof. We are assuming that $h^{-1}(X^{(0)}) = G^{(0)}$. Since T is minimal, qh is onto. Since q induces a bijection between $X^{(0)}$ and $\overline{T}^{(0)}$, the first statement follows. The second statement follows from the assumption that h is a homotopy equivalence. □

Definition 4.14. A polynomial sequence of paths in G is *short* if it is a concatenation of constant and elliptic sequences.

LEMMA 4.15. *Let $\{A_k\}$ be a short polynomial sequence in G. Then A_k is a concatenation of paths*

$$(*) \qquad A_k = V_{0,k}\hat{H}_1 V_{1,k}\hat{H}_2 \cdots \hat{H}_M V_{M,k}$$

such that

(1) *$\{V_{i,k}\}_k$ is an elliptic sequence in G,*

(2) *$\hat{H}_i$ is an h-horizontal edge of G, and*

(3) *eventually $[h(A_k)] = h(V_{0,k})H_1 h(V_{1,k})H_2 \cdots H_M h(V_{M,k})$ where $h(\hat{H}_i) = H_i$.*

Proof. By writing elements of constant sequences as concatenations of h-vertical and h-horizontal edges, we have

$$A_k = V_{0,k}\hat{H}_1 V_{1,k}\hat{H}_2 \cdots \hat{H}_M V_{M,k}$$

with all the desired properties except perhaps Item (3).

We proceed by induction on the number of elliptic sequences in the concatenation. If there are no elliptic sequences, then $\{A_k\}$ is constant and the conclusion follows. If $[h(A_k)]$ is not eventually as in Item (4) then there is an i such that, for infinitely many k, $\hat{H}_i V_{i,k}\hat{H}_{i+1}$ is elliptic. For these values of k, the path $[h(\hat{H}_i V_{i,k}\hat{H}_{i+1})]$ has a common initial and terminal endpoint z. Let x be the initial endpoint of $\hat{H}_i$ and let y be the terminal endpoint of $\hat{H}_{i+1}$. By Lemma 4.13, there is path σ in G connecting y to x such that $[h(\sigma)]$ is the trivial path at z. By Lemma 4.7, $[\hat{H}_i V_{i,k}\hat{H}_{i+1}\sigma]$ is a polynomial sequence of paths. By Lemma 4.10 and the fact that edge stabilizers in T are primitive cyclic, $[\hat{H}_i V_{i,k}\hat{H}_{i+1}\sigma]$ and hence $[\hat{H}_i V_{i,k}\hat{H}_{i+1}]$ is eventually elliptic. Thus, $V'_k := V_{i-1,k}\hat{H}_i V_{i,k}\hat{H}_{i+1}V_{i+1,k}$ is eventually elliptic. Now,

$$A_k = V_{0,k}\hat{H}_1 V_{1,k}\hat{H}_2 \cdots \hat{H}_{i-1}V'_k\hat{H}_{i+2} \cdots \hat{H}_M V_{M,k}$$

has fewer elliptic sequences in the concatenation. □

Definition 4.16. We call the expression (∗) a *stable normal form for* $\{A_k\}$. It is obtained by *erasing pairs of h-horizontal edges* $\hat{H}_i$ *and* $\hat{H}_{i+1}$ such that $[h(V_{i-1,k}\hat{H}_i V_{i,k}\hat{H}_{i+1}V_{i+1,k})]$ is vertical for infinitely many k (equivalently eventually vertical).

Remark 4.17. It follows from Lemma 4.15 that if $\{A_k\}$ is short then the length in X of $[h(A_k)]$ is eventually constant.

Definition 4.18. A polynomial sequence of paths $\{A_k\}$ in G is *long* if, given $N > 0$, there are sequences $\{B_k\}$, $\{C_k\}$, and $\{D_k\}$ such that

(1) $\{B_k\}$ and $\{D_k\}$ are short,

(2) eventually the lengths of $[h(B_k)]$ and $[h(D_k)]$ are at least N, and

(3) eventually $A_k = B_k C_k D_k$.

Lemma 4.19. *Suppose that* $\{B_k\}$ *is short. Then, the integral* $\{A_k\} = \{B_1 B_2 \cdots B_k\}$ *either is eventually long or eventually short.*

Proof. Let $B_k = V_{0,k}\hat{H}_1 V_{1,k}\hat{H}_2 \cdots \hat{H}_M V_{M,k}$ be a stable normal form. Suppose that the stable normal form for the short polynomial sequence $\{B_k B_{k+1}\}$ has N h-horizontal edges. There are two cases.

Case 1. Assume $N > M$. Since the stable normal form for $B_k B_{k+1}$ is obtained by erasing terminal h-horizontal edges from B_k and initial h-horizontal edges from B_{k+1},

- eventually $B_k = B_{1,k}B_{2,k}B_{3,k}$ where $\{B_{i,k}\}_k$ are polynomial sequences,
- $\{V_k\} := \{B_{3,k}B_{1,k+1}\}$ is eventually elliptic,
- the stable normal form for $B_{2,k}$ has a positive number of h-horizontal edges, and
- eventually

$$[h(B_k B_{k+1})] = [h(B_{1,k})][h(B_{2,k})][h(V_k)][h(B_{2,k+1})][h(B_{3,k+1})].$$

It follows that eventually

$$\begin{aligned}[h(B_k \cdots B_{k+L})] = {} & [h(B_{1,k})][h(B_{2,k})] \\ & \big([h(V_k)][h(B_{2,k+1})][h(V_{k+1})][h(B_{2,k+2})] \cdots [h(V_{k+L-1})][h(B_{2,k+L})]\big) \\ & [h(B_{3,k+L})].\end{aligned}$$

In particular, the length in X of $[h(B_k \cdots B_{k+L})]$ goes to infinity with L.

Since

$$B_1B_2\cdots B_k = (B_1\cdots B_L)(B_{L+1}B_{L+2}\cdots B_{k-L})(B_{k-L+1}\cdots B_k),$$

the integral of $\{B_k\}$ is eventually a concatenation of three polynomial sequences. The first and third are short. By choosing L large, the length in X of the stable normal form of the h-image of the first and third sequences can be made arbitrarily large. It follows that this integral is eventually long.

Case 2. Assume $N \le M$. Then, at least half of the terminal h-horizontal edges of B_k are erased with initial h-horizontal edges of B_{k+1} in putting B_kB_{k+1} into normal form. Note that in this case, M is even. Indeed, otherwise the middle h-horizontal edge of B_k is erased with the middle h-horizontal edge of B_{k+1}; but these are the same oriented edges, a contradiction. Set

$$V_k := V_{\frac{M}{2},k}\hat{H}_{\frac{M}{2}+1}\cdots V_{k-1,k}\hat{H}_kV_{M,k}V_{0,k+1}\hat{H}_1\cdots V_{\frac{M}{2}-1,k+1}\hat{H}_{\frac{M}{2}}V_{\frac{M}{2},k+1}.$$

By assumption, $\{V_k\}$ is eventually an elliptic polynomial sequence. Now

$$B_1B_2\cdots B_k = V_{0,1}\hat{H}_1\cdots V_{\frac{M}{2}-1,1}\hat{H}_{\frac{M}{2}}V_1V_2\cdots V_{k-1}\hat{H}_{\frac{M}{2}+1}V_{\frac{M}{2}+1,k}\cdots \hat{H}_MV_{M,k}.$$

We see that the integral of $\{B_k\}$ is eventually short. □

LEMMA 4.20. (1) *A polynomial sequence of paths $\{A_k\}$ in G is either short or long.*

(2) *Suppose that $\{A_{1,k}\}$, $\{A_{2,k}\}$ and $\{A_k\} := \{A_{1,k}A_{2,k}\}$ are polynomial sequences of paths in G. The overlap length in X of $[h(\overline{A}_{1,k})]$ and $[h(A_{2,k})]$ is eventually constant.*

Proof. (1) The proof is by induction on the complexity of $\{A_k\}$. Constant sequences are short. If $\{A_k\}$ is obtained by truncating and re-indexing a short (respectively long) sequence, then $\{A_k\}$ is short (respectively long). The inverse of a short (respectively long) sequence is short (respectively long).

Suppose $\{B_k\}$ and $\{C_k\}$ are short. Suppose $\{D_k\} = \{D_{1,k}D_{2,k}D_{3,k}\}$ is long with $\{D_{1,k}\}$ and $\{D_{3,k}\}$ short. Suppose $\{E_k\} = \{E_{1,k}E_{2,k}E_{3,k}\}$ is long with $\{E_{1,k}\}$ and $\{E_{3,k}\}$ short. If $\{A_k\} = \{B_kC_k\}$ then $\{A_k\}$ is short. If $\{A_k\} = \{D_kE_k\} = \{(D_{1,k})(D_{2,k}D_{3,k}E_{1,k}E_{2,k})(E_{3,k})\}$ then $\{A_k\}$ is long. If $\{A_k\} = \{B_kE_k\} = \{(B_kE_{1,k})(E_{2,k})(E_{3,k})\}$, then $\{A_k\}$ is long. Indeed, $\{B_kE_{1,k}\}$ and $\{E_{3,k}\}$ are short, and if the lengths in X of $[h(E_{1,k})]$ and $[h(E_{3,k})]$ are greater than $N + C$ where C is the eventual length of $[h(B_k)]$, then the lengths in X of $[h(B_kE_{2,k})]$ and $[h(E_{3,k})]$ are eventually greater than N. The other case $\{A_k\} = \{D_kC_k\}$ is symmetric with the case $\{A_k\} = \{B_kE_k\}$.

If $\{A_k\}$ is the integral of a short sequence, then the conclusion follows from Lemma 4.19. Suppose that $\{A_k\}$ is the integral of a long sequence $\{D_k\}$ as above. Suppose further that the lengths of $[h(D_{1,k})]$ and $[h(D_{3,k})]$ are greater than $2N$ if $k \geq K$ and that N is greater than $\mathrm{BCC}(\widetilde{qh})$. Then, eventually

$$A_k = (D_1 D_2 \cdots D_K)(D_{K+1} D_{K+2} \cdots D_{k-1} D_{1,k} D_{2,k})(D_{3,k}).$$

Further, the length in X of $[h(D_1 D_2 \cdots D_K)]$ is at least N. It follows that $\{A_k\}$ is long.

(2) There are four cases depending on whether or not $\{A_{1,k}\}$ or $\{A_{2,k}\}$ is short or long. Let $\{B_k\}$, $\{C_k\}$, $\{D_k\}$, and $\{E_k\}$ be as in the proof of (1) above. If both $\{A_{1,k}\}$ and $\{A_{2,k}\}$ are short, then so is $\{A_k\}$, and the conclusion follows from Lemma 4.15.

Choose N to be longer than $\mathrm{BCC}(\widetilde{qh})$. Then, the overlap length of $[h(\overline{A}_{1,k})]$ and $[h(A_{2,k})]$ is less than N. Suppose that $\{A_{1,k}\} = \{D_k\}$ and $\{A_{2,k}\} = \{E_k\}$. Suppose that the lengths in X of $[h(D_{3,k})]$ and $[h(E_{1,k})]$ are eventually greater than N. Then, the overlap length of $[h(\overline{A}_{1,k})]$ and $[h(A_{2,k})]$ is eventually the same as the overlap length of $[h(\overline{D}_{3,k})]$ and $[h(E_{1,k})]$. Thus, we are reduced to the case of short sequences.

The other two cases, $\{A_{1,k}\} = \{B_k\}$, $\{A_{2,k}\} = \{E_k\}$ and $\{A_{1,k}\} = \{D_k\}$, $\{A_{2,k}\} = \{C_k\}$, similarly reduce to the case of short sequences. □

PROPOSITION 4.21. *If $\{A_k\}$ is a polynomial sequence of paths (respectively closed paths) in G, then the sequence $\{\mathit{length}_X([h(A_k)])\}$ (respectively $\{\ell_T([[A_k]])\}$) of real numbers is eventually polynomial. If P is a polynomial such that eventually $P(k) = \mathit{length}_X([h(A_k)])$ (respectively $\{\ell_T([[A_k]])\}$) then the leading coefficient of P is bounded below by $m/\deg(P)!$ where m is the length of the shortest edge of T.*

Proof. We prove the statements about sequences of paths; the case of closed paths is similar and is left to the reader. The proof is by induction on the complexity of $\{A_k\}$ with the induction statement being that $\{\mathit{length}_X([h(A_k)])\}$ is eventually polynomial and that the positive constant sequences of real numbers used in the integration operations are realized as lengths of edge paths in T. This directly proves the first statement of the proposition and the second follows from Lemma 4.5(3b).

The case that $\{A_k\}$ has zero complexity follows from the first item in Remark 0.1. We now assume the desired result for all polynomial paths with complexity less than that of $\{A_k\}$.

By Lemma 4.20(1), $\{A_k\}$ either is short or long. In the former case, the induction statement follows from Remark 4.17 and the first item in Remark 0.1. We now assume that $\{A_k\}$ is long.

If $\{A_k\}$ is obtained from a sequence satisfying the inductive statement by inversion or by truncation and re-indexing, then $\{A_k\}$ satisfies the inductive statement.

Suppose that $\{A_k\}$ is the concatenation of $\{A_{1,k}\}$ and $\{A_{2,k}\}$ where both $\{A_{1,k}\}$ and $\{A_{2,k}\}$ satisfy the inductive statement, and at least one of $\{A_{1,k}\}$ and $\{A_{2,k}\}$ is not short. Let A be the eventual overlap length of $[h(\overline{A}_{1,k})]$ and $[h(A_{2,k})]$ (see Lemma 4.20(2)). Then, eventually the length in X of $[h(A_k)]$ equals

$$length_X([h(A_{1,k})]) + length_X([h(A_{2,k})]) - 2A.$$

Lemma 4.5(4) implies that $\{length_X([h(A_k)])\}$ is eventually polynomial and hence that the inductive statement is satisfied.

Finally, suppose that $\{A_k\}$ is the integral of the sequence $\{B_k\}$ where $\{B_k\}$ satisfies the induction statement. By Lemma 4.20(2), the overlap length of $[h(\overline{B}_k)]$ and $[h(B_{k+1})]$ is eventually a constant B. After re-indexing and truncating, $\{length_X([h(A_k)])\}$ is $\{2B\}$ plus the integral of $\{length_X([h(B_k)]) - 2B\}$. The latter sequence is polynomial by Lemma 4.5(4) and is eventually the length of an edge path in T by the second item in Remark 0.1 with $D = D' = B$. This completes the induction step and so also the proof of the proposition. □

Let $T \in \mathcal{T}$, $\mathcal{O} \in \mathrm{UPG}_{F_n}$, and $a \in F_n$. By Lemma 4.2 and Proposition 4.21, the sequence $\{\ell_T(\mathcal{O}^k([[a]]))\}$ is eventually polynomial. The degree $d(\mathcal{O}, T, a)$ of the polynomial is uniformly bounded by the number of strata in a UR for $\mathcal{O}$. Let $d(\mathcal{O}, T) = \max\{d(\mathcal{O}, T, a) \mid a \in F_n\}$.

THEOREM 4.22. *Suppose $T \in \mathcal{T}$ and $\mathcal{O} \in \mathrm{UPG}(F_n)$. Set $d = d(\mathcal{O}, T)$. Then, the sequence $\{T\mathcal{O}^k/k^d\}$ converges to a tree $T\mathcal{O}^\infty \in \mathcal{T}$.*

Proof. By Lemma 4.2 and Proposition 4.21 and the definition of d, the sequences

$$\{\ell_T(\mathcal{O}^k([[a]]))/k^d\}$$

converge for all $a \in F_n$ and not all of these limits are 0. Hence, the sequence $\{T\mathcal{O}^k/d^k\}$ converges to a nontrivial F_n-tree. That this tree is very small follows from the fact, proved in [CL95], that $\mathcal{T}_{VS} \cup \{\text{trivial action}\}$ is closed under limits.

By Lemma 4.5(3), if P is a polynomial of degree d such that eventually $P(k) = \ell_T(\mathcal{O}^k([[a]]))$, then the leading coefficient of P is bounded below by a nonzero constant depending only on d and the lengths of the edges of T. It follows that the collection of nonzero numbers of the form $\lim \ell_T(\mathcal{O}^k([[a]]))/k^d$ are bounded away from 0, and so $T\mathcal{O}^\infty$ is simplicial. □

Question. Does Theorem 4.22 hold for trees in $\mathcal{T}_{VS}$?

Definition 4.23. Trees T for which $d(\mathcal{O}, T) = 0$, i.e. for which the sequence

$$\{\ell_T(\mathcal{O}^k([[\gamma]]))\}_{k=0}^\infty$$

is eventually constant for every $\gamma \in F_n$, are *$\mathcal{O}$-nongrowers.* Others are *$\mathcal{O}$-growers.* If $d(\mathcal{O}, T) = 1$, then we say that T *grows linearly under $\mathcal{O}$.*

Remark 4.24. There exist nongrowers that are not fixed. An example is the F_2-tree T pictured below. The first map $\mathrm{Rose}_2 \to X$ is the marking and is given by identifying the two bold arcs. (Recall that the edges of Rose_2 are identified with the elements of the standard basis $\{e_1, e_2\}$.) All edges of X are horizontal except for the image e_1. The tree T is $\mathcal{O}$-nongrowing where $\mathcal{O}$ is represented by the automorphism given by $e_1 \mapsto e_1$, $e_2 \mapsto e_1e_2$. Such examples do not exist in $\mathrm{SL}_n(\mathbb{Z})$ or the mapping class group of a surface.

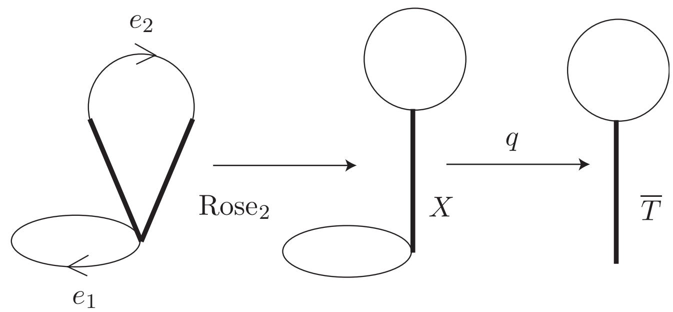

Nongrowers are also responsible for the existence of automorphisms $\mathcal{O}$ and compact sets $K \subset \mathbb{P}\mathcal{T}_{VS}$ in the complement of $\mathrm{Fix}(\mathcal{O})$ with the property that for no k is $K\mathcal{O}^k$ contained in a prescribed small neighborhood of $\mathrm{Fix}(\mathcal{O})$. Here $\mathbb{P}\mathcal{T}_{VS}$ is the space of homothety classes of very small F_n-trees. A concrete example can be described as follows. Let $\mathcal{O} \in \mathrm{Out}(F_4)$ be represented by the automorphism given by $e_1 \mapsto e_1$ and $e_i \mapsto e_ie_1$ if $i > 1$. The compact set K consists of the bi-infinite sequence $\{\dots, T_{-2}, T_{-1}, T_0, T_1, T_2, \dots\}$ together with the limiting tree T_∞. The tree $\overline{T}_k$ is pictured below. Here we have indicated the marking by the elements of F_4 on the closed edges (i.e. loops) of X. Only the edge marked e_1 is vertical. Passing to the limit as k goes to infinity amounts to opening up the loops of $\overline{T}$ which in the limit correspond to e_2 and e_3. Now notice that $T_k\mathcal{O}^k$ converges to a nonfixed nongrower (which is a tree just like T_∞ except for a permutation of $\{e_2, e_3, e_4\}$).

It is, however, true that if K is a compact subset of $\mathbb{P}\mathcal{T}_{VS}$ consisting of growers, then the accumulation set of the sequence $K\mathcal{O}^k$ is a subset of $\mathrm{Fix}(\mathcal{O})$.

4.2. *Explicit limits.* The goal of this section is to give an explicit construction of $T\mathcal{O}^\infty$ for $T \in \mathcal{T}$ in the case that $d = d(\mathcal{O}, T) \geq 1$, *an assumption we make for the rest of this section.* The case that $\mathcal{O}$ has linear growth was done by Cohen and Lustig in [CL95]. Let $f : G \to G$ be a UR representing $\mathcal{O}$. Unless otherwise stated, paths in graphs are assumed to begin and end at vertices.

Notation 4.25. For a path σ in G, $f_\#(\sigma)$ is the path $[f(\sigma)]$.

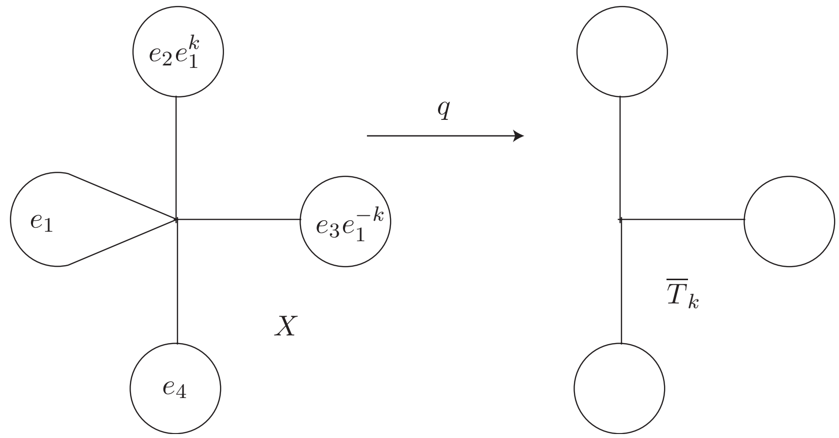

By Theorem 3.11(5), a sufficiently high iterate of a path σ has a splitting into exceptional subpaths and single edges. The following lemma implies that σ itself has a decomposition (not necessarily a splitting) of this sort.

LEMMA 4.26. *Maximal exceptional subpaths of a path do not overlap.*

Proof. Suppose that $\sigma_1 = E_i\tau^k\bar{E}_j$ is an exceptional subpath of σ and that σ_2 is an exceptional subpath of σ that intersects σ_1 in a proper nontrivial subpath of σ_1. Either the initial edge of σ_2 is contained in σ_1 and is not E_i or the terminal edge of σ_2 is contained in σ_1 and is not $\bar{E}_j$. These two cases are interchanged when σ is replaced by $\bar{\sigma}$ so that there is no loss in assuming that the former holds. The heights of the first and last edges of an exceptional path are strictly greater than the height of any other edge in the exceptional path. Thus if σ_2 contains $\bar{E}_j$ then $\bar{E}_j$ is the first or last edge of σ_2. Since no exceptional path begins with $\bar{E}_j$, σ_2 cannot extend past $\bar{E}_j$ and σ_2 is a subpath of σ_1. □

Lemma 4.26 implies that there is a well defined decomposition $\sigma = \sigma_1 \dots \sigma_m$ into single edges and maximal exceptional subpaths. We call this the *canonical decomposition of* σ.

Definition 4.27. Let $h : G \to X$ be a homotopy equivalence of G to a model for T taking vertices to vertices. For any path $\sigma \subset G$ with endpoints at vertices let $d(\sigma) = d(\sigma, X)$ be the degree of the polynomial sequence $\mathit{length}_X([h(f^k(\sigma))])$, and let

$$L_G(\sigma) = \lim_{k\to\infty} \mathit{length}_X([h(f^k(\sigma))])/k^d.$$

The notation L_G is chosen to remind the reader that L_G is a length function on paths in G. Proposition 4.21 implies that $L_G(\sigma)$ is well defined. By definition,

$L_G([f(\sigma)]) = L_G(\sigma)$ for all σ. Theorem 4.22 implies that L_G agrees with the length function induced by $T\mathcal{O}^\infty$ on circuits in G.

The problem with this definition of L_G is that it requires iteration and taking limits. For any path $\sigma \subset G$, let $L^*_G(\sigma) = \sum_{i=0}^m L_G(\sigma_i)$ where $\sigma = \sigma_1 \dots \sigma_m$ is the canonical decomposition of σ. We show below (Proposition 4.32) that $L_G = L^*_G$. Along with Lemma 4.28, this allows us to compute L_G directly once its value on a finite number of paths is known.

We say that an edge E_i is *linear* if its suffix u_i is a Nielsen path; in this case we write $u_i = \tau_i^{n_i}$ for a primitive Nielsen path τ_i and some $n_i > 0$. If τ is a primitive Nielsen path and if there is more than one linear edge whose suffix is an iterate of τ then we say that $\{E_i : \tau_i = \tau\}$ is the *linear family* associated to τ. By reordering the edges of G we may assume that the edges in a linear family are consecutively numbered, say $\{E_i : s \le i \le t\}$.

The following lemma gives $L_G(\sigma)$ when σ is a linear edge or an exceptional path.

Lemma 4.28. *Suppose that τ is a primitive Nielsen path, that E_i and E_k belong to the linear family associated to τ and that $\sigma = E_i\tau^m\bar{E}_k$. Then*

(1) *If $\ell_T([[\tau]]) = 0$ then $d(E_i) = d(E_k) = d(\sigma) = 0$ and $L_G(E_i) = L_G(E_k) = L_G(\sigma) = 0$.*

(2) *If $\ell_T([[\tau]]) > 0$ then $d(E_i) = d(E_k) = 1$ and $d(\sigma) \le 1$. In particular, if $d > 1$, then $L_G(E_i) = L_G(E_k) = L_G(\sigma) = 0$.*

(3) *If $\ell_T([[\tau]]) > 0$ and $d = 1$, then $L_G(E_i) = n_i \cdot \ell_T([[\tau]]), L_G(E_k) = n_k \cdot \ell_T([[\tau]])$ and $L_G(\sigma) = \ell_T([[\tau]]) \cdot |n_i - n_k|$.*

Proof. This is a direct consequence of the bounded cancellation lemma and the fact that for any $a \in F_n$ and $m > 0$, $\ell_T([[a^m]]) = m\ \ell_T([[a]])$. □

Lemma 4.29. *If σ has a splitting $\sigma = \sigma_1 \cdots \sigma_m$ into exceptional subpaths and single edges, then $\sigma = \sigma_1 \cdots \sigma_m$ is the canonical decomposition of σ and $L_G(\sigma) = L^*_G(\sigma)$.*

Proof. Lemma 4.26 implies that if τ is a maximal exceptional subpath of σ then $\tau = \sigma_j \dots \sigma_k$ for some $1 \le j \le k \le m$. Since exceptional paths have no nontrivial splittings, $j = k$ and $\tau = \sigma_j$. This proves that $\sigma = \sigma_1 \cdots \sigma_m$ is the canonical decomposition of σ. Lemma 4.20(2) implies that there is a uniform bound between $length_X([h(f^k(\sigma))])$ and $\sum length_X([h(f^k(\sigma_i))])$. The equality $L_G(\sigma) = L^*_G(\sigma)$ therefore follows from the assumption that $d \ge 1$. □

We say that an edge E_i is *below* the edge E_j, $i \ne j$, if the canonical decomposition of either $f(E_j)$ or $f(\bar{E}_j)$ contains a term that is not a Nielsen path and that has E_i as its initial edge. If E_i is not below any edge then it is *topmost*.

LEMMA 4.30. *$L_G(\sigma) = L^*_G(\sigma) = 0$ if σ is any of the following*:

- *a suffix*,
- *an edge that is not topmost*,
- *an edge that is fixed by f*,
- *a Nielsen path.*

Proof. Suppose that u is the suffix of an edge E. Remark 3.14 and Theorem 3.11(2) imply that u has a splitting $u = u_1 \cdots u_m$ into exceptional subpaths and single edges and that $f(E)$ has a splitting $f(E) = E \cdot u_1 \cdots u_m$. Lemma 4.29 and the fact that $L_G([f(E)]) = L_G(E)$ therefore imply the first item and that each $L_G(u_i) = L^*_G(u_i) = 0$. Since L_G and L^*_G agree on edges, to prove the second and third items we need only prove that $L_G(E) = 0$ when E is either not topmost or is fixed. The former case follows from Lemma 4.28 applied to a u_i as above and the latter from the assumption that $d \geq 1$.

If σ is a Nielsen path then $\sigma = [f^k(\sigma)]$ for all k so that σ has a splitting into exceptional subpaths and single edges. Lemma 4.29 implies that $L^*_G(\sigma) = L_G(\sigma)$. Since $d \geq 1$, $L_G(\sigma) = 0$. □

LEMMA 4.31. *Assume that $L^*_G(\alpha) = 0$.*

- *If the initial endpoint of σ equals the terminal endpoint of α then $L^*_G([\alpha\sigma]) = L^*_G(\sigma)$.*
- *If the terminal endpoint of σ equals the initial endpoint of α then $L^*_G([\sigma\alpha]) = L^*_G(\sigma)$.*

Proof. Since $L^*_G(\rho) = L^*_G(\bar{\rho})$ for all ρ, the two items are equivalent.

We prove the first item by induction on the number of terms in the canonical decomposition of α. Let $\sigma = \sigma_1 \dots \sigma_m$ be the canonical decomposition of σ. If α is a single edge there are four cases to consider. If $\alpha\sigma_1 \dots \sigma_m$ is the canonical decomposition of $[\alpha\sigma]$ then $L^*_G([\alpha\sigma]) = L^*_G(\sigma) + L_G(\alpha) = L^*_G(\sigma)$. If $\sigma_1 = \bar{\alpha}$ then $\sigma_2 \dots \sigma_m$ is the canonical decomposition of $[\alpha\sigma]$ and so $L^*_G([\alpha\sigma]) = L^*_G(\sigma) - L_G(\alpha) = L^*_G(\sigma)$. If $\alpha = \bar{E}_i$ and $\sigma_1 = E_i\tau^m\bar{E}_k$ is an exceptional path, then the canonical decomposition of $[\alpha\sigma]$ is the canonical decomposition of the Nielsen path τ^m followed by $\bar{E}_k$ followed by $\sigma_2 \dots \sigma_m$. By Lemma 4.28, $L_G(\sigma_1) = L_G(E_k) = 0$. By Lemma 4.30, $L^*_G(\tau^m) = 0$. Thus $L^*_G([\alpha\sigma]) = L^*_G(\sigma) - L_G(\sigma_1) + L^*_G(\tau^m) + L_G(E_k) = L^*_G(\sigma)$. The remaining case is that $\alpha = E_i$, that σ begins with $\tau^m\bar{E}_k$ and that $E_i\tau^m\bar{E}_k$ is an exceptional path. Then $L^*_G(\sigma) = L^*_G([\bar{\alpha}[\alpha\sigma]]) = L^*_G([\alpha\sigma])$ by the previous case. This completes the proof when α is a single edge.

If $\alpha = E_i\tau^m\bar{E}_k$ is exceptional, then there are three cases. By Lemma 4.28, $L_G(E_i) = L_G(E_k)$. By Lemma 4.30, $L^*_G(\tau^m) = 0$. If the initial edge of σ is not E_k then $L^*_G([\alpha\sigma]) = L^*_G(\alpha) + L^*_G(\sigma) = L^*_G(\sigma)$. If $\sigma_1 = E_k$ then $L^*_G([\alpha\sigma]) = L^*_G(\sigma) - L_G(E_k) + L^*_G(\tau^m) + L_G(E_i) = L^*_G(\sigma)$. Finally, if $\sigma_1 = E_k\tau^p\bar{E}_l$ is exceptional then the canonical decomposition of $[\alpha\sigma]$ is $(E_i\tau^{m+p}\bar{E}_l)\sigma_2\dots\sigma_m$ so $L^*_G([\alpha\sigma]) = L^*_G(\sigma) - L_G(E_i\tau^m\bar{E}_k) + L^*_G(E_i\tau^{m+p}\bar{E}_l) = L^*_G(\sigma)$. This completes the proof when there is only one term in the canonical decomposition of α.

In general, let α_1 be the first term in the canonical decomposition and write $\alpha = \alpha_1\alpha'$. Then $L^*_G([\alpha\sigma]) = L^*_G([\alpha_1[\alpha'\sigma]]) = L^*_G([\alpha'\sigma]) = L^*_G(\sigma)$ where the last equality is by induction. □

Proposition 4.32. *For any path $\sigma \subset G$, $L_G(\sigma) = L^*_G(\sigma)$.*

Proof. By Theorem 3.11(5), $L_G(\sigma) = L_G([f^k(\sigma)]) = L^*_G([f^k(\sigma)])$ for all sufficiently large k. It therefore suffices so show that L^*_G is $f_\#$-invariant.

We argue by induction on the height $h(\sigma)$ of σ. For $h(\sigma) = 0$ the statement is vacuously true. Assume that $height(\sigma) = h$ and that $L^*_G([f(\sigma)]) = L^*_G(\sigma)$ for paths σ with height at most $h-1$. If E_h is part of a linear family, denote this family $\{E_i : s \le i \le t\}$; otherwise let $s = h$. The edges E_i with $s \le i \le h$ are all highest edges so that there is a splitting (see Definition 3.15) of σ into subpaths of the following types: $\mu, E_i\mu, \mu\bar{E}_i$ and $E_i\mu\bar{E}_j$ where $s \le i, j \le h$ and where μ is contained in G_{s-1}. The canonical decomposition of σ is a refinement of this splitting and so it suffices to assume that σ has one of these forms.

The case that $\sigma = \mu$ follows by induction. Suppose that $\sigma = E_i\mu$. Since f is a homotopy equivalence and fixes all vertices, there is a path η_i of height at most $s-1$ such that $f_\#(\eta_i)$ is the suffix u_i of E_i. By Lemma 4.31 and the inductive hypothesis, $L^*_G(\eta_i) = L^*_G(u_i) = 0$. Since μ and $[u_i f(\mu)]$ are contained in G_{s-1}, E_i is the first term in the canonical decompositions of $\sigma = E_i\mu$ and $[f(\sigma)] = E_i[u_i f(\mu)]$. Thus $L^*_G(f_\#(\sigma)) = L^*_G(f_\#([[E_i\bar{\eta}_i][\eta_i\mu]])) = L^*_G(E_i f_\#([\eta_i\mu])) = L_G(E_i) + L^*_G(f_\#([\eta_i\mu])) = L_G(E_i) + L^*_G([\eta_i\mu]) = L_G(E_i) + L^*_G(\mu) = L^*_G(\sigma)$. The $\mu\bar{E}_i$ case follows by symmetry. If $\sigma = E_i\mu\bar{E}_j$ is exceptional, then so is $[f(\sigma)]$ and Lemma 4.28 completes the proof. Otherwise both E_i and $\bar{E}_j$ are terms in the canonical decomposition of both σ and $[f(\sigma)]$. In this case, one writes $\sigma = [[E_h\bar{\eta}_i][\eta_i\mu\bar{\eta}_j][\eta_j\bar{E}_h]]$ and argues as in the $\sigma = E_h\mu$ case. □

Our remaining task is to explicitly construct a tree whose length function agrees with L^*_G.

Lemma 4.33. *Let $\hat{G}$ be the subgraph consisting of edges E with $L_G(E) = 0$. If $d \ge 2$, then $T\mathcal{O}^\infty$ is the tree obtained from G by collapsing the components of $\hat{G}$ to points and assigning each remaining edge E the length $L_G(E)$.*

Proof. By Lemma 4.28 each exceptional subpath of σ is contained in $\hat{G}$. One may therefore compute $L^*(\sigma)$ by writing σ as a concatenation of edges and adding up the L_G-length of each edge. □

For the remainder of the section we assume that $d = 1$. To each edge E of G, assign a formal length $L_G(E)$. By reordering the edges of G we may assume that the edges in a linear family with positive formal length are consecutively numbered, say $\{E_i : s \leq i \leq t\}$, and that within such a linear family, formal lengths are increasing; equivalently, $n_j > n_i$ if $s \leq i < j \leq t$.

If E_i and E_j are distinct edges in a linear family then $length_G(E_j\bar{E}_i) > L^*(E_j\bar{E}_i)$ by Lemma 4.28. To create a graph G' that assigns the correct lengths to exceptional paths of positive formal length, we fold edges within linear families $\{E_i : s \leq i \leq t\}$ of positive formal length as follows. First fold part of $\bar{E}_t$ over all of $\bar{E}_{t-1}$. Denote the unfolded part of E_t by E'_t and assign it a formal length equal to $L_G(E_t) - L_G(E_{t-1})$. Next fold part of $\bar{E}_{t-1}$ over all of $\bar{E}_{t-2}$ denoting the unfolded part of E_{t-1} by E'_{t-1} and assign it a formal length equal to $L_G(E_{t-1}) - L_G(E_{t-2})$. Continue this down the linear family with the last fold defining E'_{s+1}. Repeat this for all linear families with positive formal length. An edge $E \subset G$ that is not subdivided in this process determines an edge in G' that we label E' and assign a formal length equal to $L_G(E)$. Let $g : G \to G'$ be the total folding map. If E_i is a member of a linear family with positive length $\{E_i : s \leq i \leq t\}$ then $g(E_i) = E'_iE'_{i-1}\dots E'_s$. We will not use f to induce a homotopy equivalence of G' but simply use G' as the basis for the construction of $T\mathcal{O}^\infty$.

For each linear edge $E_i \subset G$, let $\eta_i = E_i\tau_i\bar{E}_i$ be the *primitive Nielsen path determined by* E_i. Denote $[g(\eta_i)]$ by η'_i and let τ'_i be the subpath satisfying $\eta'_i = E'_i\tau'_i\bar{E}'_i$. If $g(E_i)$ is a single edge E'_i, then $\tau'_i = [g(\tau_i)]$. Otherwise E_i is an element of an exceptional family $\{E_i : s \leq i \leq t\}$ and $i > s$. In that case $\tau'_i = E'_{i-1}\dots E'_s\tau_i\bar{E}'_s\dots\bar{E}'_{i-1} = \eta'_{i-1}$.

Remark 4.34. By Theorem 3.11(5), a path $\sigma \subset G$ is a Nielsen path if and only if it is a concatenation of subpaths, each of which is either a fixed edge or $[\eta_i^k] = E_i\tau_i^k\bar{E}_i$ for some integer k.

LEMMA 4.35. *Suppose that σ is a path in G and that $\sigma = \sigma_1\dots\sigma_m$ is its canonical decomposition. Denote $[g(\sigma)]$ by σ' and $[g(\sigma_i)]$ by σ'_i. Then*

- *$\sigma' = \sigma'_1\dots\sigma'_m$ is a decomposition into subpaths.*
- *The maximal subpaths $\{\alpha'_k\}$ of σ' of the form $\eta_l'^m$ are disjoint.*
- *Each α'_k is contained in some σ'_j.*
- *For each j, the sum of the formal lengths of the edges of σ'_j that are not contained in any α'_k equals $L_G(\sigma_j)$.*

Proof. We begin by enumerating the possible values of σ'_j. If σ_j is a single edge E_i then $\sigma'_j = E'_i$ unless E_i is contained in a linear family $\{E_j : s \le j \le t\}$ and $i > s$. In that case, $\sigma'_j = E'_i E'_{i-1} \dots E'_s$. If $\sigma_j = E_i \tau_s^m \bar{E}_k$ where $s \le k \le i \le t$ then $\sigma'_j = E'_i E'_{i-1} \dots E'_{k+1} E'_k \dots E'_s \tau_i^m \bar{E}'_s \dots \bar{E}'_k = E'_i E'_{i-1} \dots E'_{k+1} [{\eta'_k}^m]$. Given these possibilities, the third item implies the second and fourth items.

If the first item fails then some $\bar{\sigma}'_j$ and σ'_{j+1} have a common initial edge. Examining the values above and keeping in mind that the initial edges of $\bar{\sigma}_j$ and σ_{j+1} are distinct, we see that the only possibility is that $\sigma_j = E_i$ and σ_{j+1} begins with $\bar{E}_s$ for some linear family $\{E_j : s \le j \le t\}$ with $s \le i \le t$. But then there is an exceptional subpath that overlaps with σ_j and σ_{j+1} in contradiction to the definition of canonical subpath. This proves the first item.

The homotopy equivalence g induces a bijection on paths with endpoints at vertices. Given a subpath $\mu' \subset \sigma'$ let $\mu \subset \sigma$ be the smallest path (not necessarily with endpoints at vertices) satisfying $[g(\mu)] = \mu'$. Suppose that $\mu' = [{\eta'_l}^m]$. If E_l is not part of a linear family then μ must have endpoints at vertices and so, by uniqueness, must be $[{\eta_l}^m]$. In particular, μ is contained in some σ_j and so μ' is contained in some σ'_j. If E_l is part of a linear family then it may be that one or both endpoints of μ are contained in the interior of edges in the same linear family as E_l. In this case the path ν obtained from μ by including the edges that contain the endpoints of μ has image of the form $\bar{E}'_j {\tau'_s}^m \bar{E}'_k$. By uniqueness ν is exceptional and we conclude as before that μ' is contained in a single σ'_j. This proves the third item and so also the second and fourth items. □

Definition 4.36. Let $\hat{G}'$ be the subgraph of G' consisting of edges with formal length zero and let T^0 be the tree determined from G' by collapsing the components of $\hat{G}'$ to points and making the length of each remaining edge equal to its formal length. For $i > 0$ we will inductively define T^i working our way up the strata of G. If E_i is not a linear edge with positive formal length, then $g(E_i) = E'_i$ is a single edge and we define $T^i = T^{i-1}$. Otherwise, define T^i to be the tree obtained from T^{i-1} by pulling τ'_i over E'_i [BF91, p. 452]. Remark 4.34 implies that τ'_i is elliptic in T^{i-1} and so this operation is well defined. Denote the tree obtained at the end of the process by T'.

Lemma 4.37. *If $d = 1$ then $T\mathcal{O}^\infty = T'$.*

Proof. Given a circuit σ contained in G, denote $[g(\sigma)]$ by σ' and let $\{\alpha'_k\}$ be as in Lemma 4.35. By Lemma 4.35, it suffices to show that $\ell_{T'}([[\sigma']])$ equals the sum of the formal lengths of the edges of σ' that are not contained in any α'_k. But this is exactly how T' was constructed. □

Remark 4.38. Nontrivial edge stabilizers of $T\mathcal{O}^\infty$ all arise from Definition 4.36. We record for future use the fact that each nontrivial edge stabilizer of T' is generated by a conjugate of a root of a linear suffix of f.

4.3. *Primitive subgroups.* From the homology, it is clear that if the elements of a free factor system are permuted under a unipotent outer automorphism, then each is invariant, and the restriction is unipotent. We will show moreover that a periodic free factor (or even a vertex stabilizer of a tree in $\mathcal{T}_{VS}$) is invariant. Our argument uses only that vertex stabilizers are primitive.

Recall that a subgroup H of a group J is primitive if, for all $a \in J$ and all $i \neq 0$, $a^i \in H$ implies that $a \in H$. An element a of J is *primitive* if $\langle a \rangle$ is a primitive subgroup of J.

Lemma 4.39. *Let H be a finitely generated nontrivial primitive subgroup of F_n. Then the normalizer $N(H)$ of H in F_n is H.*

Proof. Let T be a minimal free simplicial F_n-tree and let T_H be a minimal H-invariant subtree of T. Let $\gamma \in N(H)$. Then $\gamma(T_H) = T_H$ and so the axis of γ is in T_H and projects to a circuit in T_H/H. Thus, a power of γ is in H. Since H is primitive, γ is in H. □

Lemma 4.40. *Let $\mathcal{H}$ be a subgroup of* $\mathrm{Out}(F_n)$ *and let H be a finitely generated primitive subgroup of F_n whose conjugacy class is $\mathcal{H}$-invariant. Then, the restriction map $\rho_H : \mathcal{H} \to \mathrm{Out}(H)$ is well-defined. Further, if $\mathcal{H} \subset \mathrm{PG}(F_n)$, then $\mathcal{H}|_H := \rho_H(\mathcal{H}) \subset \mathrm{PG}(H)$.*

Proof. The first statement is an easy consequence of Lemma 4.39. The second statement follows from the fact that the inclusion of a finitely generated subgroup into F_n is a quasi-isometry, see for example [Sho91]. □

Proposition 4.41. *Suppose that $\mathcal{O} \in \mathrm{UPG}(F_n)$ and that $H \subseteq F_n$ is a primitive finitely generated subgroup. If $\mathcal{O}^k([[H]]) = [[H]]$ for some $k > 0$, then $\mathcal{O}([[H]]) = [[H]]$. Furthermore, if $\Phi \in Aut(F_n)$ is a lift of $\mathcal{O}$ with $\Phi^k(H) = H$ then $\Phi(H) = H$.*

The statement is false without the primitivity assumption as the following example shows: $F_2 = \langle e_1, e_2 \rangle$, $\Phi(e_1) = e_1$, $\Phi(e_2) = e_1 e_2$, $H = \langle e_1^2, e_2 \rangle$, $k = 2$.

Proof. By Proposition 3.16, we may assume that $\mathrm{rank}(H) > 1$. Let $f : G \to G$ be a UR for $\mathcal{O}$ such that f is linear on each edge of G. By $p : G' \to G$ denote the covering space of G corresponding to H. There is a lift $g : G' \to G'$ of f^k. We claim that, after replacing g by a power if necessary, g fixes a vertex. Indeed, by linear algebra, some power g^m of g will have negative Lefschetz number. Any fixed point v of negative index of g^m composed with the retraction to the core is fixed under g^m. The claim now follows from the observations that $p(v)$ is a fixed point for a positive power of f, that a positive power of a UR is also a UR, and that the only fixed points of a UR that is linear on edges are vertices and fixed edges.

We now use v and $p(v)$ as base points. Let σ be a closed path in G' based at v. The sequence $\{[g^i(\sigma)]\}$ forms a sequence of lifts of a subsequence of the sequence $\{[f^j(p(\sigma))]\}$. The latter is eventually a polynomial sequence (Lemma 4.2), and hence by Lemma 4.9 eventually the based closed path $[f^j(p(\sigma))]$ lifts to a based closed path in G'. Applying this to closed paths σ representing generators of $\pi_1(G', v)$ we conclude that f^j eventually lifts to G'. Thus, eventually $\mathcal{O}^j([[H]]) = [[H]]$, and the claim follows.

For the "furthermore" part of the proposition let v be the base point and choose g so that v is fixed. □

Definition 4.42. A circuit in a graph is *primitive* if it is not a proper power of another circuit.

LEMMA 4.43. *Let G be a finite connected core graph with oriented edges. Suppose that $f : G \to G$ is a cellular homeomorphism preserving orientations and inducing a nontrivial permutation of the edges. Then either*

(1) *there is a primitive circuit $\sigma = E_1E_2 \cdots E_m$ that is nontrivially rotated by f, i.e. there is a p, $0 < p < m$, such that $f(E_i) = E_{p+i}$ (subscripts are taken* mod m), *or*

(2) *there is an f-fixed vertex such that f nontrivially permutes the edges containing v.*

Proof. First suppose that $v, f(v), \dots, f^m(v) = v$ is a nontrivial orbit of vertices with m minimal. Choose an embedded path τ connecting v and $f(v)$. If $\tau f(\tau) \cdots f^m(\tau)$ is essential, then $[[\tau f(\tau) \cdots f^m(\tau)]]$ is a circuit as in (1). If not, then a vertex in τ is as in (2).

If f fixes the vertices of G, then the common vertex of a nontrivial orbit of edges is as in (2). □

PROPOSITION 4.44. *Suppose that $\mathcal{O} \in \mathrm{UPG}(F_n)$ and that $H \subseteq F_n$ is a primitive finitely generated subgroup whose conjugacy class is fixed by $\mathcal{O}$. Then the restriction* (*see Lemma* 4.40) *$\mathcal{O}|_H \in \mathrm{UPG}(H)$.*

Proof. It is obvious that $\mathcal{O}|_H \in \mathrm{PG}(H)$. We argue that the action on homology is unipotent. Let $f : G \to G$ be a UR for $\mathcal{O}$. By $p : G' \to G$ denote the covering space of G corresponding to H and let $f' : G' \to G'$ be a lift of f. By C denote the core of G'. Let $\rho : G' \to C$ be the nearest point retraction. If C does not contain any lifts of the highest edge $E \subset G$, then we may argue by induction on the number of strata. Therefore we assume that C contains lifts of E. From C form a finite graph $\hat{C}$ by collapsing all edges that do not project to E. The map $\rho f'$ induces a cellular isomorphism $\hat{f} : \hat{C} \to \hat{C}$. Since C is finite, $\hat{f}$ has finite order. The main step of the proof is to argue that $\hat{f} = \mathrm{id}$.

If $\hat{f}$ is not the identity, then, according to Lemma 4.43, there are two cases. Suppose first that there is a primitive circuit $\sigma = E_1 E_2 \cdots E_m$ that is nontrivially rotated by $\hat{f}$. Here each E_i is a lift of E or of $\overline{E}$, and, say, $\hat{f}(E_i) = E_{p+i}$, $0 < p < m$ (indices are taken mod m).

Choose a path in C of the form $E_1 \tau E_2$ where τ does not cross any lifts of E or $\overline{E}$. The sequence $\{[f'^k(\tau)]\}_{k=1}^{\infty}$ projects to an eventually polynomial sequence. Further, for infinitely many values of k (those in the same congruence class modulo the order of p in m) these paths have common initial and common terminal endpoints. It follows from Lemma 4.9 that for large k and any i there is a path that joins E_{ip+1} and E_{ip+2} and projects to the same path as $[f'^k(\tau)]$. Repeat this construction for every $\hat{f}$-orbit of consecutive edges in σ to obtain a primitive circuit in C that projects to a proper power. This contradicts the primitivity assumption and shows that $\hat{f} = id$ in this case.

The second case is that there is an $\hat{f}$-fixed vertex v of $\hat{C}$ such that $\hat{f}$ nontrivially permutes the edges containing v. Let $E_1, E_2, \dots, E_m$ be a minimal $\hat{f}$-orbit of edges with initial vertex v. Let $E'_1, E'_2, \dots, E'_m$ be the corresponding lifts to C with initial vertices $v'_1, v'_2, \dots, v'_m$. Note that these vertices are distinct because p is an immersion. We may now repeat the argument given in the first case with τ a path connecting v_1 and v_2 that does not cross a lift of E or $\overline{E}$, and ultimately reach the same contradiction.

If σ is any circuit in C representing a cycle, then $(f'_* - id)(\sigma)$ is a cycle supported in the cores of the components of

$$C \setminus (\cup\{\text{interiors of lifts of } E\}).$$

Inductively, it follows that a high power of $f'_* - id$ kills σ. Thus, $\mathcal{O}|_H \in \mathrm{UPG}(H)$. □

4.4. *Unipotent automorphisms and trees.* See Section 2 for definitions.

PROPOSITION 4.45. *Let* $\mathcal{O} \in \mathrm{UPG}(F_n)$ *and* $T \in \mathcal{T}_{VS}$. *Suppose that,* $\mathrm{Point}(T)$ *and* $\mathrm{Arc}(T)$ *are* $\mathcal{O}$*-invariant. Then,*

(1) *each element of* $\mathrm{Arc}(T)$ *is* $\mathcal{O}$*-invariant,*

(2) *each element of* $\mathrm{Point}(T)$ *is* $\mathcal{O}$*-invariant, and*

(3) *if* $[[V]] \in \mathrm{Point}(T)$, *then* $\mathcal{O}|_V \in \mathrm{UPG}(V)$.

Proof. By Theorem 2.3, $\mathrm{Arc}(T)$ is a finite set. Item (1) follows by Proposition 3.16.

Similarly, each of the finitely many elements of $\mathrm{Point}(T)$ is $\mathcal{O}$-periodic, and hence is $\mathcal{O}$-invariant by Proposition 4.41. The restriction of $\mathcal{O}$ is unipotent by Proposition 4.44. □

Remark 4.46. If $T \in \mathcal{T}_{VS}$ is fixed by $\mathcal{O} \in \mathrm{Out}(F_n)$, then the hypotheses of Proposition 4.45 are satisfied.

LEMMA 4.47. *Let $f : G \to G$ be a cellular homeomorphism of a connected finite graph. Suppose that f fixes all valence one vertices and induces a unipotent map of $H_1(G;\mathbb{Z})$. Then either f is homotopic rel vertices to the identity or there is a homeomorphism G to S^1 that conjugates f to a rotation.*

Proof. After a homotopy relative to vertices, we may assume that every point is periodic. If G is a circle, the claim is clear. We may therefore assume that either G has valence one vertices or negative Euler characteristic. In either case $\mathrm{Fix}(f) \neq \emptyset$. We will assume that $\mathrm{Fix}(f) \neq G$ and derive a contradiction.

After replacing f by an iterate if necessary, we may assume that all nonfixed points have the same period, say $k > 1$. Choose a shortest arc α in G that intersects $\mathrm{Fix}(f)$ exactly in its (possibly equal) endpoints $\partial\alpha$. Suppose that there exists $v \in \alpha \setminus \partial\alpha$ such that $f(v) \in \alpha$. Since v is not a fixed point, either v or $f(v)$ is not the midpoint of α. But then there is path in $\alpha \cup f(\alpha)$ that intersects $\mathrm{Fix}(f)$ exactly in its boundary and that is shorter than α. We conclude that $\alpha, f(\alpha), \dots f^{k-1}(\alpha)$ intersect only in their endpoints. It follows that there is a nonfixed periodic class under the action of f^k on the homology of the subgraph $\cup_{i=0}^{k-1} f^i(\alpha)$. This contradicts Corollary 3.3. □

PROPOSITION 4.48. *Assume $n > 1$. Suppose $\mathcal{O} \in \mathrm{UPG}(F_n)$ fixes $T \in \mathcal{T}$. Let $\Phi \in \mathrm{Aut}(F_n)$ be a lift of $\mathcal{O}$ and let $f_\Phi : T \to T$ be a Φ-equivariant isometry. Then f_Φ fixes all orbits of vertices and directions.*

Proof. The map f_Φ induces a periodic homeomorphism $\overline{f}_\Phi$ of the quotient graph $\overline{T} = T/F_n$. If $x \in T$ is a vertex such that $\mathrm{Fix}_T(\mathrm{Stab}_T(x)) = \{x\}$, then the image of x in $\overline{T}$ is f_Φ-fixed by Proposition 4.45. In particular, $\overline{f}_\Phi$ fixes all valence 1 vertices. Since the induced action on homology of $\overline{T}$ is unipotent, by Lemma 4.47, $\overline{f}_\Phi$ either is the identity or is conjugate to a rotation of the circle. The latter is impossible since $\mathrm{Elliptic}(T) \neq \{1\}$. □

LEMMA 4.49. *Let $\mathcal{O} \in \mathrm{UPG}(F_n)$, $S \in \mathcal{T}_{VS}$, and $a \in F_n$. Suppose $\ell_{S\mathcal{O}^\infty}(a) > 0$. Then there is a K such that $\ell_S(\mathcal{O}^{K'}(a)) > 0$ for all $K' > K$.*

Proof. This is immediate by the definition of limits. □

PROPOSITION 4.50. *Let $f : G \to G$ be a* UR *for $\mathcal{O} \in \mathrm{UPG}(F_n)$ and let $T \in \mathcal{T}_{VS}$ be $\mathcal{O}$-nongrowing. If $a \in F_n$ is elliptic in $T\mathcal{O}^\infty$, then a is elliptic in T.*

Proof. We first comment that, by the definition of nongrowing, if $x \in F_n$ is elliptic in $T\mathcal{O}^\infty$ then there is $k(x)$ such that $\mathcal{O}^k([[x]])$ is elliptic in T for $k > k(x)$.

Let $V = \langle a_1, \ldots, a_p \rangle$ be a point stabilizer of $T\mathcal{O}^\infty$ containing a. By the above comment, there is a k such that, for all pairwise products $a_i a_j$, we have $\mathcal{O}^k([[a_i a_j]])$ is elliptic in T. By Theorem 2.1, $\mathcal{O}^k([[V]])$ is elliptic in T.

Since V is a point stabilizer of the $\mathcal{O}$-fixed tree $T\mathcal{O}^\infty$, $\mathcal{O}([[V]]) = [[V]]$ by Theorem 2.3 and Proposition 4.41. So $[[V]]$, and hence a, is elliptic in T. □

5. A Kolchin theorem for unipotent automorphisms

The rest of the paper is devoted to the proof of our main theorem.

THEOREM 5.1. *For every finitely generated unipotent subgroup $\mathcal{H}$ of* Out(F_n), *there is a tree in $\mathcal{T}$ with all edge stabilizers trivial that is fixed by all elements of $\mathcal{H}$. Furthermore, there is such a tree with exactly one orbit of edges and if $\mathcal{F}$ is any maximal proper $\mathcal{H}$-invariant free factor system then* T *can be chosen so that $\mathcal{F}(T) = \mathcal{F}$.*

We fix notation as follows. Let $\mathcal{F}$ be a maximal $\mathcal{H}$-invariant proper free factor system. Choose a generating set

$$\mathcal{H} = \langle \mathcal{O}_1, \mathcal{O}_2, \ldots, \mathcal{O}_k \rangle$$

and for $1 \le l \le k$ let $f_l : G^l \to G^l$ be a UR for $\mathcal{O}_l$ with a filtration element $G^l_{r(l)}$ such that $\mathcal{F}(G^l_{r(l)}) = \mathcal{F}$. Choose $T_0 \in \mathcal{T}$ with trivial edge stabilizers and with $\mathcal{F}(T_0) = \mathcal{F}$.

Definition 5.2. The *bouncing sequence* associated with the above choices is the sequence of simplicial trees

$$\{T_0, T_1, T_2, \cdots\}$$

in $\mathcal{T}$ defined by

$$T_i = T_{i-1}\mathcal{O}_i^\infty$$

where subscripts of the $\mathcal{O}_i$'s are taken mod k (see Theorem 4.22).

We prove Theorem 5.1 by showing that the bouncing sequence is eventually constant and that the limit tree satisfies the conclusions of the theorem. As intermediate steps, we show that T_i is $\mathcal{O}_{i+1}$ nongrowing for all sufficiently large i (Proposition 5.7) and that the edge stabilizers of T_i are trivial for all sufficiently large i (Lemma 5.10). It turns out that the bouncing sequence is constant. This is reflected, for example, in Lemma 5.10 where having established properties of T_i for large i, the properties then hold for all i.

Before starting the proof, we give some examples of what can happen if $\mathcal{F}$ is not assumed to be maximal.

Example 5.3. Recall that $F_2 = \langle e_1, e_2\rangle$. Set $\mathcal{H} = \langle \mathcal{O}\rangle$ with $\mathcal{O}$ represented by the automorphism Φ given by $\Phi(e_1) = e_1$, $\Phi(e_2) = e_2e_1$, and let $T_0 \in \mathcal{T}_{CV}$. Then, $T_0\mathcal{O}^\infty$ is the tree T from Example 2.7. The tree T is $\mathcal{O}$-fixed, and so the bouncing sequence is eventually constant. However, T has nontrivial edge stabilizers, and in this case the iteration scheme fails to discover a tree as in the conclusion of Theorem 5.1. Here $\mathcal{F}$ is trivial. We can discover a larger invariant proper free factor system, namely $\{[[\langle e_1\rangle]]\}$, by looking at the edge stabilizer.

Example 5.4. Recall $F_3 = \langle e_1, e_2, e_3\rangle$. Set $\mathcal{H} = \langle \mathcal{O}_1, \mathcal{O}_2\rangle$, where $\mathcal{O}_i$ is represented by the automorphism Φ_i given by $\Phi_1(e_1) = e_1$, $\Phi_1(e_2) = e_2e_1$, $\Phi_1(e_3) = e_3$, $\Phi_2(e_1) = e_1$, $\Phi_2(e_2) = e_2$, and $\Phi_2(e_3) = e_2^{-1}e_1e_2e_3$. Notice that the basis $\{e_1, e_2, e_2e_3\}$ is better adapted to Φ_2 since $\Phi_2(e_2e_3) = e_1e_2e_3$. Consider the trees in $\mathcal{T}$ whose models are pictured below.

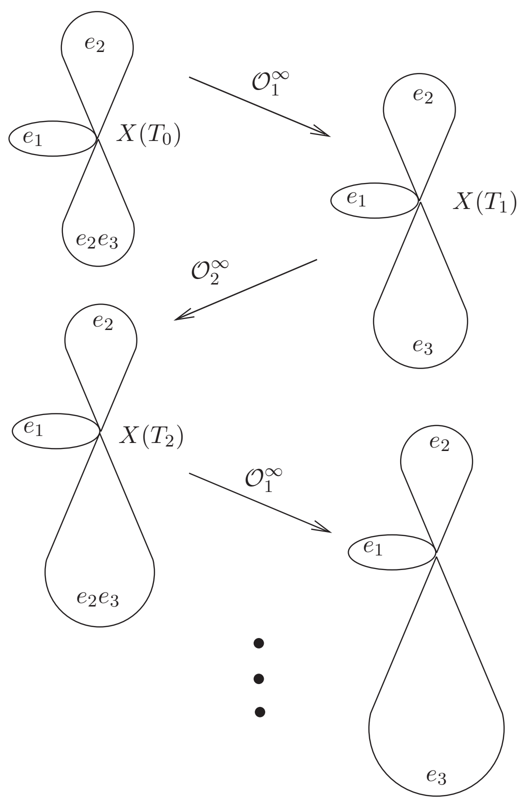

In the illustration, group elements indicate the image of the markings, and the only vertical edges are the ones corresponding to e_1. All edge stabilizers are trivial and in each case the indicated tree has infinite cyclic vertex stabilizers (always conjugate to $\langle e_1\rangle$). The tree T_1 is the limit $T_0\mathcal{O}_1^\infty$, T_2 is the limit $T_1\mathcal{O}_2^\infty$, etc. The tree T_2 is combinatorially isomorphic to T_0, i.e. T_0 and T_2 belong to the same open cone of $\mathcal{T}$. Notice, however, that T_0 and T_2 are not homothetic: the ratio $\ell(e_2)/\ell(e_2e_3)$ is smaller in T_2 than in T_0. The bouncing sequence indeed *bounces* between two open cones in $\mathcal{T}$, so it does not stabilize. All trees in the sequence are nongrowers under all elements of $\mathcal{H}$. The ratios $\ell_{T_i}(e_2)/\ell_{T_i}(e_2e_3)$ converge to 0. Here $\mathcal{F}$ is $\{[[\langle e_1\rangle]]\}$. We see that the edge marked e_2 gets relatively shorter and shorter in the bouncing sequence. This tells us how to enlarge $\mathcal{F}$ to a larger $\mathcal{H}$-invariant free factor system, namely $\{[[\langle e_1, e_2\rangle]]\}$.

5.1. *$\mathcal{F}$ contains the suffixes of all nonlinear edges.*

PROPOSITION 5.5. *If E is a nonlinear edge of G^l with suffix u, then u is contained in $G^l_{r(l)}$.*

Proof. To simplify notation we write f, G and G_r for f_l, G^l and $G^l_{r(l)}$.

Suppose that u is not contained in G_r. If E is crossed by the suffix u' of an edge E' then E' is not a linear edge and u' is not contained in G_r. We may therefore assume that E is the edge in the highest stratum. Since G_r is f-invariant and f fixes all vertices, $[f^k(u)]$ is not contained in G_r for any k. Thus, the eigenray $R = E \cdot u \cdot [f(u)] \cdot [f^2(u)] \cdot \ \ldots$ is not carried by $\mathcal{F}$.

The edge E determines a splitting of F_n as either a free product or an HNN-extension. Let $\mathcal{F}_E$ denote the resulting free factor system. Since E is not an edge of G_r, $\mathcal{F} \sqsubset \mathcal{F}_E$. Also, $\mathcal{F} \neq \mathcal{F}_E$ since $\mathcal{F}_E$ carries R, but $\mathcal{F}$ does not.

The argument breaks up into two cases.

Case 1. $\mathcal{F}_E$ carries $\mathcal{H}R$. In this case, the smallest free factor system $\mathcal{F}'$ carrying $\mathcal{F}$ and $\mathcal{H}R$ is proper (since $\mathcal{F}' \sqsubset \mathcal{F}_E$), is $\mathcal{H}$-invariant (since both $\mathcal{F}$ and $\mathcal{H}R$ are), and satisfies $\mathcal{F} \sqsubset \mathcal{F}'$ properly (since R is not carried by $\mathcal{F}$). This contradicts the choice of $\mathcal{F}$.

Case 2. $\mathcal{F}_E$ does not carry $\mathcal{H}R$. In this case, we will show that $\mathcal{H}$ contains an element of exponential growth. There is an element of $\mathcal{H}$ such that, when represented as a homotopy equivalence $g : G \to G$, $[g(R)]$ crosses infinitely many E's. The idea is that the image of a path containing E's under a high power of f contains long initial subpaths of R, and the image under g of a path with long initial subpaths of R contains lots of E's. This feedback gives rise to exponential growth. We now make this more precise.

Let R^* denote an initial subpath of R with the property that $[g(R^*)] = STU$ is the concatenation of three paths such that the lengths of S and U are at least $\mathrm{BCC}(g)$ and such that the number of times that T crosses E and $\overline{E}$ is at least five. Let M be the length of $[g(R^*)]$. Choose N so that, for all paths τ starting with E of length no more than M, $[f^N(\tau)]$ starts with ER^* (see Proposition 3.18). We claim that the element of $\mathcal{H}$ represented by

gf^N has exponential growth. Indeed, since F_n and the universal cover of G are quasi-isometric, it is enough to find a circuit σ in G such that the length of $[[(gf^N)^i(g(\sigma))]]$ grows exponentially in i. We show that σ can be taken to be any circuit containing ER^* as a subpath. If this is the case, then $[[g(\sigma)]]$ contains $[g(R^*)]$ except that perhaps initial and terminal subpaths of length less than $\mathrm{BCC}(g)$ may have been lost. In particular, $[[g(\sigma)]]$ crosses five E's and $\overline{E}$'s separated by a distance of no more than M. It follows that the highest edge splitting of $[[g(\sigma)]]$ induced by initial vertices of the E's (and terminal vertices of the $\overline{E}$'s) contains at least two subpaths crossing E's or $\overline{E}$'s and with length at most M. By our choice of N, $[[f^N g(\sigma)]]$ contains two disjoint subpaths of the form ER^* or its inverse. So, $[[gf^N g(\sigma)]]$ contains two disjoint copies of $[g(R^*)]$ or its inverse except for a loss of initial and terminal subpaths of length less than $\mathrm{BCC}(g)$ and so contains at least two disjoint subpaths each crossing five E's or $\overline{E}$'s that are separated by a distance of no more than M. This pattern continues and the number of such subpaths containing three E's or $\overline{E}$'s at least doubles with each application of gf^N. □

Corollary 5.6. (1) *If $i = l \mod k$ then T_{i-1} is $\mathcal{O}_i$-growing if and only if there is a linear edge E in G^l whose suffix u has positive length in T_{i-1}; in this case the growth is linear.*

(2) *If $\sigma \subset G_l$ is a circuit and $l_{T_i}(\sigma) > 0$ then there is a suffix u as in (1) such that for every $N > 0$, $[[f^k(\sigma)]]$ contains $[f^N(u)]$ as a subpath for all sufficiently large k.*

Proof. Proposition 5.5 implies that for any circuit σ contained in G^l there exists a constant K such that each $[[f_l^k(\sigma)]]$ has a decomposition into at most K subpaths, each of which is either a single edge, a path contained in $G^l_{r(l)}$, or of the form u_j^m for some fixed suffix u_j of f_l. Up to a uniform bound, the only terms that contribute to $\ell_{Ti-1}([[f_l^k(\sigma)]])$ are those of the form u_j^m and these contribute $m \cdot \ell_{Ti-1}(u_j)$. This proves (1).

For (2), we use the notation of Definition 4.27. There is no loss in assuming that the canonical decomposition of σ is a splitting. Proposition 5.5 implies that $d(E) = 0$ for all nonlinear edges. Lemma 4.29 then implies that at least one of the terms in the canonical decomposition of σ is a linear edge of positive length or an exceptional path of positive length. In either case, (2) follows. □

5.2. *Bouncing sequences stop growing.*

Proposition 5.7. *T_i is eventually $\mathcal{O}_{i+1}$-nongrowing.*

Proof. To simplify notation we assume that $i = 0 \mod k$. Let $\mathcal{U}$ be the (finite) set of suffixes of f_1 that are fixed by f_1. Set $K = |\mathcal{U}|$. We will show that there are at most K values of i such that T_i is $\mathcal{O}_1$-growing.

Suppose to the contrary that $T_{i_0}, T_{i_1}, \cdots, T_{i_K}$ are $\mathcal{O}_1$-growing with $i_0 < i_1 < \cdots < i_K$ and each $i_l = 0 \bmod k$.

SUBLEMMA 5.8. *There are* $w_i \in \mathcal{H}$, $1 \le i \le K$, *and* $u_i \in \mathcal{U}$, $0 \le i \le K$ *such that, given* $M \ge 0$, $[[f_1^B(w_i(u_i))]]$ *contains* $[u_{i-1}^M]$ *as a subpath for all large* B.

Proof of Sublemma 5.8. By Corollary 5.6, there is $u_K \in \mathcal{U}$ such that $\ell_{T_{i_K}}(u_K) > 0$. Since $T_{i_K} = T_{i_{K-1}+1}\mathcal{O}_2^\infty \cdots \mathcal{O}_{i_K - i_{K-1}}^\infty$, we may approximate T_{i_K} by $T_{i_{K-1}+1}w_K$ where $w_K = \mathcal{O}_2^{N_2} \cdots \mathcal{O}_{i_K-i_{K-1}}^{N_{i_K-i_{K-1}}}$ for suitably chosen N_j. In particular, we may assume that $\ell_{T_{i_{K-1}+1}w_K}(u_K) > 0$. In other words, $\ell_{T_{i_{K-1}+1}}(w_K(u_K)) > 0$. Corollary 5.6 then provides a suffix u_{K-1}, such that $\ell_{T_{i_{K-1}}}(u_{K-1}) > 0$ and such that $[[f_1^B(w_K(u_K))]]$ contains $[u_{K-1}^M]$ as a subpath for all large B. The argument may be repeated starting with u_{K-1}, etc. The sublemma follows. □

We now continue with the proof of Proposition 5.7. Two of the u_i's produced in the sublemma are equal, say $u_0 = u_K$. We shall show that there exists an element in $\mathcal{H}$ of exponential growth, a contradiction that will establish the proposition.

Let C be as in Lemma 3.17 for the UR f_1, and let A be such that the length in G of $[w_i(u_i)^A]$ is larger than twice BCC(w_i) (with w_i realized as homotopy equivalence on G). Choose B so that the circuit $[[f^B w_i(u_i)]]$ contains u_{i-1}^{C+2+A} as a subpath.

We claim that $\mathcal{O}_1^B w_1 \cdots \mathcal{O}_1^B w_K$ has exponential growth. Indeed, we will show that if σ is any path in G containing L occurrences of u_i^{C+2+A} whose interiors are disjoint, then $[f_1^B w_i(\sigma)]$ contains $2L$ occurrences of u_{i-1}^{C+2+A} with disjoint interiors. After all, when we apply w_i to σ, we obtain for each occurrence of u_i^{C+2+A} an occurrence of $[w_i(u_i^{C+2})]$ (at most $[w_i(u_i^A)]$ is lost by our choice of A). After applying f_1^B, by Lemma 3.17 we see L occurrences of $[f_1^B w_i(u_i^2)]$ with disjoint interiors. Finally, by our choice of B, each $[f_1^B w_i(u_i^2)]$ contains two copies of u_{i-1}^{C+2+A} with disjoint interiors. □

5.3. *Bouncing sequences never grow.* Recall (Section 2.2) that, for a simplicial tree T, Arc(T) denotes the set of conjugacy classes of stabilizers of nondegenerate arcs of T.

LEMMA 5.9. *Suppose that* $T \in \mathcal{T}$, *that* $\mathcal{O} \in$ UPG(F_n), *and that* T *is* $\mathcal{O}$*-nongrowing. Set* $T' = T\mathcal{O}^\infty$. *Then*, Arc(T') $\subset$ Arc(T). *Further, elements of* Arc(T') *are* $\mathcal{O}$*-invariant.*

Proof. Let $[[\langle e \rangle]] \in$ Arc(T'). We claim that there is an arc $[x, y]$ in T' and elements $a, b \in F_n$ such that

- $[x, y] \cap \mathrm{Fix}_{T'}(e)$ is nondegenerate,
- a and b are elliptic in T',
- $\mathrm{Fix}_{T'}(a) \cap [x, y] = \{x\}$, and
- $\mathrm{Fix}_{T'}(b) \cap [x, y] = \{y\}$.

Indeed, since the axes of hyperbolic elements cover T' (Theorem 2.2) there is an element $c \in F_n$ such that $\mathrm{Axis}_{T'}(c)$ has nondegenerate overlap with the arc $\mathrm{Fix}_{T'}(e)$. If we choose m large enough so that $c^m([x, y])$ is disjoint from $[x, y]$ then we may take a to be $c^{-m}ec^m$, b to be $c^m ec^{-m}$, x to be the point in $\mathrm{Fix}_{T'}(a)$ minimizing distance to $\mathrm{Fix}_{T'}(e)$, and y the point in $\mathrm{Fix}_{T'}(b)$ minimizing the distance to $\mathrm{Fix}_{T'}(e)$. This establishes the claim.

Keeping in mind that $\ell_{T'}(ab)$ is twice the distance in T' between $\mathrm{Fix}_{T'}(a)$ and $\mathrm{Fix}_{T'}(b)$, we see that $\ell_{T'}(ab) > \ell_{T'}(ae)+\ell_{T'}(be)$. Since T is $\mathcal{O}$-nongrowing, eventually we have $\ell_{T'}(ab) = \ell_T(\Phi^i(ab))$, etc. where $\Phi \in \mathrm{Aut}(F_n)$ represents $\mathcal{O}$. Therefore, eventually $\Phi^i(a)$, $\Phi^i(b)$, and $\Phi^i(e)$ are elliptics in T, and $\ell_T(\Phi^i(ab)) > \ell_T(\Phi^i(ae))+\ell_T(\Phi^i(be))$. Hence, $\Phi^i(e)$ is eventually a stabilizer of a nondegenerate arc of T. Since by Theorem 2.3 there are only finitely many conjugacy classes of arc stabilizers in T, it follows that the sequence $\{[[\Phi^i(e)]]\}$ takes only finitely many values, and is therefore constant by Proposition 3.16. The lemma follows. □

LEMMA 5.10. *For all $1 \le l \le k$ and for all $i \ge 0$,*

(1) *$\mathcal{F}$ contains the f_l-suffix of every edge in G^l.*

(2) *T_i is an $\mathcal{O}_{i+1}$ nongrower.*

(3) *The edge stabilizers of T_i are trivial.*

(4) *$\mathcal{F}(T_i) = \mathcal{F}$.*

Proof. The main step in the proof is to show that (2), (3) and (4) hold eventually, which is to say, for all sufficiently large i. By Proposition 5.7 there is a largest $s \ge 0$ such that T_s is an $\mathcal{O}_{s+1}$ grower. By Remark 4.38 all nontrivial arc stabilizers of T_{s+1} are generated by conjugates of roots of linear suffixes of f_s. Lemma 5.9 implies that if $\langle e\rangle$ is a nontrivial edge stabilizer of T_i for $i \ge s$ then e is conjugate to a root of a linear suffix of f_s. Lemma 5.9 also implies that the sequence $\{\mathrm{Arc}(T_i)\}$ is eventually constant and that if $\mathrm{Arc}(T_i)$ is not eventually trivial then there is linear suffix u of f_s such that $[u]$ is $\mathcal{H}$-invariant. The free factor system given by the highest edge of G^s contains both $\mathcal{F}$ and $[u]$. Therefore, the smallest free factor system that contains both $\mathcal{F}$ and $[u]$ is proper, and it is also $\mathcal{H}$-invariant (since both $\mathcal{F}$ and $[u]$ are), and it properly contains $\mathcal{F}$ (since it carries $[u]$, while $\mathcal{F}$ does not). This contradicts the choice of $\mathcal{F}$ and shows that edge stabilizers of T_i are eventually trivial.

By Proposition 4.50, $\{\mathrm{Elliptic}(T_i)\}$ eventually forms a nonincreasing sequence. Since the edge stabilizers of T_i are eventually trivial, the collection of nontrivial vertex stabilizers of T_i may be recovered from $\mathrm{Elliptic}(T_i)$ as the collection of maximal subgroups of F_n in the set $\mathrm{Elliptic}(T_i)$. So, eventually the sequence $\{\mathcal{F}(T_i)\}$ is a decreasing sequence of free factor systems. By Lemma 2.10, this sequence is eventually constant, hence eventually $\mathcal{H}$-invariant and hence eventually $\mathcal{F}(T_i) = \mathcal{F}$.

Having established (2) and (4) for large i, Corollary 5.6 implies (1) and then (2) for all i. Lemma 5.9 and induction then imply (3) for all i. The proof given above that (4) holds for large i now shows that (4) holds for all i. □

5.4. *Finding Nielsen pairs.*

Definition 5.11. Let $T \in \mathcal{T}$ have trivial edge stabilizers, and let $\mathcal{H}$ be a unipotent subgroup of $\mathrm{Out}(F_n)$. Assume that conjugacy classes of vertex stabilizers of T are $\mathcal{O}$-invariant for all $\mathcal{O} \in \mathcal{H}$. We say that distinct nontrivial vertex stabilizers V and W of T form a *Nielsen pair* for $\mathcal{H}$ if, for all $\mathcal{O} \in \mathcal{H}$ and all lifts Φ of $\mathcal{O}$ to $\mathrm{Aut}(F_n)$ there exists $a \in F_n$ such that $\Phi(V) = V^a$ and $\Phi(W) = W^a$. (It suffices to check this for one lift.)

Here is an alternative description. Let $\tilde{\mathcal{V}}^{(2)}$ denote the set of unordered pairs of distinct nontrivial vertex stabilizers of T. There is a natural diagonal action of $\mathrm{Aut}(F_n)$ on $\tilde{\mathcal{V}}^{(2)}$. This action descends to an action of $\mathrm{Out}(F_n)$ on $\mathcal{V}^{(2)} := \tilde{\mathcal{V}}^{(2)}/\mathrm{Inner}(F_n)$. If V and W are distinct nontrivial vertex stabilizers, then the corresponding element of $\mathcal{V}^{(2)}$ is denoted $[[V, W]]$. The pair V, W is a Nielsen pair for $\mathcal{H}$ if $[[V, W]]$ is a fixed point for the action of $\mathcal{H}$ on $\mathcal{V}^{(2)}$.

For example, if T is fixed by $\mathcal{H}$ and V, W are nontrivial stabilizers of neighboring vertices, then V and W form a Nielsen pair. The following Lemma 5.12 is an immediate consequence of the definition.

LEMMA 5.12. *Let T and $\mathcal{H}$ be as in Definition* 5.11.

- *If T' is another simplicial F_n-tree that has the same vertex stabilizers as T, then two vertex stabilizers V and W form a Nielsen pair in T if and only if they form a Nielsen pair in T'.*
- *If $\mathcal{H} = \langle \mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_k \rangle$ and two vertex stabilizers V and W of T form a Nielsen pair for $\langle \mathcal{O}_i \rangle$ for all i, then they form a Nielsen pair for $\mathcal{H}$.* □

PROPOSITION 5.13. *Let $\mathcal{H} = \langle \mathcal{O}_1, \dots, \mathcal{O}_k \rangle$ be a unipotent subgroup of $\mathrm{Out}(F_n)$, and let $T \in \mathcal{T}$ be such that*

- *T has trivial edge stabilizers,*

- *$\mathcal{F}(T) = \mathcal{F}$ where $\mathcal{F}$ is maximal and $\mathcal{H}$-invariant,*
- *T is $\mathcal{O}_i$-nongrowing for all i.*

Then T contains a Nielsen pair for $\mathcal{H}$.

The rest of this section is devoted to the proof of Proposition 5.13, which appears ahead, after some preparation. By $h_i : G^i \to G^i$ denote a UR for $\mathcal{O}_i$ such that

(1) $\mathcal{F} = \mathcal{F}(G^i_{r(i)})$ for some filtration element $G^i_{r(i)}$, and

(2) if E is an edge outside $G^i_{r(i)}$, then $h_i(E) = Eu$ for some closed path u (depending on i and E) in $G^i_{r(i)}$.

Such a representative exists. Indeed, by Theorem 3.11 there is a UR h_i satisfying everything except perhaps the condition that u is contained in $G^i_{r(i)}$. It follows from Corollary 5.6 applied to $\mathcal{H} = \langle \mathcal{O}_i \rangle$, $T_0 = T$, and $f_i = h_i$ that any h_i satisfies this last condition as well.

Using Lemma 5.12, we shall detect that two vertex stabilizers V and W of T form a Nielsen pair for $\mathcal{H}$ by examining for every i whether they form a Nielsen pair for $\langle \mathcal{O}_i \rangle$ in the tree S_i obtained from the universal cover of G^i by collapsing all edges that project to $G^i_{r(i)}$.

Edge paths P in G^i are of the form $\nu_0 H_1 \nu_1 H_2 \cdots H_p \nu_p$ where each H_j is an edge not in $G^i_{r(i)}$ and each ν_j is a path in $G^i_{r(i)}$. We call the subpaths ν_j *vertical elements* and the letter H is chosen for *horizontal* (with G^i as a model for the tree obtained from the universal cover of G^i by collapsing edges that project to $G^i_{r(i)}$). Some of the ν_j's could be trivial paths. When P is such a path, then the iterates $h_i^N(P)$ have a similar form $\nu_0^{(N)} H_1 \nu_1^{(N)} H_2 \cdots H_p \nu_p^{(N)}$. For each j, the sequence $\{\nu_j^{(N)}\}$ is seen to be eventually polynomial by application of Lemmas 4.2 and 4.7 to the pieces of the splitting of $h_i^N(P)$ at the endpoints of H_k where suffixes do not develop. We say that a vertical element ν_j is *inactive* if $\nu_j^{(N)}$ is independent of N. Otherwise, ν_j is *active*. Of course, h_i and the edge path P are implicit in these definitions. Even trivial ν_j's could be active. It follows from Lemma 4.13 applied to these same pieces that 'inactive' is equivalent to 'eventually inactive'.

When $i \neq j$ there is a homotopy equivalence $\phi_{ij} : G^i \to G^j$ given by markings. We may assume that this map sends vertices to vertices and restricts to a homotopy equivalence $G^i_{r(i)} \to G^j_{r(j)}$. Let C be a constant larger than the BCC of any lift of ϕ_{ij} to universal covers. Let ν be a vertical element in a path P in G^i. We can transfer P to another G^j using ϕ_{ij} and tightening. The path $[\phi_{ij}(\nu)]$ has length bounded above and below by a linear function in the length of ν, and then at most $2C$ is added or subtracted. In particular, if the length of

ν is larger than some constant $C_0 > 2C$, then ν *induces* a well-defined vertical element in G^j. Short ν's can disappear and new short vertical elements can appear in $[\phi_{ij}(P)]$.

Choose constants $C_1 \leq C_2 \leq \cdots \leq C_{7k}$ such that if a vertex element ν has length $\leq C_i$ $(0 \leq i < 7k)$ and is transferred to some other graph, then the induced vertex element has length $\leq C_{i+1}$. Also, fix $\epsilon \in (0, 1/14k)$.

LEMMA 5.14. *For a sufficiently large integer $m > 0$, the following statements hold.*

(1) *Let $N_i = 2^{2^{(7k-i+1)m}}$, and let $I_{i,l}$ be the interval*
$$\left[(1 - l\epsilon)N_i, (1 + l\epsilon)N_i^m\right]$$
for $i = 1, 2, \ldots, 7k$, $l = 1, 2, \ldots, 14k$. Then $I_{i,1} \subset I_{i,2} \subset \cdots \subset I_{i,14k}$ and the intervals $I_{i,14k}$ are pairwise disjoint for $i = 1, 2, \ldots, 7k$. Furthermore, the intervals $I_{i,14K}$ are disjoint from $[0, C_{7k}]$.

(2) *If a vertex element ν in an edge path P in G^i is active and has length $\leq (1 + 14k\epsilon)N_{i+1}^m$ (which is the right-hand endpoint of $I_{i+1,14k}$), then the h_i-iterated vertex element $\nu^{(N_i)}$ has length in $I_{i,1}$.*

(3) *If a vertex element ν in an edge path P in G^i has length in $I_{p,l}$ $(l < 14k)$, then, after transferring to G^j, ν induces a vertex element whose length belongs to $I_{p,l+1}$.*

(4) *If a vertex element ν in an edge path P in G^i has length in $I_{j,l}$ and if $i > j$ and $l < 14k$, then the iterated vertex element $\nu^{(N_i)}$ in $h_i^{N_i}(P)$ has length in $I_{j,l+1}$.*

We think of the first index in intervals $I_{i,l}$ as measuring the order of magnitude of lengths of vertex elements. The second index is present only for technical reasons: there is a slight loss when transferring from one graph to another (3), and when applying "lower magnitude maps" (4).

Proof of Lemma 5.14. To see that the right-hand endpoint of $I_{i+1,14k}$ is to the left of the left-hand endpoint of $I_{i,14k}$ we have to show that
$$(1 + 14k\epsilon)2^{2^{(7k-i)m}+m} < (1 - 14k\epsilon)2^{2^{(7k-i+1)m}}$$
i.e. that
$$2^{[2^{(7k-i+1)m} - 2^{(7k-i)m} - m]} > \frac{1 + 14k\epsilon}{1 - 14k\epsilon}.$$
That the latter inequality holds for large m follows from the observation that the exponent of the left-hand side
$$2^{(7k-i)m}(2^m - 1) - m$$
goes to infinity as $m \to \infty$.

It follows from Theorem 3.11(5) and (6) that there are polynomials Q_i and R_i with nonnegative coefficients such that whenever ν is an active vertex element in a path P in G^i, then the length of $\nu^{(N)}$ is in the interval $[N - R_i(\ell_{G^i}(\nu)), (1 + \ell_{G^i}(\nu))Q_i(N)]$. The proof now reduces to the fact that exponential functions grow faster than polynomial functions. For example, (2) follows from the inequalities

$$N_i - R_i((1 + 14k\epsilon)N_{i+1}^m) > (1 - 14k\epsilon)N_i$$

and

$$(1 + (1 + 14k\epsilon))N_{i+1}^m Q_i(N_i) < (1 + 14k\epsilon)N_i^m.$$

If we assume without loss of generality that $R_i(x) = x^d$ then the first inequality simplifies to

$$\frac{N_i}{N_{i+1}^{m+d}} > \frac{(1 + 14k\epsilon)^d}{14k\epsilon}.$$

Again, the left-hand side amounts to $2^{\exp}$ with

$$exp = 2^{(7k-i)m}(2^m - m - d)$$

and it goes to infinity as $m \to \infty$. The proof of the second inequality and of the other claims in the lemma are similar. (For (3) use the fact that there is a linear function L such that if σ is a vertex element of a path P' induced by a vertex element ν of a path P, then the length of σ is bounded by $L(\mathit{length}(\nu))$.)

□

We will argue that if there are no $\mathcal{H}$-Nielsen pairs in T, then the element $\mathcal{O}_{7k}^{N_{7k}} \cdots \mathcal{O}_2^{N_2}\mathcal{O}_1^{N_1} \in \mathcal{H}$ has exponential growth.

Start with a circuit P_1 in G^1 that is not contained in $G^1_{r(1)}$ and all of whose vertex elements have length $\leq C_1$. This circuit is the *first generation*. Then apply $h_1^{N_1}$ to obtain $h_1^{N_1}(P_1)$ and transfer this new circuit via ϕ_{12} to G^2. The resulting circuit P_2 is the *second generation*. Then apply $h_2^{N_2}$ and transfer to G^3 to obtain the *third generation* circuit P_3, etc. The circuit P_{7k} whose generation is $7k$ lives in G^{7k}. Then repeat this process cyclically: apply $h_{7k}^{N_{7k}}$ and transfer to G^1 to get a circuit P_{7k+1} of $(7k+1)^{\text{st}}$ generation etc.

Suppose that ν is a vertex element of some P_i. If $\nu^{(N_i)}$ has length $\geq C_0$, then $\nu^{(N_i)}$ induces a well-defined vertex element ν' in P_{i+1}. We say that ν *gives rise* to ν'.

We will now label some of the vertex elements of the P_i's with positive integers. Consider maximal (finite or infinite) chains $\nu_1, \nu_2, \cdots$ of vertex elements such that ν_i gives rise to ν_{i+1}. In particular, there is an integer s such that ν_i is a vertex element of P_{i+s} for $i \geq 1$. If the length of the chain is $\geq 7k$, then label ν_i by the integer i. If the chain has $< 7k$ vertex elements, we will leave all of them unlabeled. All labels > 1 in P_i correspond to unique labels in P_{i-1}. A *birth* is the introduction of label 1. A *death* is an occurrence of a labeled vertex element that does not give rise to a vertex element in the next generation. Any labeled vertex element can be traced backwards to its birth. Traced forward, any labeled vertex element either eventually dies, or lives forever (and the corresponding label goes to infinity).

LEMMA 5.15. *If a vertex element ν in some P_i is not labeled, then ν is h_i-inactive and its length is $\leq C_{7k}$.*

Proof. The first element ν_1 of a maximal chain $\nu_1, \nu_2, \ldots, \nu_s$, $s < 7k$, must have length $\leq C_1$. Indeed, assume not. Say ν_1 is a vertex element in P_{i+1}. By the choice of P_1 we must have $i \geq 1$. With a transfer to G^i, ν_1 induces a vertex element ν' of length $> C_0$. Now $\nu' = \sigma^{(N_i)}$ and σ gives rise to ν_1, so the chain was not maximal.

If all ν_i's are inactive, then the claim about the length follows from the definition of constants C_i. If ν_i is the first active element of the chain, then ν_{i+1} has length in $I_{i,2}$ by Lemma 5.14(2). With each generation the second index of the interval increases by two until $7k$ generations are complete (by (3) and (4)) or its length increases to some $I_{j,2}$ with $j < i$ by (2), and its life continues at least $7k$ more generations. This contradicts $s < 7k$. □

LEMMA 5.16. *If two vertex elements in P_i are labeled with no labeled vertex elements between them, then either at least one of them dies in the next $< k$ generations, or a birth occurs between them in the next $< k$ generations.*

Proof. If not, then the path between two such vertex elements is a Nielsen path (i.e., its lift to T connects two vertices whose stabilizers form a Nielsen pair). □

LEMMA 5.17. *Consider the cyclically ordered set of labels in each P_i.*

(1) *If two labels are adjacent, at least one is $< 3k$.*

(2) *If two labels have one label between them, then at least one is $< 4k$.*

(3) *If two labels have two labels between them, then at least one is $< 5k$.*

(4) *If two labels have three labels between them, then at least one is $< 6k$.*

Proof. Let a and b be two adjacent labels in some P_i with $a, b \geq 3k$ and assume that i is the smallest such i. Consider the ancestors of the two labels. According to Lemma 5.16, a death must occur between the two in some P_{i-s} with $s < k$. Thus in P_{i-s} we have labels $\cdots (a-s) \cdots x \cdots (b-s) \cdots$ and $x \geq 7k$. The dots between $(a-s)$ and $(b-s)$ are vertex elements that die before reaching P_i, and their labels are therefore $\geq 6k$. By our choice of i we conclude that x is the only label between $(a-s)$ and $(b-s)$. Now consider further ancestors of $(a-s)$, x, and $(b-s)$. Again by Lemma 5.16, a death must occur between vertex elements labeled $(a-s)$ and x in some P_{i-s-t} with $t<k$. We thus have two adjacent labels $\geq 5k$ in P_{i-s-t}, contradicting the choice of i.

Now suppose that in some P_i we have labels $\cdots axb\cdots$ and $a,b\geq 4k$. By (1) we must have $x<3k$. If a death occurs between a and x, or between b and x, in the previous k generations, then we obtain a contradiction to (1). If not, then by Lemma 5.16 we conclude that $x<k$ and then we have adjacent labels $a-x-1$ and $b-x-1$ in P_{i-x-1} contradicting (1).

Proofs of (3) and (4) are analogous. □

Proof of Proposition 5.13. Suppose that there are no $\mathcal{H}$-Nielsen pairs in T. Let $C_0, C_1, \dots, C_{7k}$ and ϵ be constants as explained above. Let m be an integer satisfying Lemma 5.14, and consider the labeling of vertex elements in paths P_i as above. The fact that $\mathcal{O}_{7k}^{N_{7k}}\cdots\mathcal{O}_2^{N_2}\mathcal{O}_1^{N_1}\in\mathcal{H}$ grows exponentially now follows from the observation that the number of labels in P_{i+k} is at least equal to the number of labels in P_i multiplied by $5/4$. Indeed, consider the labels in P_i that will die before reaching P_{i+k}. All such labels have to be $\geq 6k$ (since a vertex element cannot die before reaching the ripe old age of $7k$). By Lemma 5.17, any two such labels have at least three labels a, b, and c between them. By Lemma 5.16, there will be at least one birth between a and b and at least one birth between b and c between generations $i+1$ and $i+k$. Thus the number of deaths is at most a quarter of the number of labels in P_i, and the number of births is at least twice the number of deaths. The above inequality follows. □

5.5. *Distances between the vertices.*

Lemma 5.18. *Let V and W be two vertex stabilizers of T_0 and let d_j denote the distance between the vertices in T_j fixed by V and W. If V and W form a Nielsen pair for $\mathcal{H}$, then $d_0=d_1=d_2=\cdots$.*

Proof. Choose nontrivial elements $v\in V$ and $w\in W$. The distance between the vertices in T_j fixed by V and W equals $\frac{1}{2}\ell_{T_j}(vw)$ and the distance in T_{j+1} is analogously $\frac{1}{2}\ell_{T_{j+1}}(vw)$. The latter number can be computed as $\frac{1}{2}\ell_{T_j}(\hat{\mathcal{O}}_{j+1}^N(v)\hat{\mathcal{O}}_{j+1}^N(w))$ for large N, where $\hat{\mathcal{O}}_{j+1}$ denotes a lift of $\mathcal{O}_{j+1}$ to $\mathrm{Aut}(F_n)$ (since T_j is $\mathcal{O}_{j+1}$-nongrowing). This in turn equals the distance in T_j between the vertices fixed by $\hat{\mathcal{O}}_{j+1}^N(V)$ and $\hat{\mathcal{O}}_{j+1}^N(W)$. But that equals the distance between the vertices fixed by V and W since V and W form a Nielsen pair for $\langle\mathcal{O}_{j+1}\rangle$. □

Lemma 5.19. *Let $D_j\subset\mathbb{R}$ denote the set of distances between two distinct vertices in T_j with nontrivial stabilizer. Then*

(1) *D_j is discrete.*

(2) $D_j \supseteq D_{j+1}$.

(3) *There are finitely many F_n-equivalence classes of paths P joining two vertices of T_j with nontrivial stabilizer and with length*$(P) = \min D_j$.

(4) *If V and W are two nontrivial vertex stabilizers of T_j such that the distance between the corresponding vertices is* $\min D_j$, *then V and W form a Nielsen pair for $\langle \mathcal{O}_j \rangle$.*

(5) $\min D_j \leq \min D_{j+1}$, *and*

(6) *if* $\min D_j = \min D_{j+1}$ *then any two nontrivial vertex stabilizers V and W in T_{j+1} realizing the minimal distance also realize minimal distance in T_j.*

Proof. (1) Every element of D_j is a real number that can be represented as a linear combination of (finitely many) edge lengths of T_j with nonnegative integer coefficients. Hence D_j is discrete.

(2) Every element of D_{j+1} has the form $\frac{1}{2}\ell_{T_j}(\hat{\mathcal{O}}_{j+1}^N(v)\hat{\mathcal{O}}_{j+1}^N(w))$ (see the proof of Lemma 5.18) and hence occurs also as an element of D_j.

(3) Let P be such a path. The quotient map $T_j \to T_j/F_n$ is either injective on P or identifies only the endpoints of P, hence there are only finitely many possible images of P in the quotient graph. If two such paths have the same image, then they are F_n-equivalent.

(4) Since $\mathcal{O}_j$ fixes T_j, for any lift $\hat{\mathcal{O}}_j \in \mathrm{Aut}(F_n)$ of $\mathcal{O}_j$ we can choose an $\hat{\mathcal{O}}_j$-invariant isometry $\phi : T_j \to T_j$. By Proposition 4.48 and Lemma 4.47, ϕ induces the identity in the quotient graph. Therefore the immersed path P joining the vertices fixed by V and W is mapped by ϕ to a translate of itself (we are using the fact that all interior vertices of P have trivial stabilizer).

(5) is a consequence of (2).

(6) Choose a lift $\hat{\mathcal{O}}_{j+1} \in \mathrm{Aut}(F_n)$ of $\mathcal{O}_{j+1}$. For large N, the distance in T_{j+1} between the vertices fixed by V and W has the form $\frac{1}{2}\ell_{T_j}(\hat{\mathcal{O}}_{j+1}^N(v)\hat{\mathcal{O}}_{j+1}^N(w))$. It follows that for large N the immersed path P_N joining vertices in T_j fixed by $\hat{\mathcal{O}}_{j+1}^N(V)$ and $\hat{\mathcal{O}}_{j+1}^N(W)$ has length $\min D_j$. By (4), V and W form a Nielsen pair for h_j and therefore the paths P_N are translates of each other and have length $\min D_j$. □

5.6. *Proof of Theorem* 5.1. We are now ready for the proof of Theorem 5.1 which is reformulated as follows.

THEOREM 5.20. *Let $\mathcal{H} = \langle \mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_k \rangle$ be a unipotent subgroup of* $\mathrm{Out}(F_n)$. *By $\mathcal{F}$ denote a maximal $\mathcal{H}$-invariant proper free factor system. Let $T_0 \in \mathcal{T}$ have trivial edge stabilizers and satisfy $\mathcal{F}(T_0) = \mathcal{F}$. Then, the* *bouncing sequence that starts with T_0 is eventually constant. The stable value T is a simplicial tree with trivial edge stabilizers, a single orbit of edges and $\mathcal{F}(T) = \mathcal{F}$.*

Proof. By Lemma 5.10, the sequence consists of nongrowers, the vertex stabilizers are independent of the tree in the sequence, and all edge stabilizers are trivial. By Proposition 5.13, eventually all trees contain Nielsen pairs for $\mathcal{H}$. By Lemma 5.18 it follows that the numbers $\min D_j$ of Lemma 5.19 are bounded above and hence stabilized. Say $\min D_{j+1} = \min D_{j+2} = \cdots = \min D_{j+k}$. Let V and W be two nontrivial vertex stabilizers in T_{j+k} that realize $\min D_{j+k}$. By Lemma 5.19, the vertex stabilizers V and W form a Nielsen pair for every $\langle \mathcal{O}_i \rangle$, and hence for $\mathcal{H}$. Let P be the embedded path joining the corresponding vertices. Since P realizes the minimum distance between vertices with nontrivial stabilizer, its projection into the quotient graph is an embedding except perhaps at the endpoints (cf. the proof of Lemma 5.19(3)). If P projects onto the quotient graph, then this quotient graph has one edge and T_{j+k} is fixed by $\mathcal{H}$. If P does not project onto the quotient graph, we obtain a contradiction by collapsing P and its translates and thus constructing an $\mathcal{H}$-invariant proper free factor system strictly larger than $\mathcal{F}$. $\square$

6. Proof of the main theorem

In this section we show that Theorem 5.1 implies Theorem 1.1.

Recall from the introduction that for a marked graph G with a filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_K = G$ the set of upper triangular homotopy equivalences of G up to homotopy relative to the vertices is denoted by $\mathcal{Q}$.

LEMMA 6.1. *$\mathcal{Q}$ is a group under the operation induced by composition.*

Proof. Since the composition of upper triangular homotopy equivalences is clearly upper triangular, it suffices to show that if f is upper triangular, then there exists an upper triangular g such that $fg(E_i)$ and $gf(E_i)$ are homotopic rel endpoints to E_i for $1 \le i \le K$. We define $g(E_i)$ inductively starting with $g(E_1) = E_1$. Assume that g is defined on G_{i-1} and that $fg(E_j)$ and $gf(E_j)$ are homotopic rel endpoints to E_j for each $j < i$. If $f(E_i) = v_i E_i u_i$, define $g(E_i) = v_i' E_i u_i'$ where u_i' equals $\overline{g(u_i)}$ and v_i' equals $\overline{g(v_i)}$. Since v_i is a path in G_{i-1} with endpoints at vertices, $fg(v_i)$ is homotopic rel endpoints to v_i. Thus $f(v_i')$ is homotopic rel endpoints to $\overline{v}_i$ and $v_i f(v_i')$ is homotopic rel endpoints to the trivial path. A similar argument shows that $u_i f(u_i')$ is homotopic rel endpoints to the trivial path and hence that $fg(E_i) = f(v_i') v_i E_i u_i f(u_i')$ is homotopic rel endpoints to E_i. A similar argument showing that $gf(E_i)$ is homotopic rel endpoints to E_i completes the proof. $\square$

Proof that Theorem 5.1 *implies Theorem* 1.1. The proof is by induction on n. The $n = 1$ case is obvious so we may assume that Theorem 1.1 holds for rank less than n. By Theorem 5.1 there is an $\mathcal{H}$ invariant free factor system $\mathcal{F}'$ represented by either one free factor of rank $n-1$ or two free factors whose rank adds to n. Moreover, if $\mathcal{F}$ is an $\mathcal{H}$-invariant proper free factor system we may assume that $\mathcal{F} \sqsubset \mathcal{F}'$.

We will give the argument in the case that $\mathcal{F}' = \{[[F^1]], [[F^2]]\}$ where $F_n = F^1 * F^2$. The remaining case is analogous; details for both cases can be found in the first part of the proof of Lemma 2.3.2 of [BFH00].

The free factor system $\mathcal{F}$ induces free factor systems $\mathcal{F}^1$ and $\mathcal{F}^2$ of F^1 and F^2. By the inductive hypothesis, there are filtered marked graphs K^i with filtration elements realizing $\mathcal{F}^i$ and there are lifts of $\mathcal{H}|_{F^i}$ to $\mathcal{Q}^i$, the group of upper triangular homotopy equivalences of K^i up to homotopy relative to vertices. Define G to be the graph obtained from the disjoint union of K^1 and K^2 by adding an edge E with initial endpoint at a vertex $v_1 \in K^1$ and terminal endpoint at a vertex $v_2 \in K^2$. We may assume that F^i is identified with $\pi_1(K^i, v_i)$. Collapsing E to a point gives a homotopy equivalence of G to a graph whose fundamental group is naturally identified with $F^1 * F^2 = F_n$ and so provides a marking on G. A filtration on $K^1 \cup K^2$ is obtained by taking unions of filtration elements of $\mathcal{F}^1$ and of $\mathcal{F}^2$. Adding E as a final stratum produces a filtration of G in which $\mathcal{F}$ is realized by a filtration element.

For each $\mathcal{O} \in \mathcal{H}$, let $f_i \in \mathcal{Q}^i$ be the lift of $\mathcal{O}|_{F^i}$ and let Φ_i be an automorphism representing $\mathcal{O}$ whose restriction to F^i agrees with the automorphism induced by f_i under the identification of F^i with $\pi_1(K^i, v_i)$. Then $\Phi_1 = i_c\Phi_2$ for some $c \in F_n$. Represent c by a closed path γ based at v_1 and define $f : G \to G$ to agree with f_i on K_i and by $f(E) = [\gamma E]$. Then $f : G \to G$ is a topological representative of $\mathcal{O}$ and Corollary 3.2.2 of [BFH00] implies that, up to homotopy relative to vertices, $f(E) = \bar{u}_1 E u_2$ for some closed paths $u_i \subset K_i$. In other words f represents an element in the group $\mathcal{Q}$ of upper triangular homotopy equivalences of G up to homotopy relative to vertices.

It remains to arrange that $\mathcal{O} \mapsto f$ defines a homomorphism from $\mathcal{H}$ to $\mathcal{Q}$. It is convenient to subdivide E into edges E_i with common initial endpoint at the midpoint of E and with terminal endpoint at v_i. Thus $f(E_i) = E_i u_i$ and the fundamental groups of $K^i \cup E_i$ are identified with F^i. The automorphism $f_\#$ of F_n induced by f preserves both F^1 and F^2; this uniquely determines both $f_\#|_{F^1}$ and $f_\#|_{F^2}$. Replacing u_i with a different path $v_i \subset K^i$ replaces $f_\#|_{F^i}$ with $i_a f_\#|_{F^i}$ where a is the nontrivial element of F_n represented by the closed path $v_i\bar{u}_i$. If K^i has rank greater than one then $f_\#|_{F^i} \neq i_a f_\#|_{F^i}$ and so u_i is uniquely determined by $f_\#|_{F^i}$. If K^i has rank one then $f_\#|_{F^i}$ is independent of the choice of u_i and we always choose u_i to be trivial. In either case $\mathcal{O} \mapsto f$ defines a homomorphism from $\mathcal{H}$ to $\mathcal{Q}$. □

University of Utah, Salt Lake City, UT
E-mail address: bestvina@math.utah.edu

Rutgers University, Newark, Newark, NJ
E-mail address: feighn@newark.rutgers.edu

Lehman College, The City University of New York, Bronx, NY
E-mail address: michael@lehman.cuny.edu

References

[AB87] R. Alperin and H. Bass, Length functions of group actions on Λ-trees, in *Combinatorial Group Theory and Topology* (Alta, Utah, 1984), 265–378, Princeton Univ. Press, Princeton, NJ, 1987.

[Ali02] E. Alibegović, Translation lengths in Out(F_n) (Dedicated to John Stallings on the occasion of his 65th birthday), *Geom. Dedicata* **98** (2002), 87–93.

[BF91] M. Bestvina and M. Feighn, Bounding the complexity of simplicial group actions on trees, *Invent. Math.* **103** (1991), 449–469.

[BFH97a] M. Bestvina, M. Feighn, and M. Handel, Lamentations, trees, and irreducible automorphism of free groups, *Geom. Funct. Anal.* **7** (1997), 215–244.

[BFH97b] ———, Erratum to "Laminations, trees, and irreducible autormorphism of free groups", *Geom. Funct. Anal.* **7** (1997), 1143.

[BFH00] ———, The Tits alternative for Out(F_n): I. Dynamics of exponentially growing autormorphisms, *Ann. of Math.* **151** (2000), 517–623.

[BFH04] ———, Solvable subgroups of Out(F_n) are virtually abelian, *Geom. Dedicata* **104** (2004), 71–96.

[BH92] M. Bestvina and M. Handel, Train tracks and automorphisms of free groups, *Ann. of Math.* **135** (1992), 1–51.

[BLM83] J. S. Birman, A. Lubotzky, and J. McCarthy, Abelian and solvable subgroups of the mapping class groups, *Duke Math. J.* **50** (1983), 1107–1120.

[CL95] M. M. Cohen and M. Lustig, Very small groups actions on **R**-trees, *Proc. London Math. Soc.* **55** (1987), 571–604.

[Coo87] D. Cooper, Automorphisms of free groups have finitely generated fixed point sets, *J. Algebra* **111** (1987), 571–604.

[CM87] M. Culler and M. W. Morgan, Group actions on **R**-trees, *Proc. London Math. Soc.* **55** (1987), 571–604.

[CV86] M. Culler and K. Vogtmann, Moduli of graphs and automorphisms of free groups, *Invent. Math.* **84** (1986), 91–119.

[GL95] D. Gaboriau and G. Levitt, The rank of actions on **R**-trees, *Ann. Sci. École Norm. Sup.* **98** (1995), 549–570.

[SW79] P. Scott and T. Wall, Topological methods in group theory, in *Homological Group Theory* (*Proc. Sympos.*, *Durham*, 1977), 137–203, Cambridge Univ. Press, Cambridge, 1979.

[Ser80] J-P. Serre, *Trees* (Translated from French by John Stillwell), Springer-Verlag, New York, 1980.

[Ser92] ———, *Lie Algebras and Lie Groups* (1964 lectures given at Harvard University), second edition, Springer-Verlag, New York, 1992.

[Sho91] H. Short, Quasiconvexity and a theorem of Howson's, in *Group Theory from a Geometrical Viewpoint* (Trieste, 1990), 168–176, World Sci. Publ. Co., River Edge, NJ, 1991.

(Received October 4, 1996)
(Revised July 26, 2002)