Solvable subgroups of $Out(F_n)$ are virtually abelian

Mladen Bestvina, Mark Feighn, and Michael Handel*

June 1999

1 Introduction

Let F_n denote a free group of rank n. The group $Out(F_n)$ contains mapping class groups of compact surfaces and maps onto $GL(n, \mathbb{Z})$. It is perhaps not surprising that $Out(F_n)$ behaves at times like a mapping class group and at times like a linear group. J. Birman, A. Lubotzky, and J. McCarthy [BLM83] showed that solvable subgroups of mapping class groups are finitely generated and virtually abelian. Of course, $GL(3, \mathbb{Z})$ contains the Heisenberg group which is solvable but not virtually abelian. In this paper, we show that, with respect to the nature of solvable subgroups, $Out(F_n)$ behaves more like mapping class groups.

Theorem 1.1. *Every solvable subgroup of $Out(F_n)$ has a subgroup of index at most 3^{5n^2} that is finitely generated and free abelian.*

The rank of an abelian subgroup of $Out(F_n)$ is bounded by $vcd(Out(F_n)) = 2n - 3$ for $n > 1$ [CV86].

Since $Aut(F_n)$ embeds in $Out(F_{n+1})$, solvable subgroups of $Aut(F_n)$ are also virtually abelian. Theorem 1.1 complements [BFH96b] where we show that $Out(F_n)$ satisfies the Tits Alternative, i.e. that subgroups of $Out(F_n)$ are either virtually solvable or contain a free group of rank 2.

H. Bass and A. Lubotzky [BL94] have shown that solvable subgroups of $Out(F_n)$ are virtually polycyclic. In particular, they are finitely generated. We include an independent proof of this fact for completeness and because the ingredients of our proof are needed for the proof of Theorem 1.1.

We begin with a solvable subgroup $\mathcal{H}$ of $\mathrm{Out}(F_n)$. After passing to a subgroup of finite index if necessary, there is (Theorem 7.0.1 of [BFH96a]) an exact sequence

$$1 \to \mathcal{H}_0 \to \mathcal{H} \xrightarrow{\Omega} \mathbb{Z}^b \to 1$$

*All three authors gratefully acknowledge the support of the National Science Foundation.

where $\mathcal{H}_0$ is UPG (defined below). We will construct an embedding $\Phi : \mathcal{H}_0 \to \mathbb{Z}^r$ that virtually extends to a homomorphism $\Phi' : \mathcal{H} \to \mathbb{Z}^r$. The direct sum of Ω and Φ' is then an embedding of a finite index subgroup of $\mathcal{H}$ into $\mathbb{Z}^{b+r}$ showing that $\mathcal{H}$ is finitely generated and virtually free abelian.

Our approach is motivated by the special case that $\mathcal{H}$ is realized as a subgroup of the mapping class group of a compact surface S. The surface S decomposes into a union of annuli $A_1, \ldots, A_r$ and subsurfaces S_i; virtually every $\eta \in \mathcal{H}$ is represented by a homeomorphism $f : S \to S$ that restricts to a Dehn twist on each A_j and that preserves each S_i. If $\eta \in \mathcal{H}_0$ then each $f|S_i$ is the identity. The homomorphisms Φ and Φ' are defined by taking their j^{th} coordinates to be the number of twists that occurs across A_j.

We study an element of $\mathrm{Out}(F_n)$ by considering its lifts to $\mathrm{Aut}(F_n)$. More precisely, given $\eta \in Out(F_n)$ we focus on those automorphisms Θ that represent η and whose induced homeomorphisms of the Cantor set C_∞ of ends of F_n fixes at least three points. (See Section 2 for definitions.) This includes, but is not limited to, those Θ whose fixed subgroup has rank at least two.

In the course of proving Theorem 1.1 we prove the following result which is of independent interest.

Proposition 1.2. *Every abelian subgroup $\mathcal{H} \subset Out(F_n)$ has a virtual lift $\tilde{\mathcal{H}} \subset Aut(F_n)$. If x is a non-trivial primitive element of F_n whose conjugacy class is fixed by each element of $\mathcal{H}$, then $\tilde{\mathcal{H}}$ can be chosen so that each element of $\tilde{\mathcal{H}}$ fixes x.*

The paper is organized as follows. In Section 2 we establish notation and record known results for future reference. In Sections 3 and 4 we prove that $\mathcal{H}_0$ is finitely generated and free abelian and that the above exact sequence is virtually central. In Section 5 we prove Proposition 1.2 and in Section 6 we prove Theorem 1.1.

2 Notation and Preliminaries

2.1 Automorphisms and Lifts

We assume that F_n is identified with $\pi_1(R_n, *)$ where R_n is the rose with n petals and with vertex $*$. Let $\tilde{R}_n$ be the universal cover of R_n and let $\tilde{*}$ be a preferred lift of $*$. The space of ends of $\tilde{R}_n$ is a Cantor set that we denote C_∞.

A *marked graph* is a graph G with a preferred vertex v along with a homotopy equivalence $\tau : (R_n, *) \to (G, v)$ that identifies $\pi_1(G, v)$ with $\pi_1(R_n, *)$ and so also with F_n.

The universal cover Γ of G is a tree whose space of ends is denoted by $\mathcal{E}(\Gamma)$. We assume that a preferred lift $\tilde{v} \in \Gamma$ of $v \in G$ has been chosen. The marking homotopy equivalence lifts to an equivariant map $\tilde{\tau} : (\tilde{R}_n, \tilde{*}) \to (\Gamma, \tilde{v})$ that induces a homeomorphism from C_∞ to $\mathcal{E}(\Gamma)$ (see, for example, Section 2.2 of [BFH96a]). We use this homeomorphism to identify $\mathcal{E}(\Gamma)$ with C_∞ and so for the rest of this paper refer to the space of ends of Γ as C_∞.

The preferred vertex $\tilde{v} \in \Gamma$ provides an identification of the group $\mathcal{T}$ of covering translations of Γ with $\pi_1(G, v)$ and so with F_n. We sometimes write $T_x : \Gamma \to \Gamma$ for the covering translation corresponding to $x \in F_n$. The action of $\mathcal{T}$ on Γ extends to an action (of F_n) on C_∞ by homeomorphisms. If $T \in \mathcal{T}$ is non-trivial, then the homeomorphism $\hat{T} : C_\infty \to C_\infty$ determined by T fixes exactly two points; the line in Γ that connects these two points is the only T-invariant line in Γ and is called the axis of T.

A homotopy equivalence of a marked graph $f : G \to G$ determines an outer automorphism of $\pi_1(G, v)$ and so an element $\eta \in Out(F_n)$; we say that $f : G \to G$ represents η. To specify an automorphism of $\pi_1(G, v)$ induced by f, we must choose a path $\rho \subset G$ from v to $f(v)$. The path ρ also specifies a lift $\tilde{f} : \Gamma \to \Gamma$ of $f : G \to G$. Namely, let $\tilde{\rho}$ be the lift of ρ with initial endpoint $\tilde{v}$ and let $\tilde{f}$ be the lift of f that carries $\tilde{v}$ to the terminal endpoint of $\tilde{\rho}$. In this way, we have a bijection between lifts $\tilde{f} : \Gamma \to \Gamma$ of $f : G \to G$ and elements $\Phi \in Aut(F_n)$ that represent $\eta \in Out(F_n)$. As a particular case, the inner automorphism $i_x(y) = xyx^{-1}$ representing the identity element in $Out(F_n)$ corresponds to $T_x : \Gamma \to \Gamma$. In general $\tilde{f}T_x = T_{\Phi(x)}\tilde{f}$.

Each $\tilde{f} : \Gamma \to \Gamma$ extends to a homeomorphism $\hat{f} : C_\infty \to C_\infty$. Denote the group of homeomorphisms of C_∞ by $Homeo(C_\infty)$. The composition $\Phi \mapsto \tilde{f} \mapsto \hat{f}$ is independent of the choices of G, $f : G \to G$ and $\tilde{v}$; it defines an injective homomorphism from $Aut(F_n)$ to $Homeo(C_\infty)$. (See, for example, Section 3.2 of [BFH96a].) We will sometimes write $\hat{\Phi}$ instead of $\hat{f}$. If $\tilde{L} \subset \Gamma$ is a line with endpoints $P, Q \in C_\infty$, we denote the line with endpoints $\hat{f}(P)$ and $\hat{f}(Q)$ by $\tilde{f}_\#(\tilde{L})$. If $\tilde{L}$ is the axis of T_x then $\tilde{f}_\#(\tilde{L})$ is the axis of $T_{\Phi(x)}$.

Lemma 2.1. *Suppose that $n > 1$, that $f : G \to G$ is a homotopy equivalence of a marked graph, that $\tilde{f} : \Gamma \to \Gamma$ is a lift of $f : G \to G$ and that $\Phi : F_n \to F_n$ is the automorphism corresponding to $\tilde{f}$. Define $C = Fix(\hat{f})$ and $\mathbb{F} = Fix(\Phi)$. Define $\mathcal{T}(C)$ to be $\{T \in \mathcal{T} : \hat{T}(C) = C\}$ if $C \neq \emptyset$ and to be the empty set otherwise. Then the following are equivalent:*

(i) *$T_x \in \mathcal{T}(C)$.*

(ii) *$x \in \mathbb{F}$.*

(iii) *At least one (and hence both) endpoints of the axis of T_x are contained in C.*

(iv) *T_x commutes with $\tilde{f}$.*

(v) *$\hat{T}_x$ commutes with $\hat{f}$.*

Moreover, if $\mathbb{F}$ is non-trivial, then it is its own normalizer in F_n.

Proof of Lemma 2.1. The equivalence of (ii) and (iv) follows from $\tilde{f}T_x = T_{\Phi(x)}\tilde{f}$. The automorphism form of (v) is that $\hat{i}_x$ commutes with $\hat{\Phi}$. If $\Phi(x) = x$, then i_x commutes with Φ; thus (ii) implies (v). Since $i_x \Phi i_x^{-1}$ and Φ represent the same outer automorphism, they differ by an inner automorphism i_y. If (v) holds then $\hat{i}_y = \hat{T}_y$ is the identity. Since $n > 1$, T_y is the identity and Φ commutes with i_x. This latter condition is equivalent to (ii) so we have established the equivalence of (ii) and (v).

Let P and Q be, respectively, the initial and terminal endpoints of the axis $\tilde{L}$ of T_x. Then $\hat{f}(P)$ and $\hat{f}(Q)$ are respectively, the initial and terminal endpoints of the axis $\tilde{f}_\#(\tilde{L})$ of $T_{\Phi(x)}$. If $\hat{f}$ fixes P or Q, then $f_\#(\tilde{L})$ and $\tilde{L}$ share either an initial endpoint or a terminal endpoint and so must be equal. It follows that $\tilde{f}$ fixes both P and Q and that $\Phi(x)$ and x are both positive powers of the same primitive element of F_n. Since Φ is an automorphism, $\Phi(x) = x$. Thus (iii) implies (ii). For the converse, note that P and Q are, respectively, the unique source and sink for the action of $\hat{T}_x$ on C_∞. Any homeomorphism of C_∞ that commutes with $\hat{T}_x$ must therefore fix both P and Q so (v) implies (iii). We have now shown that (ii), (iii), (iv) and (v) are equivalent.

If $C = \emptyset$, then (i) and (iii) are equivalent. Assume then that $C \neq \emptyset$. It is immediate from the definitions that (v) implies (i). Assume that (i) holds and that P and Q are as above. If C contains a point R that is neither P nor Q, then C also contains $\lim_{k\to\infty} \hat{T}_x^k(R) = Q$. This proves that (i) implies (iii) and so completes the proof that the five conditions are equivalent.

Suppose that $T_x \in \mathcal{T}(C)$ and that the forward endpoint R of the axis of T_y is not in C. The endpoints of the axis of $T_y^n T_x T_y^{-n}$ are $\hat{T}_y^n(P)$ and $\hat{T}_y^n(Q)$, which, for sufficiently large n are close to R and therefore not in C. Given the equivalence of (i) and (iii), this proves that $\mathcal{T}(C)$ is its own normalizer in $\mathcal{T}$ and hence that $\mathbb{F}$ is its own normalizer in F_n. $\square$

We conclude this subsection by recording some facts for future reference. If $\tilde{\mathcal{H}} \subset Aut(F_n)$ is a lift of $\mathcal{H} \subset Out(F_n)$ then we denote the corresponding lift to $Homeo(C_\infty)$ by $\hat{\mathcal{H}}$. For any subgroup $\mathbb{F} \subset F_n$, the closure in C_∞ of the endpoints of the axes of elements in $\mathbb{F}$ is denoted by $C(\mathbb{F})$ and is naturally identified with the space of ends of $\mathbb{F}$. Although the UPG property referred to in the following lemma is not defined until the next section, it is convenient to place this result here. The part of the lemma that refers to UPG subgroups is quoted from [BFH96b] so there is no danger of circular reasoning.

A subgroup $\mathbb{F}$ of F_n is *invariant*, up to conjugacy, by a subgroup $\mathcal{H}$ of $Out(F_n)$ if each $\eta \in \mathcal{H}$ is represented by $\Phi \in Aut(F_n)$ satisfying $\Phi(\mathbb{F}) = \mathbb{F}$.

Lemma 2.2. *Suppose that $\mathcal{H}$ is a subgroup of $Out(F_n)$ and that $\mathbb{F}$ is a subgroup of F_n that is $\mathcal{H}$-invariant, up to conjugacy, and that is its own normalizer. Then*

1. *There is a well-defined restriction $\mathcal{H}|\mathbb{F} \subset Out(\mathbb{F})$.*

2. *If $\mathcal{H}$ is UPG, then $\mathcal{H}|\mathbb{F}$ is UPG.*

3. *If $\mathbb{F}$ has rank at least two, then any lift $\widetilde{\mathcal{H}|\mathbb{F}} \subset Aut(\mathbb{F})$ of $\mathcal{H}|\mathbb{F}$ extends uniquely to a lift $\tilde{\mathcal{H}} \subset Aut(F_n)$.*

4. *If $\widetilde{\mathcal{H}|\mathbb{F}} \subset Aut(\mathbb{F})$ is a lift of $\mathcal{H}|\mathbb{F}$ and $\hat{\Phi}$ fixes at least three points for each $\Phi \in \widetilde{\mathcal{H}|\mathbb{F}}$, then $\widetilde{\mathcal{H}|\mathbb{F}}$ extends uniquely to a lift $\tilde{\mathcal{H}} \subset Aut(F_n)$.*

Proof of Lemma 2.2. Given $\eta \in \mathcal{H}$, choose $\Phi \in Aut(F_n)$ that represents η and that preserves $\mathbb{F}$. Define $\eta|\mathbb{F}$ to be the element of $Out(\mathbb{F})$ determined by $\Phi|\mathbb{F}$. Since $\mathbb{F}$ is its own normalizer, two automorphisms of F_n that preserve $\mathbb{F}$ are conjugate by an element of F_n if and only if they are conjugate by an element of $\mathbb{F}$. Thus $\eta|\mathcal{H}$ is independent of the choice of Φ and is well defined. If $\mathbb{F}$ has rank at least two, then $\mathbb{F}$ is not contained in the fixed subgroup of any non-trivial inner automorphisms of F_n. In particular, if $\Phi_1, \Phi_2 \in Aut(F_n)$ represent η and restrict to the same element of $Aut(\mathbb{F})$, then $\Phi_1 = \Phi_2$. Part (3) follows easily. For part (4), suppose that $\Phi_1, \Phi_2 \in Aut(F_n)$ represent η and that $\hat{\Phi}_1$ and $\hat{\Phi}_2$ have at least three common fixed points. There is a covering translation T such that $\hat{\Phi}_1\hat{\Phi}_2^{-1} = \hat{T}$. Since $\hat{T}$ fixes at least three points, T is the identity and $\Phi_1 = \Phi_2$. This implies (4). Part (2) is Lemma 4.13 of [BFH96b]. □

2.2 Relative train track maps for UPG outer automorphism

We study an element $\eta \in Out(F_n)$ through its lifts $\Phi \in Aut(F_n)$. We analyze the Φ's by representing η as a homotopy equivalence $f : G \to G$ of a marked graph with particularly nice properties and studying the corresponding $\tilde{f}$'s. In this subsection we recall some properties of $f : G \to G$.

A *filtered graph* is a marked graph along with a filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_K = G$ where each (not necessarily connected) subgraph G_i is obtained from G_{i-1} by adding a single oriented edge E_i. We reserve the words path and loop for immersions of the interval and the circle respectively. We do not distinguish between two paths (or two loops) that differ only by an orientation preserving reparameterization. If ρ is a map of the interval or the circle into G or Γ, then $[\rho]$ is the unique path or loop that is homotopic to ρ rel endpoints if any. We say that a homotopy equivalence $f : G \to G$ *respects the filtration* if each $f(E_i) = E_i u_{i,f}$ for some loop $u_{i,f} \subset G_{i-1}$. In particular, such an f fixes each vertex of G. We denote the set of vertices of G by $\mathcal{V}$ and the group (Lemma 6.2 of [BFH96b]) of homotopy classes, relative to vertices, of filtration respecting homotopy equivalences of G by $FHE(G, \mathcal{V})$. There is a natural map from $FHE(G, \mathcal{V})$ to $Out(F_n)$. We say that $\eta \in Out(F_n)$ is UPG (for unipotent with polynomial growth) if it is in the image of $FHE(G, \mathcal{V})$ for some G. We say that a subgroup of $Out(F_n)$ is UPG if each of its elements is. The main theorem of [BFH96b] states that every finitely generated UPG subgroup $\mathcal{H}_0$ lifts to a subgroup $\mathcal{K}$ of $FHE(G)$ for some G. We say that $\mathcal{K}$ is a *Kolchin representative* of $\mathcal{H}_0$.

In general we will use $\mathcal{K}$ to denote a subgroup of $FHE(G, \mathcal{V})$.

We say that a line $\tilde{L}$ in Γ *has height* l if it projects to a line in G whose image is contained in G_l but not in G_{l-1}. Equivalently $\tilde{L}$ has height l if $\tilde{L}$ crosses (in either direction) an edge that projects to E_l but does not cross an edge that projects to E_j for $j > l$. If $\tilde{L}$ has height l and $\tilde{E}$ is an edge of Γ that projects to E_l and is crossed by $\tilde{L}$ then we say that the initial vertex of $\tilde{E}$ is a *highest edge splitting vertex for* $\tilde{L}$. Subdividing $\tilde{L}$ at all of its highest edge splitting vertices defines a decomposition $\tilde{L} = \ldots \tilde{\sigma}_{-1} \cdot \tilde{\sigma}_0 \cdot \tilde{\sigma}_1 \ldots$ that we refer to as the *highest edge splitting of* $\tilde{L}$. If $\tilde{f} : \Gamma \to \Gamma$ is a lift of $f \in FHE(G, \mathcal{V})$ and if $\tilde{L}$ is $\tilde{f}_\#$-invariant, then Lemma 5.2 of [BFH96a] implies that $[\tilde{f}(\tilde{\sigma}_j)] = \tilde{\sigma}_{j+r}$ for some r and all j. Roughly speaking, $\tilde{f}$ acts on $\tilde{L}$ by translating the pieces in the highest edge splitting by r units.

The definitions of height and highest edge splitting extend in the obvious way to rays or finite segments whose endpoints are vertices.

We will need the following fixed point results.

Lemma 2.3. *Suppose that G is a filtered graph.*

- *For any distinct $P_1, P_2, P_3 \in C_\infty$, there is a vertex $\tilde{v} \in \Gamma$ with the following property: if $f \in FHE(G, \mathcal{V})$ and $\hat{f}$ fixes P_1, P_2 and P_3, then $\tilde{f}$ fixes $\tilde{v}$.*
- *If $f \in FHE(G, \mathcal{V})$ and $\tilde{f} : \Gamma \to \Gamma$ is fixed point free, then $\hat{f}$ fixes exactly two points.*

Proof of Lemma 2.3. For the first item, let $\tilde{L}_{i,j} \subset \Gamma$ be the line connecting P_i to P_j and let $l_{i,j}$ be the height of $\tilde{L}_{i,j}$. Assuming without loss that $l_{1,2} \geq l_{1,3}, l_{2,3}$, let $\tilde{L} = \tilde{L}_{1,2}$ and let $\tilde{v}$ be any highest edge splitting vertex of $\tilde{L}$. Suppose that $f \in FHE(G, \mathcal{V})$ and that $\hat{f}$ fixes each P_i. If $\tilde{f}$ does not fix $\tilde{v}$ then $\tilde{f}$ translates the highest edge splitting vertices of $\tilde{L}$ away from one endpoint of $\tilde{L}$, say P_1, and toward the other, P_2. It follows that the highest edge splitting of $\tilde{L}$ is bi-infinite and hence that $l_{1,2} = l_{1,3} = l_{2,3}$. This implies that $\tilde{f}$ translates the highest edge splitting vertices of $\tilde{L}_{2,3}$ away from P_3 and toward P_2. But now $\tilde{f}$ translates the highest edge splitting vertices of $\tilde{L}_{1,3}$ away from P_1 and away from P_3. This contradiction verifies the first item.

We assume now that $f \in FHE(G, \mathcal{V})$ and that $\tilde{f}$ is fixed point free. By the first item, it suffices to show that $\hat{f}$ fixes at least two points. Implicit in Proposition 6.21 of [BFH96a] is the existence of a half-infinite ray $\tilde{R}_+ \subset \Gamma$ with highest edge splitting $\tilde{R}_+ = \tilde{\sigma}_0 \cdot \tilde{\sigma}_1 \ldots$ such that $[\tilde{f}(\tilde{\sigma}_j)] = \tilde{\sigma}_{j+r}$ for some $r > 0$ and all $j \geq 0$. Let $\tilde{v}_j$ be the initial vertex of $\tilde{\sigma}_j$. Since f fixes each vertex in G, $\tilde{f}$ restricts to a bijection of vertices and there are unique vertices $\tilde{v}_j$ so that $\tilde{f}(\tilde{v}_j) = \tilde{v}_{j+r}$ for all $j \in \mathbb{Z}$. Let $\tilde{\sigma}_j$ be the path connecting $\tilde{v}_j$ to $\tilde{v}_{j+1}$ for $j < 0$ and note that $[\tilde{f}(\tilde{\sigma}_j)] = \tilde{\sigma}_{j+r}$ for all j. If l is the height of $\tilde{R}_+$, then each $\tilde{\sigma}_j$ contains exactly one lift of E_l and these lifts are distinct for distinct j. (The $j < 0$ case follows from the $j \geq 0$ case which holds by construction.) It follows that $\ldots \tilde{\sigma}_{-1} \cdot \tilde{\sigma}_0 \cdot \tilde{\sigma}_1 \ldots$ is an embedded line whose endpoints are both fixed by $\hat{f}$. □

Lemma 2.4. *Up to conjugation by covering translations, each homotopy equivalence $f : G \to G$ of a marked graph has only finitely many lifts $\tilde{f} : \Gamma \to \Gamma$ such that $\hat{f} : C_\infty \to C_\infty$ fixes at least three points.*

Proof of Lemma 2.4. Given $P_1, P_2, P_3 \in Fix(\hat{f})$, let $\tilde{v} \in \Gamma$ be the unique vertex contained in each of the lines connecting two of the P_i's. The Bounded Cancellation Lemma [Coo87] implies that there is a bound, depending on f but not on the choice of the lift $\tilde{f}$ or the choice of the points P_i, to the length of the path $\tilde{\sigma}$ connecting $\tilde{v}$ to $\tilde{f}(\tilde{v})$. In particular, the projected image σ takes on only finitely many values as we vary the lift $\tilde{f}$ and the choice of points P_i.

Suppose that $\tilde{f}_i$, $i = 1, 2$, are lifts of f and that there exist lifts $\tilde{v}_i$ of a vertex v and $\tilde{\sigma}_i$ of a path σ such that $\tilde{\sigma}_i$ is the path connecting $\tilde{v}_i$ to $\tilde{f}_i(\tilde{v}_i)$. The covering translation $T : \Gamma \to \Gamma$ that carries $\tilde{v}_1$ to $\tilde{v}_2$ satisfies $\tilde{f}_2 T = T\tilde{f}_1$ and so conjugates $\tilde{f}_1$ to $\tilde{f}_2$. $\square$

3 Property A

In this section we prove that a finitely generated, solvable UPG subgroup $\mathcal{H}_0$ embeds in some $\mathbb{Z}^r$ and hence that every solvable UPG subgroup is free abelian. Some of the arguments proceed by induction on the skeleta G_i of a filtered graph G. For this reason, we do not assume in this section that $\mathcal{K}$ is a Kolchin representative of some $\mathcal{H}_0$ but only that *$\mathcal{K}$ is a finitely generated solvable subgroup of $FHE(G, \mathcal{V})$ with the following feature: if E_i is not a loop then there exists $f \in \mathcal{K}$ such that $u_{i,f}$ is non-trivial.* In particular, we do not assume that $\mathcal{K}$ injects into $Out(\pi_1(G))$, we allow G to have valence one vertices and we allow G to be disconnected.

Remark 3.1. Suppose that $\mathcal{K}$ is a Kolchin representative for $\mathcal{H}_0$ and that E_i is not a loop. If $u_{i,f}$ is trivial for all $f \in \mathcal{K}$ then collapsing E_i to a point produces a new Kolchin representative for $\mathcal{H}_0$ with one fewer edge. It is therefore always possible to satisfy the above italicized property when choosing Kolchin representatives.

Remark 3.2. If $u_{i,f}$ is non-trivial for some $f \in \mathcal{K}$, then the terminal vertex of E_i is contained in a loop in G_{i-1} and so must have valence at least two in G_{i-1}. Thus if E_i is the first edge to contain a vertex v, then either E_i is a loop or E_i has v as initial vertex. In either case, v is the initial vertex of E_i.

The group of lifts $\tilde{f} : \Gamma \to \Gamma$ [respectively $\hat{f} : C_\infty \to C_\infty$] of elements of $\mathcal{K}$ has a natural projection to $\mathcal{K}$. By an *action of $\mathcal{K}$ on Γ [respectively C_∞] by lifts* we mean a section of this projection. In other words, an action by lifts is an assignment $f \mapsto \tilde{f}$ [respectively $f \mapsto \hat{f}$] that respects composition (up to homotopy relative to the vertices).

For each edge $E_i \subset G$ choose, once and for all, a lift $\tilde{E}_i^* \subset \Gamma$. Define $s_i(f) : \Gamma \to \Gamma$ to be the unique lift of $f \in \mathcal{K}$ that fixes the initial endpoint of $\tilde{E}_i^*$ and note that *s_i is an action of $\mathcal{K}$ on Γ by lifts.* Denote the terminal endpoints of E_i and $\tilde{E}_i^*$ by v_i and $\tilde{v}_i$ respectively.

Definition 3.3. If E_i is not a component of G_i denote the component of G_{i-1} that contains v_i by B_i and the copy of the universal cover of B_i that contains $\tilde{v}_i$ by Γ_{i-1}. For each $f \in \mathcal{K}$, Γ_{i-1} is $s_i(f)$-invariant.

We say that $\mathcal{K}$ (and the choice of the $\tilde{E}_i^*$'s) satisfies *Property A* (for abelian) if there is a collection of lines $\tilde{L}_i \subset \Gamma_{i-1}$, one for each E_i that is not a component of G_i, such that :

- $\tilde{L}_i$ is $s_i(f)_\#$-invariant for all $f \in \mathcal{K}$.
- $\tilde{v}_i$ is a highest edge splitting vertex for $\tilde{L}_i$.
- If $\tilde{L}_i$ and $\tilde{L}_j$ differ by a covering translation, then $\tilde{L}_i = \tilde{L}_j$ and $\tilde{v}_i = \tilde{v}_j$.

If $\tilde{L}_i$ projects to an primitive loop α_i and if $u_{i,f}$ is non-trivial for some $f \in \mathcal{K}$, then we say (abusing notation slightly) that α_i is an *essential axis* and that E_i is an *essential edge (associated to α_i)*. The set of essential axes is denoted $A(\mathcal{K})$ and the set of essential edges is denoted $E(\mathcal{K})$. If the essential axis α is associated to m_α essential edges, then we say that m_α is the *multiplicity* of α. In the analogy with the mapping class group of a compact surface, $A(\mathcal{K})$ is the set of reducing curves and $\mathcal{K}$ can be chosen so that each $m_\alpha = 1$.

Example 3.4. Let G be the graph with vertices w_1 and w_2 and edges $E_1, \dots, E_6$ such that: E_1, E_2 and E_3 are loops based at w_1; E_4 and E_5 are loops based at w_2; and E_6 has initial endpoint at w_1 and terminal endpoint at w_2. Define $f_1 : G \to G$ by $E_2 \mapsto E_2E_1$, $E_3 \mapsto E_3E_1$ and $E_j \mapsto E_j$ otherwise. Define $f_2 : G \to G$ by $E_5 \mapsto E_5E_4$, $E_6 \mapsto E_6E_5$ and by $E_j \mapsto E_j$ otherwise. Then f_1 and f_2 generate an abelian subgroup $\mathcal{K}$ of $FHE(G, \mathcal{V})$.

For $i = 1$ and 4, E_i is a component of G_i. Choose a lift $\tilde{w}_1 \in \Gamma$ of w_1. Define $\tilde{E}_2^*$ and $\tilde{E}_3^*$ to be the lifts of E_2 and E_3 that terminate at $\tilde{w}_1$, define $\tilde{v}_2 = \tilde{v}_3 = \tilde{w}_1$, and define $\tilde{L}_2 = \tilde{L}_3$ to be the line that projects to E_1 and that contains $\tilde{w}_1$. Choose a lift $\tilde{w}_2 \in \Gamma$ of w_2. Define $\tilde{E}_5^*$ and $\tilde{E}_6^*$ to be the lifts of E_5 and E_6 that terminate at $\tilde{w}_2$ and define $\tilde{v}_5 = \tilde{v}_6 = \tilde{w}_2$. The line $\tilde{L}_5$ projects to E_4 and contains $\tilde{w}_2$. Let $R_2 \subset G$ be the half-infinite line $E_5E_5E_4E_5E_4^2E_5E_4^3\dots$, let $R_1 \subset G$ be the half-infinite line $\dots E_5\bar{E}_4^3E_5\bar{E}_4^2E_5\bar{E}_4$, let $\tilde{R}_1$ be the lift of R_1 terminating at $\tilde{w}_2$ and let $\tilde{R}_2$ be the lift of R_2 initiating at $\tilde{w}_2$. Then $\tilde{L}_6 = \tilde{R}_1\tilde{R}_2$ does not project to a loop in G. We conclude that $\mathcal{K}$ satisifies Property A and that the essential axes of $\mathcal{K}$ are E_1, with multiplicity two and associated essential edges E_2 and E_3, and E_4, with multiplicity one and associated essential edge E_5.

For each essential axis α, choose an essential edge E_i associated to α and fix notation (that is independent of the choice of E_i) as follows: $\tilde{v}_\alpha = \tilde{v}_i \in \tilde{L}_i$; T^α is a primitive covering translation with axis $\tilde{L}_i$; and $s_\alpha(f)$ is the lift of $f \in \mathcal{K}$ that fixes $\tilde{v}_\alpha$.

Lemma 3.5. **(1)** *s_α defines an action of $\mathcal{K}$ on Γ by lifts.*

(2) $s_i(f) = s_\alpha(f)$ *if and only if* $u_i(f)$ *is trivial.*

(3) $s_\alpha(f)_\#(\tilde{L}_i) = \tilde{L}_i$*; in particular,* $s_\alpha(f)$ *commutes with* T^α*.*

Proof of Lemma 3.5. (1) and (2) are immediate from the definitions. If $\tilde{L}_i = \ldots \tilde{\sigma}_{-1} \cdot \tilde{\sigma}_0 \cdot \sigma_1 \ldots$ is the highest edge splitting of $\tilde{L}_i$, then $[s_i(f)(\tilde{\sigma}_j)] = \tilde{\sigma}_{j+r}$ for some r and all j. Let $\sigma_j \subset G$ be the image of $\tilde{\sigma}_j \subset \Gamma$. Since $\tilde{L}_i$ is the axis of a covering translation, the splitting is periodic and there is a smallest positive integer k such that $[f^k(\sigma_j)] = \sigma_j$ for all j. Since f represents a UPG outer automorphism, $k = 1$ (see Theorem 5.1.8 of [BFH96a] or the proof of Proposition 4.5 of [BFH96b]).We conclude that $f_\#(\sigma_j) = \sigma_j$ for all j and hence that $[s_\alpha(f)(\tilde{\sigma}_j)] = \tilde{\sigma}_j$ for all j. This implies (3). □

The following lemma justifies our notation.

Lemma 3.6. *If* $\mathcal{K}$ *satisfies Property A, then there is an injective homomorphism* $\Phi_{\mathcal{K}} : \mathcal{K} \to \mathbb{Z}^r$ *where* r *is the cardinality of* $E(\mathcal{K})$*. In particular,* $\mathcal{K}$ *is free abelian.*

Proof of Lemma 3.6. If E_i is an essential edge with essential axis α then $s_i(f)$ and $s_\alpha(f)$ are lifts of f that commute with T^α and so differ by an iterate of T^α. Define a homomorpism $\phi^i_{\mathcal{K}} : \mathcal{K} \to \mathbb{Z}$ by $s_i(f) = (T^\alpha)^{\phi^i_{\mathcal{K}}(f)} s_\alpha(f)$ and define $\Phi_{\mathcal{K}}$ to be the product of the $\phi^i_{\mathcal{K}}$'s. By Lemma 3.5(2), f is in the kernel of $\Phi_{\mathcal{K}}$ if and only $u_{i,f}$ is trivial for each essential edge E_i.

It suffices to show that if $f \in \mathcal{K}$ is not the identity map and if $i > 1$ is the smallest positive integer for which $u_{i,f}$ is non-trivial, then E_i is an essential edge. Since $u_{j,f}$ is trivial for all $j < i$, $f|G_{i-1}$ is the identity and $s_i(f)$ is the restriction to Γ_{i-1} of a non-trivial covering translation $T : \Gamma \to \Gamma$. The line $\tilde{L}_i \subset \Gamma_{i-1}$ is T-invariant so must be the axis of T. Thus $E_i \in E(\mathcal{K})$ as desired. □

The following lemma produces a pair of fixed points in C_∞ or equivalently a fixed line in Γ.

Lemma 3.7. *Suppose that* $\Phi_1, \ldots, \Phi_m \in Aut(F_n)$ *commute and that their images in* $Out(F_n)$ *generate a* UPG *subgroup. Then* $\cap_{j=1}^m Fix(\hat{\Phi}_j)$ *contains at least two points.*

Proof of Lemma 3.7. Choose a Kolchin representative $\mathcal{K}$ of the UPG subgroup. There are elements $f_j \in \mathcal{K}$ and commuting lifts $\tilde{f}_j : \Gamma \to \Gamma$ so that $\hat{f}_j = \hat{\Phi}_j$ for $j = 1, \ldots, m$.

If $\tilde{f}_1$ is fixed point free then Lemma 2.3(2) implies that $\mathrm{Fix}(\hat{f}_1)$ is a pair of points $\{P, Q\}$. Since $\hat{f}_j$ commutes with $\hat{f}_1$, $\hat{f}_j$ setwise preserves $\{P, Q\}$ and we need only show that $(\tilde{f}_j)_\#$ does not reverse the orientation on the line $\tilde{L}$ connecting P and Q. Let $\tilde{L} = \ldots \tilde{\sigma}_{-1} \cdot \tilde{\sigma}_0 \cdot \tilde{\sigma}_1 \ldots$ be the highest edge splitting of $\tilde{L}$. If $(\tilde{f}_j)_\#$ reverses the orientation on $\tilde{L}$, then for some j, $[\tilde{f}(\tilde{\sigma}_j)]$ equals $\tilde{\sigma}_j$ with its orientation reversed. The projected image $\sigma_j \subset G$ determines a conjugacy class in F_n that is periodic but not fixed under the action of the outer automorphism determined by f. This contradicts Proposition 4.5 of [BFH96b].

Suppose next that $Fix(\tilde{f}_1) \neq \emptyset$ and that $\tilde{f}_1$ does not commute with any covering translations. The fixed point set of f_1 is a union of vertices and entire edges, so $\tilde{f}_1$ must fix at least one vertex. Each edge of G has at most one lift to Γ whose initial endpoint is fixed by $\tilde{f}_1$. Indeed, if there were two, then the covering translation that carried one to the other would commute with $\tilde{f}_1$. By Remark 3.2, there is at least one edge $\tilde{E}_k$ whose initial endpoint is fixed by $\tilde{f}_1$; we may assume that k is minimal. Let $\tilde{v}$ be the initial endpoint of $\tilde{E}_k$. By the same argument as for edges, no non-trivial translate of $\tilde{v}$ is fixed by $\tilde{f}_1$. Since the $\tilde{f}_j$'s permute the fixed points of $\tilde{f}_1$ and permute the translates of $\tilde{v}$, each $\tilde{f}_j$ must fix $\tilde{v}$. Let B' be the component of G_{k-1} that contains the terminal endpoint of E_k. (If B' did not exist, then E_k would be a loop that was fixed by f_1 and the terminal endpoint of $\tilde{E}_k$ would be a fixed translate of $\tilde{v}$.) Denote the copy of the universal cover of B' that contains the terminal endpoint of $\tilde{E}_k$ by $\Gamma'_{k-1} \subset \Gamma$. Since each f_j maps an initial segment of E_k to an initial segment of E_k, Γ'_{k-1} is $\tilde{f}_j$-invariant for each j. By our choice of k, the restriction $\tilde{f}_1|\Gamma'_{k-1}$ is fixed point free. We can now repeat the argument of the first case on the restriction of the $\tilde{f}_j$'s to Γ'_{k-1}.

Let $C = Fix(\hat{f}_1)$, $\mathbb{F} = Fix(\Phi_1)$ and $\mathcal{T}(C) = \{T \in \mathcal{T} : \hat{T}(C) = C\}$. The space of ends of $\mathbb{F}$ naturally embeds in C as the closure of the set of endpoints of axes for elements of $\mathcal{T}(C)$. (See Lemma 2.1.) It remains to consider the case that $\mathcal{T}(C)$ and hence $\mathbb{F}$ is non-trivial. Since the Φ_j's commute, $\mathbb{F}$ is Φ_j invariant and C is $\hat{f}_j$-invariant for all j. Lemma 2.2 implies that the image of the $\Phi_j|\mathbb{F}$'s in $Out(\mathbb{F})$ generate a UPG subgroup.

If $m = 1$, then the lemma follows from the fact that C contains at least two points. We now argue by induction, with the induction statement being that the lemma holds for all $m' < m$. By the inductive hypothesis, there exist $P, Q \in C$ that are fixed by the $m-1$ maps $\hat{f}_2|C, \ldots, \hat{f}_m|C$. Since $\hat{f}_1|C$ is the identity, P and Q are fixed by each $\hat{f}_j$. □

Remark 3.8. We will apply Lemma 3.7 using the language of lifts rather than automorphims. The precise statement is: if $\tilde{f}_j : \Gamma \to \Gamma$ $(j = 1, \ldots m)$ are commuting lifts of elements of $FHE(G, \mathcal{V})$, then $\cap_{j=1}^m Fix(\hat{f}_j)$ contains at least two points.

Corollary 3.9. *Every finitely generated solvable UPG subgroup $\mathcal{H}_0$ has a Kolchin representative that satisfies Property A; in particular, $\mathcal{H}_0$ is free abelian.*

Proof of Corollary 3.9. Choose a Kolchin representative $\mathcal{K} \subset FHE(G, \mathcal{V})$ for $\mathcal{H}_0$. We work our way up the strata E_i, modifying $\mathcal{K}$ so that it satisfies Property A. The inclusion $G_i \subset G$ defines a restriction homomorphism $R_i : FHE(G, \mathcal{V}) \to FHE(G_i, \mathcal{V} \cap G_i)$; denote $R_i(\mathcal{K})$ by $\mathcal{K}_i$. Since E_1 is a component of G_1, $\mathcal{K}_1$ has Property A; we may assume by induction that $\mathcal{K}_{i-1}$ satisfies Property A and is therefore abelian.

If E_i is a component of G_i then $\mathcal{K}_i$ satisfies Property A. We may therefore assume that Γ_{i-1} and s_i are defined. The ends of Γ_{i-1} define a subset $C_\infty^* \subset C_\infty$.

Let $R : \mathcal{K}_i \to \mathcal{K}_{i-1}$ be the restriction homomorphism. We suppose at first that R is an isomorphism, and hence that $\mathcal{K}_i$ is abelian. Choose $g_1, \dots, g_m \in \mathcal{K}$ so that $g_1|G_i, \dots, g_m|G_i$ generate $\mathcal{K}_i$. Since the $g_j|G_i$'s commute, s_i is an action and $s_i(g_j)|\Gamma_{i-1}$ depends only on $g_j|G_i$, the $s_i(g_j)|\Gamma_{i-1}$'s commute. By Remark 3.8, there exist $P_i, Q_i \in C_\infty^*$ that are fixed by each $\hat{s}_i(g_j)$ and hence by $\hat{s}_i(f)$ for each $f \in \mathcal{K}$. The line $\tilde{L}_i$ connecting P_i to Q_i is $s_i(f)_\#$-invariant for each $f \in \mathcal{K}$.

If $\tilde{L}_i$ and $\tilde{L}_j$ differ by a covering translation for some $j < i$, then after replacing $\tilde{E}_i^*$ by a translate if necessary, we may assume that $\tilde{L}_i = \tilde{L}_j$. If $\tilde{L}_i = \tilde{L}_j$ for some $j < i$, let $\tilde{w}_i = \tilde{v}_j$; otherwise let $\tilde{w}_i$ be any highest edge splitting vertex in $\tilde{L}_i$. If $\tilde{w}_i = \tilde{v}_i$ then $\mathcal{K}_i$ satisfies Property A and we have finished the induction step. If not, we modify G and Γ by an operation that simply reattaches E_i by 'sliding' its terminal endpoint along the path in G that is the image of the path in Γ connecting $\tilde{v}_i$ to $\tilde{w}_i$. This operation is fully described in Section 5.4 of [BFH96a] (see in particular the definition of sliding, Lemma 5.4.2 and the text preceding the statement of Lemma 5.4.2.) and is performed simultaneously to each $f \in \mathcal{K}$. This does not change $s_i(f)|\Gamma_{i-1}$ and respects the group structure of $\mathcal{K}$. A new Kolchin representative (still called $\mathcal{K}$) is produced that agrees with the old one on G_{i-1} and has the additional feature that $\tilde{w}_i$ is the terminal vertex of the new $\tilde{E}_i^*$. Property A is now verified and we have completed the inductive step.

It remains to consider the case that the kernel K of R is non-trivial. Since each element of K restricts to the identity on G_{i-1} : (i) the assignment $f^* \mapsto u_{i,f^*}$ defines a homomorphism from K into the free group $\pi_1(G, v_i)$; and (ii) the restriction of $s_i(f^*)$ to Γ_{i-1} is a covering translation. Since $\mathcal{K}$, and hence K, is solvable, the image of the homomorphism of (i) is isomorphic to Z. In other words, there exists a closed path γ based at v_i such that $u_{i,f^*} = [\gamma^{k(f^*)}]$ for all $f^* \in K$. If $\tilde{\gamma}$ is the lift of γ that begins at $\tilde{v}_i$ and T is the covering translation of Γ_{i-1} that carries the initial endpoint of $\tilde{\gamma}$ to the terminal endpoint of $\tilde{\gamma}$, then $s_i(f^*)|\Gamma_{i-1} = T^{k(f^*)}|\Gamma_{i-1}$ for all $f^* \in K$. If $s_i(f^*)|\Gamma_{i-1}$ is not the identity, then $Fix(\hat{s}_i(f^*)|C_\infty^*$ is the endpoint set $\{P, Q\}$ of the axis $\tilde{L}_i$ of T.

Choose $f^* \in K$ such that $s_i(f^*)|\Gamma_{i-1}$ is not the identity. For each $f \in \mathcal{K}_i$ there exists $f_0^* \in K$ such that $ff^*f^{-1} = f_0^*$. Then $\{P, Q\} = Fix(\hat{s}_i(f_0^*)|C_\infty^*) = \hat{s}_i(f)(Fix(\hat{s}_i(f^*)|C_\infty^*)) = \hat{s}_i(f)(\{P, Q\})$. As we argued in the proof of Lemma 3.7, this implies that $\hat{s}_i(f)$ fixes P and Q. We have now verified that $\tilde{L}_i$ is $s_i(f)_\#$-invariant for each $f \in \mathcal{K}$. The proof now concludes as in the previous case. ☐

4 Abelian subgroups are finitely generated

In this section we prove that an abelian UPG subgroup $\mathcal{H}_0$ is finitely generated. Section 3 and the exact sequence $1 \to \mathcal{H}_0 \to \mathcal{H} \xrightarrow{\Omega} \mathbb{Z}^b \to 1$ then implies that every solvable subgroup of $Out(F_n)$ is finitely generated. This fact was originally proved by H. Bass and A. Lubotzky [BL94]. We also show that the above sequence is a virtually central extension.

If $\mathcal{K}$ is a Kolchin representative of $\mathcal{H}_0$ that satisfies Property A, if $f \in \mathcal{K}$ and if $\Phi_{\mathcal{K}} : \mathcal{K} \to \mathbb{Z}^r$ is the embedding of Lemma 3.6, then the coordinates of $\Phi_{\mathcal{K}}(f)$ are defined in terms of the $s_i(f)$'s and the $s_\alpha(f)$'s. Our goal is to recognize these lifts of f by their induced actions on C_∞, thus removing their dependence on the choice of $\mathcal{K}$.

We begin by placing an additional restriction (the second item below) on our Kolchin graphs and reiterating our assumption from Section 3 (the first item below).

Lemma 4.1. *For every finitely generated UPG subgroup $\mathcal{H}_0$ there is a Kolchin representative $\mathcal{K}$ with the following properties.*

- *If E_i is not a loop then $u_{i,f}$ is non-trivial for some $f \in \mathcal{K}$.*
- *Every vertex of G is the initial vertex of at least two edges.*

Proof of Lemma 4.1. Start with any Kolchin representative $\mathcal{K}$. After restricting to a subgraph if necessary, we may assume that G has no valence one vertices. We will show that if either of the two properties fail, then we can replace G by a graph with fewer edges. This process terminates after finitely many steps to produce the desired Kolchin representative.

If the first property fails, then E_i is pointwise fixed for each $f \in \mathcal{K}$ and we may collapse it to a point. Suppose that the first property holds, that v is a vertex and that E_i and E_{j_0} are the first and second edges that contain v. By Remark 3.2, v is the initial endpoint of E_i. If E_i is not a loop, then v is a valence one vertex in G_{j_0-1} and so is not contained in a loop in G_{j_0-1}. In this case, v can not be the terminal vertex of E_{j_0} unless it is also the initial vertex of E_{j_0}.

Suppose then that E_i is a loop and that $E_{j_0}, \dots, E_{j_m}$ are the other edges that contain v. If the second property fails, then $E_{j_0}, \dots, E_{j_m}$ are non-loops with v as terminal endpoint and $u_j(f) = E_i^{k_j(f)}$ for $j = j_0, \dots, j_m$. Redefine $\mathcal{K}$ by replacing each $k_j(f)$ with $k_j(f) - k_{j_0}(f)$. This can be achieved on the graph level by postcomposing f with a map that is homotopic to the identity and that drags v $j_0(f)$ times around E_i in the reverse direction. This has no effect on the outer automorphism determined by f and respects the group structure on $\mathcal{K}$. The edge E_{j_0} is now fixed by each $f \in \mathcal{K}$ and so can be collapsed to point. $\square$

Definition 4.2. We say that a Kolchin representative $\mathcal{K}$ for $\mathcal{H}_0$ is an *abelian Kolchin representative* if it satisfies Property A and the conclusions of Lemma 4.1.

We now turn to the task of defining axes without reference to a Kolchin representative. Recall that in the analogy with the mapping class group of a compact surface M, $A(\mathcal{K})$ corresponds to the set of reducing curves in the minimal reduction. Such reducing curves are completely characterized [HT85] as follows : the free homotopy class determined by a closed curve $\alpha \subset M$ is an element of the minimal reducing set for the mapping class represented by a homeomorphism $h : M \to M$ if and only if for some (and hence each) covering translation $T : \tilde{M} \to \tilde{M}$ corresponding to α, there are two lifts $\tilde{h}_1, \tilde{h}_2 : \tilde{M} \to \tilde{M}$ of h that commute with T and whose extensions $\hat{h}_i$ over the 'circle at infinity' fix at least three points. The analogous result in the non-geometric case is given in Corollary 4.7.

Definition 4.3. Suppose that $x \in F_n$ and that G is a marked graph. For any $\psi \in Out(F_n)$, define the set of interesting lifts of ψ (with respect to x), denoted by either $IL(\psi, T_x)$ or $IL(\psi, x)$, to be $\{\hat{\Phi} : C_\infty \to C_\infty :$

- $\Phi \in Aut(F_n)$ represents ψ.
- $\hat{\Phi}$ commutes with $\hat{T}_x$ (or equivalently Φ fixes x).
- $\hat{\Phi}$ fixes at least three points$\}$.

The following lemma considers the effect of conjugation on $IL(\psi, T_x)$. Part (1) is a special case of part (2) with $\Phi = i_y$ and is included to make specific references easier. The proof of the lemma is immediate from the definitions.

Lemma 4.4. *Suppose that G is a marked graph, that $x \in F_n$ and that $\psi \in Out(F_n)$.*

(1) *For any $y \in F_n$, $\hat{T}_y(IL(\psi, T_x))\hat{T}_y^{-1} = IL(\psi, T_y T_x T_y^{-1}) = IL(\psi, T_{yxy^{-1}})$.*

(2) *For any $\eta \in Out(F_n)$ and $\Theta \in Aut(F_n)$ representing η, $\hat{\Theta}(IL(\psi, x))\hat{\Theta}^{-1} = IL(\eta\psi\eta^{-1}, \Theta(x))$.*

Recall that $\alpha \in A(\mathcal{K})$ has an associated primitive covering translation T^α whose axis equals $\tilde{L}_i$ for each essential edge E_i associated to α. The next lemma states that the actions s_α and s_i produce interesting lifts.

Lemma 4.5. *Suppose that $\mathcal{K}$ is an abelian Kolchin representative, that α is an essential axis and that E_i is an essential edge associated to α. Then $\bigcap_{f \in \mathcal{K}} Fix(\hat{s}_\alpha(f))$ and $\bigcap_{f \in \mathcal{K}} Fix(\hat{s}_i(f))$ each contain at least one point that is not an endpoint of $\tilde{L}_i$. In particular, if f represents $\psi \in \mathcal{H}_0$, then $\hat{s}_\alpha(f), \hat{s}_i(f) \in IL(\psi, T^\alpha)$.*

Proof of Lemma 4.5. To prove the $\hat{s}_\alpha(f)$ case we will find a point $R_\alpha \in C_\infty$ that is fixed by each $\hat{s}_\alpha(f)$. Since $\mathcal{K}$ satisfies the conclusions of Lemma 4.1, there is an edge $\tilde{E}_j \subset \Gamma$ with initial vertex $\tilde{v}_i = \tilde{v}_\alpha$ and with projection $E_j \subset G$ not equal to α. If $u_{j,f}$ is trivial for all $f \in \mathcal{K}$, then E_j is a loop that is fixed by each $f \in \mathcal{K}$. In this case, choose R_α to be an endpoint of the line that projects to E_j and that contains $\tilde{v}_i$. If $u_{j,f}$ is non-trivial for some $f \in \mathcal{K}$, then $\tilde{L}_j$ is defined. There is a covering translation T that carries $\tilde{v}_i$ to the initial endpoint of $\tilde{E}_j^*$; define $\tilde{L}_j' = T^{-1}(\tilde{L}_j)$. Since $s_\alpha(f)$ and $T^{-1}s_j(f)T$ are lifts of f that fix $\tilde{v}_i$, they are equal. In this case, choose R_α to be an endpoint of $\tilde{L}_j'$.

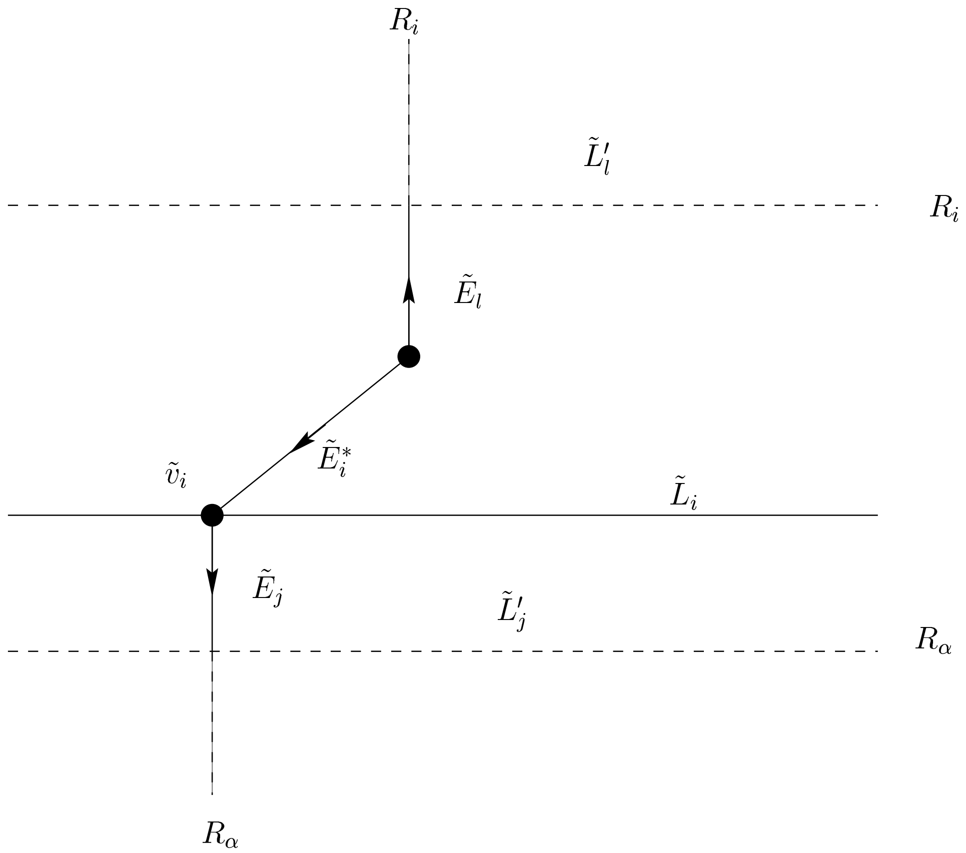

The proof for $s_i(f)$ is similar. Since $\mathcal{K}$ satisfies the conclusions of Lemma 4.1, there exists an edge $\tilde{E}_l \neq \tilde{E}^*_i$, with the same initial endpoint as $\tilde{E}^*_i$. A common fixed point R_i for each $\hat{s}_i(f)$ is found as in the previous case using $\tilde{E}_l$ in place of $\tilde{E}_j$. The details are left to the reader. $\square$

If $f \in \mathcal{K}$ represents $\psi \in \mathcal{H}_0$, then we use $\hat{s}(f)$ and $\hat{s}(\psi)$ interchangably. The next lemma and corollary allow us to define the set of axes for $\mathcal{H}_0$ independently of any Kolchin representative.

Lemma 4.6. *Suppose that $\mathcal{K}$ is an abelian Kolchin representative and that $\tilde{L}$ is a line with endpoints $P, Q \in C_\infty$. Suppose further that:*

- *$\tilde{f} : \Gamma \to \Gamma$ is a lift of some $f \in \mathcal{K}$*
- *$Fix(\hat{f})$ contains P, Q and at least one other point.*
- *$\tilde{f}$ does not fix the highest edge splitting vertices of $\tilde{L}$.*

Then there exists i and a covering translation T such that $T(\tilde{L}) = \tilde{L}_i$ and $\tilde{f} = T^{-1} s_i(f) T$.

Proof of Lemma 4.6. By Lemma 2.3, $Fix(\tilde{f}) \neq \emptyset$. Choose a non-trivial path $\tilde{\sigma} \subset \Gamma$ that intersects $Fix(\tilde{f})$ only in its initial vertex, say $\tilde{p}_1$, and that has its terminal vertex, say $\tilde{v}$, in $\tilde{L}$. (We do not exclude the possibility that $\tilde{\sigma} \subset \tilde{L}$.) The first edge of $\tilde{\sigma}$ is a lift $\tilde{E}_i'$ of some E_i. Let T be the covering translation that carries $\tilde{E}_i'$ to $\tilde{E}_i^*$. Since $\tilde{f}$ fixes the initial endpoint, $\tilde{p}_1$, but not the terminal endpoint, $\tilde{p}_2$, of $\tilde{E}_i'$, $u_{i,f}$ is non-trivial and $\tilde{L}_i$ is defined; let $\tilde{L}_i' = T^{-1}(\tilde{L}_i)$. Since $T^{-1}s_i(f)T$ and $\tilde{f}$ both fix $\tilde{p}_1$, they must be equal. In particular, $\hat{f}$ fixes the endpoints of $\tilde{L}_i'$.

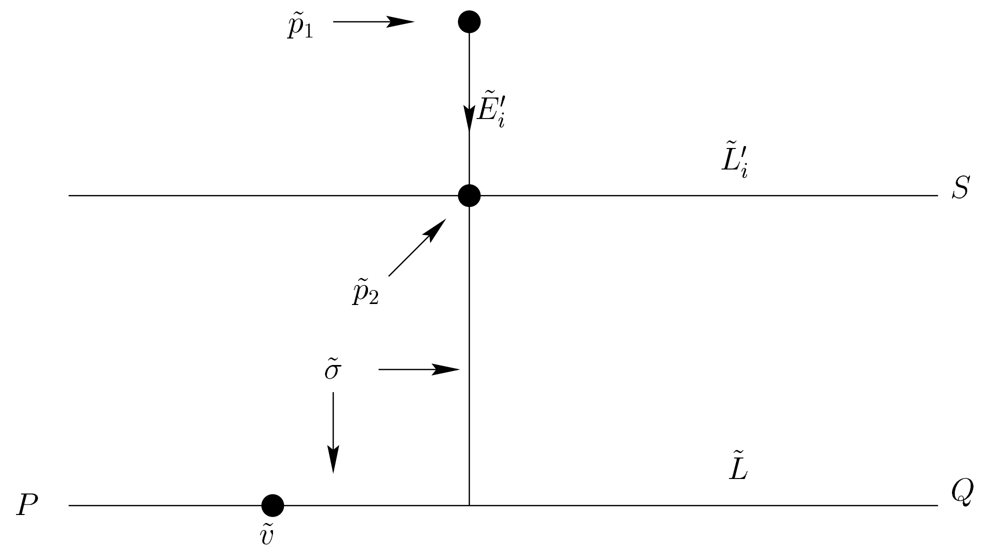

It suffices to show that $\tilde{L}_i' = \tilde{L}$. If not, then there is an endpoint, say S, of $\tilde{L}_i'$ that is neither P nor Q. Lemma 2.3 implies that the highest edge splitting vertices of either $\tilde{L}_{PS}$ (= the line connecting P to S) or $\tilde{L}_{SQ}$ are fixed. But $\tilde{L}_{PS}$ consists of a ray of $\tilde{L}$, a ray of $\tilde{L}_i'$ and perhaps a finite segment of $(\tilde{\sigma} \setminus \tilde{E}_i)$. The finite segment contains no fixed vertices. By construction, the highest edge splitting vertices of the two rays are not fixed. Thus the highest edge splitting vertices of $\tilde{L}_{PS}$ are not fixed. The symmetric argument applies to $\tilde{L}_{SQ}$ and yields the desired contradiction. $\square$

Corollary 4.7. *If $\mathcal{K}$ is an abelian Kolchin representative of $\mathcal{H}_0$, then:*

(1) *The axis of a covering translation T projects to an essential axis of $\mathcal{K}$ if and only if $IL(\psi, T)$ contains at least two elements for some $\psi \in \mathcal{H}_0$.*

(2) *For each $\alpha \in A(\mathcal{K})$ and each $\psi \in \mathcal{H}_0$, $IL(\psi, T^\alpha) = \{\hat{s}_i(\psi) : E_i$ is an essential edge associated to $\alpha\} \cup \{\hat{s}_\alpha(\psi)\}$*

Proof of Corollary 4.7. If E_i is an essential edge associated to $\alpha \in A(\mathcal{K})$ then there exists $f \in \mathcal{K}$ such that $u_i(f)$ is non-trivial. If $\psi \in \mathcal{H}_0$ is represented by f, then $\hat{s}_i(\psi) \neq \hat{s}_\alpha(\psi)$ and Lemma 4.5 implies that $IL(\psi, T^\alpha)$ contains at least two elements. For any covering translation T whose axis projects onto α, there is a covering translation S such that $T = ST^\alpha S^{-1}$. The only if part of (1) now follows from Lemma 4.4.

Conversely suppose that $\tilde{L}$ is the axis of a covering translation T and that for some $\psi \in \mathcal{H}_0$, $IL(\psi, T)$ contains more than one element. At least one of the elements in $IL(\psi, T)$ can be written as $\hat{f}$ where $\tilde{f} : \Gamma \to \Gamma$ does not fix the highest edge splitting vertices of $\tilde{L}$. Lemma 4.6 implies that $\tilde{L}$ differs by a covering translation from some $\tilde{L}_i$ and so projects to an essential axis. This completes the proof of (1).

Lemma 4.5 implies that $IL(\psi, T^\alpha) \supset \{\hat{s}_i(\psi) : E_i$ is an essential edge associated to $\alpha\} \cup \{\hat{s}_\alpha(\psi)\}$. For the reverse inclusion, suppose that $\tilde{f} : \Gamma \to \Gamma$ is a lift of $f \in \mathcal{K}$ such that $\hat{f} \in IL(\psi, T^\alpha)$. If $\tilde{f}$ fixes the highest edge splitting vertices of the axis of T^α then $\tilde{f} = s_\alpha(f)$. If not, then Lemma 4.6 (applied to the axis of T^α) implies that $\tilde{f} = T^{-1} s_i(f) T$ for some i and some covering translation T that carries the axis of T^α to the axis of T_i. Since $\mathcal{K}$ satisfies Property A, $T^\alpha = T_i$. Thus $T = (T^\alpha)^k$ for some k and $\tilde{f} = s_i(f)$ as desired. □

We are now able to define axes independently of Kolchin representatives.

Definition 4.8. The conjugacy class $[x]$ of a primitive element $x \in F_n$ is said to be *an axis for* $\mathcal{H}_0$ if $IL(\psi, x)$ contains more than one element for some element $\psi \in \mathcal{H}_0$. Lemma 4.4(1) implies that this definition depends on the conjugacy class $[x]$ and not on x itself. Define $A(\mathcal{H}_0)$ to be the set of axes for $\mathcal{H}_0$.

Given $a = [x] \in A(\mathcal{H}_0)$ and $\psi \in \mathcal{H}_0$, define $D_a(\psi) = \{d \in \mathbb{Z} : \hat{\Phi}_1 \hat{\Phi}_2^{-1} = \hat{T}_x^d$ for some $\hat{\Phi}_1, \hat{\Phi}_2 \in IL(\psi, T_x)\}$; equivalently $D_a(\psi) = \{d \in \mathbb{Z} : \Phi_1 \Phi_2^{-1} = i_x^d$ for some $\hat{\Phi}_1, \hat{\Phi}_2 \in IL(\psi, x)\}$. Lemma 4.4(1) implies that implies that $D_a(\psi)$ depends only on $[x]$ and not on x. Define $D(\psi)$ to be the union of the $D_a(\psi)$'s.

Corollary 4.9. *Suppose that $\mathcal{K}$ is an abelian Kolchin representative for $\mathcal{H}_0$, that $A(\mathcal{H}_0) = \{a_1 \dots a_m\}$ and that $f \in \mathcal{K}$ represents ψ. Then*

(1) *$A(\mathcal{K}) = \{\alpha_1, \dots, \alpha_m\}$ where α_i is the loop in G corresponding to a_i.*

(2) *For each essential edge E_i, $s_i(f) = (T^\alpha)^d s_\alpha(f)$ for some $d \in D(\psi)$.*

Proof of Corollary 4.9. This is an immediate consequence of Corollary 4.7 and the definitions. □

Lemma 4.10. *Every abelian UPG subgroup $\mathcal{H}_0$ is finitely generated.*

Proof of Lemma 4.10. Choose an increasing sequence $\mathcal{H}_1 \subset \mathcal{H}_2 \subset \dots$ of finitely generated subgroups whose union is $\mathcal{H}_0$. The cardinality of $A(\mathcal{H}_j)$ and the multiplicities of the elements of $A(\mathcal{H}_j)$ are uniformly bounded (by the maximum number of edges in a marked graph with the property that the terminal vertex of each edge has valence at least three). Since $A(\mathcal{H}_j) \subset A(\mathcal{H}_{j+1})$, we may assume after passing to a subsequence, that $A(\mathcal{H}_j)$ and the multiplicities are independent of j. In particular, there is a fixed r so that $\Phi_j : \mathcal{H}_j \to \mathbb{Z}^r$ where Φ_j is the embedding of Lemma 3.6 and where $\mathcal{H}_j$ has been identified with a Kolchin representative $\mathcal{K}_j$. After passing to a further sequence we may assume that the rank q of $\mathcal{H}_j$ is independent of j.

Fix $\psi \in \mathcal{H}_1$. Corollary 4.9 and the definition of Φ_j imply that the coordinates of $\Phi_j(\psi)$ are elements of $D(\psi)$. In particular, $\Phi_j(\psi)$ takes on only finitely many values as j varies.

Let $\psi_1, \ldots, \psi_q$ be generators of $\mathcal{H}_1$. After passing to a subsequence we may assume that $\Phi_j(\psi_i)$ is independent of j for $i = 1 \ldots q$. The lattice $\Phi_j(\mathcal{H}_1)$ is therefore independent of j. It is contained with finite index in a maximal lattice L of rank q. The lattice $\Phi_j(\mathcal{H}_j)$ has rank q and contains $\Phi_j(\mathcal{H}_1)$ so is contained in L. In particular, the index of $\mathcal{H}_1$ in $\mathcal{H}_j$ is uniformly bounded. It follows that $\mathcal{H}_j = \mathcal{H}_{j+1}$ for all sufficiently large j.

□

We close this section by showing that $1 \to \mathcal{H}_0 \to \mathcal{H} \xrightarrow{\Omega} \mathbb{Z}^b \to 1$ is a virtually central extension; the proof is a variation on that of Lemma 4.10.

Lemma 4.11. *$A(\mathcal{H}_0)$ is $\mathcal{H}$ invariant. There is a finite index subgroup of $\mathcal{H}$ whose actions on $\mathcal{H}_0$ and on $A(\mathcal{H}_0)$ are trivial.*

Proof of Lemma 4.11. Lemma 4.4(2) implies that $A(\mathcal{H}_0)$ is $\mathcal{H}$ invariant. Since $A(\mathcal{H}_0)$ is finite, there is a finite index subgroup of $\mathcal{H}$ for which the action on $A(\mathcal{H}_0)$ is trivial.

We assume now that the action of $\mathcal{H}$ on $A(\mathcal{H}_0)$ is trivial. Let $\mathcal{K}$ be an abelian Kolchin representative for $\mathcal{H}_0$ and let $\alpha \in A(\mathcal{K})$ correspond to $[x] \in A(\mathcal{H})$. Each $\eta \in \mathcal{H}$ is represented by an automorphism Φ that fixes x. Since conjugation by $\hat{\Phi}$ carries $IL(\psi, x)$ to $IL(\eta\psi\eta^{-1}, \Phi(x)) = IL(\eta\psi\eta^{-1}, x)$ and since Φ commutes with i_x, $D(\psi) = D(\eta\psi\eta^{-1})$. Let $\Phi_{\mathcal{K}} : \mathcal{H}_0 \to \mathbb{Z}^r$ be the embedding of Lemma 3.6. Since the coordinates of $\Phi_{\mathcal{K}}(\eta\psi\eta^{-1})$ are elements of $D(\eta\psi\eta^{-1}) = D(\psi)$, $\Phi_{\mathcal{K}}(\eta\psi\eta^{-1})$ takes on only finitely many values as η varies over $\mathcal{H}$ and $\psi \in \mathcal{H}_0$ is fixed. Since $\Phi_{\mathcal{K}}$ is an embedding, $\eta\psi\eta^{-1}$ takes on only finitely many values. After passing to a finite index subgroup we may assume that $\eta\psi\eta^{-1}$ is independent of η. In other words, the action of $\mathcal{H}$ by conjugation on ψ is virtually trivial. After applying this argument to a finite generating set for $\mathcal{H}_0$, we see that the action of $\mathcal{H}$ on $\mathcal{H}_0$ is virtually trivial. □

5 Proof of Proposition 1.2

To get the proof of Proposition 1.2 started we will need the following lemma.

Lemma 5.1. *If $n \geq 2$ and $\eta \in Out(F_n)$, then there exists $\Phi \in Aut(F_n)$ representing a positive power of η such that $\hat{\Phi}$ fixes at least three points. If x is a non-trivial primitive element of F_n whose conjugacy class is fixed by η, then we may choose Φ to fix x.*

The proof of Lemma 5.1 is a straightforward, but somewhat technical, application of relative train track maps. G. Levitt and M. Lustig give an alternate treatment of this lemma in Proposition I.5 of [LL99].

A relative train track map $f : G \to G$ is a homotopy equivalence of a marked graph with certain useful features. We recall here some facts about relative train track maps; complete details can be found in [BH92] and [BFH96a].

A direction is the germ of an oriented edge. There is a well defined self map Tf on the set of directions. Namely if E_i and E_j determine directions d_i and d_j and if E_j is the initial edge of $f(E_i)$ then $Tf(d_i) = d_j$. We say that a direction is fixed if it is fixed under the action of Tf. A turn is an unordered pair (d_1, d_2) of directions initiating at the same vertex. The turn is degenerate if $d_1 = d_2$ and non-degenerate otherwise. The turn is illegal if $(Tf)^k(d_1) = (Tf)^k(d_2)$ for some $k > 0$ and is legal otherwise.

A relative train track map $f : G \to G$ has a filtration by invariant subgraphs $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_K = G$. The subcomplex H_i is the union of the edges that are contained in G_i but not in G_{i-1} and is called the i^{th} stratum. There are two kinds of strata that contain periodic points; they are called exponentially growing and non-exponentially-growing.

If $\sigma \subset G$ is a loop and x is a point in the image of σ, then we can view σ as an immersed path σ_x with common initial and terminal endpoint at x.

Lemma 5.2. *Given $\eta \in Out(F_n)$, and perhaps given $x \in F_n$ whose conjugacy class is fixed by η, there exists a relative train track map $f : G \to G$ representing an iterate of η with the following properties:*

(i) *Every periodic vertex and periodic direction is fixed.*

(ii) *Every fixed vertex in a stratum H_i has at least one fixed direction in H_i.*

(iii) *Each non-exponentially-growing stratum H_i consists of a single E_i satisfying $f(E_i) = E_i u_i$ where u_i is a path in G_{i-1}. If u_i is trivial then E_i is a loop.*

(iv) *If x is given, then there is a lift $\tilde{f} : \Gamma \to \Gamma$ that commutes with T_x and that fixes a vertex $\tilde{v}$ in the axis of T_x.*

(v) *If E determines a fixed direction but is not a fixed loop, then $[f^k(E)]$ is a proper initial segment of $[f^{k+1}(E)]$ for all $k \geq 0$.*

(vi) *Every fixed vertex is either the endpoint of an exponentially growing edge or is the initial endpoint of at least two non-exponentially-growing edges.*

Proof of Lemma 5.2. Start with an improved relative train track map $f : G \to G$ representing an iterate of η provided by Theorem 5.1.5 of [BFH96a]. Property (v) follows from basic splitting properties: Lemma 4.2.1 of [BFH96a] for exponentially growing strata and condition (ne-ii) of Theorem 5.1.5 of [BFH96a] for non-exponentially-growing strata. If x is given, choose a lift $\tilde{f} : \Gamma \to \Gamma$ that commutes with T_x. The key to property (iv) is the fact the axis of T_x has a canonical splitting that is analogous to the highest edge splitting discussed earlier in the paper. This splitting is given by Lemma 4.2.6 of [BFH96a] if the highest stratum crossed by the axis is exponentially growing, and by condition (ne-iii) of Theorem 5.1.5 of [BFH96a] otherwise. The splitting is periodic so after replacing $\tilde{f}$ by some $T_x^k \tilde{f}^l$, each splitting vertex is fixed. Property (iii) is contained in conditions (ne-i) and (ne-ii) of Theorem 5.1.5 of [BFH96a]. Properties (i) and (ii) follow by taking iterates. Property (vi) is proved exactly as in Lemma 4.1. $\square$

Proof of Lemma 5.1. Let $f : G \to G$ be a relative train track map for an iterate of η as in Lemma 5.2. Suppose at first that x is given. Let $\tilde{f} : \Gamma \to \Gamma$ and $\tilde{v}$ be as in part (iv) of Lemma 5.2; denote the image of $\tilde{v}$ in G by v. There is at least one edge E with initial vertex v that determines a fixed direction. Let $\tilde{E}$ be the lift of E based at $\tilde{v}$.

If E is not a fixed loop, then by (v) of Lemma 5.2, the finite paths $\tilde{E} \subset [\tilde{f}(\tilde{E})] \subset [\tilde{f}^2(\tilde{E})] \subset \cdots \subset [\tilde{f}^k(\tilde{E})]$ form an increasing nested sequence whose union $\tilde{\tau}$ converges to a fixed point $Q \in C_\infty$ for $\hat{f}$. We claim that $\tilde{\tau}$ is not contained in the axis of any covering translation. It will follow that Q can not be an endpoint of the axis of T_x so $\hat{f}$ fixes at least three points. Let H_i be the stratum containing E. If E is non-exponentially-growing, then the claim follows from the fact that the first edge of $\tilde{\tau}$ projects to $E \subset H_i$ while the rest of $\tilde{\tau}$ projects into G_{i-1}. If E_i is exponentially growing and $\tilde{\tau}$ were contained in the axis of T_y, then the loop in G corresponding to y would be contained in G_i and would cross no illegal turns in H_i. No such loop can be fixed by the action $f_\#$ of f on loops in G (see for example Lemma 4.2.6 of [BFH96a]) in contradiction to the fact that $\hat{f}$ fixes an endpoint of the axis of T_y.

If each edge determining a fixed direction at v is a fixed loop, then condition (vi) of Lemma 5.2 implies that there are at least two fixed loops based at v. Each fixed loop lifts to a line through $\tilde{v}$ whose endpoints are fixed by $\hat{f}$. This completes the proof when x is given.

Suppose now that x is not given. If η acts periodically on some conjugacy class in F_n, then some iterate of η fixes a conjugacy class and we can choose $x \in F_n$ to be any element in that conjugacy class. We are then reduced to the previous case. If there are no periodic conjugacy classes, then the first stratum must be exponentially growing. By Lemma 2.2, we may assume that $G = G_1$ is a single exponentially growing stratum.

It suffices to show that some fixed vertex v in G has at least three fixed directions. Indeed, we can then choose $\tilde{f}$ to fix $\tilde{v}$ and produce fixed points Q_1, Q_2 and Q_3 for $\hat{f}$ as we did above using $\tilde{\tau}$. By condition (i) of Lemma 5.2, it suffices to show that there is a vertex w and directions d_1, d_2 and d_3 at w such that $(Tf)^k(d_1), (Tf)^k(d_1)$ and $(Tf)^k(d_3)$ are distinct for all k. In other words, it suffices to show that some vertex is contained in at least two legal turns. If this were not the case, then there would be only finitely many non-repeating 'legal' loops in G (i.e. loops that cross only legal turns). But the image of a non-repeating legal loop is a non-repeating legal loop and η does not act periodicially on any conjugacy class. This contradiction completes the proof. $\square$

Proof of Proposition 1.2. We will use the following equivalent statements interchangably: (i) $\mu \in Out(F_n)$ fixes the conjugacy class of x; (ii) μ is represented by $\Theta \in Aut(F_n)$ that fixes x; (iii) μ is represented by Θ such that $\hat{\Theta}$ commutes with $\hat{T}_x$. We will also use the fact that $\hat{i}_x = \hat{T}_x$.

We argue by induction on the rank of the free abelian group $\mathcal{H}$. If $\mathcal{H}$ has rank one with generator μ, then $\tilde{\mathcal{H}}$ is determined by a choice of $\Theta \in Aut(F_n)$ representing μ that fixes x if x is given. We now assume that the rank of $\mathcal{H}$ is at least two and that the lemma holds for all ranks less than that of $\mathcal{H}$.

By Lemma 5.1, there is an element $\eta \in \mathcal{H}$ and $\Phi \in Aut(F_n)$ representing η such that $\hat{\Phi}$ fixes at least three points and such that Φ fixes x if x is given. In Steps 1 and 2 below we show that, after passing to a subgroup of $\mathcal{H}$ with finite index if necessary, every $\mu \in \mathcal{H}$ is represented by $\Theta \in Aut(F_n)$ that commutes with Φ and that fixes x if x is given.

(Step 1: Choosing Θ when x is not given). Suppose at first that x is not given. Let $IL(\eta) = \{\Psi \in Aut(F_n) : \Psi$ represents η and $\hat{\Psi}$ fixes at least three points$\}$. There is an equivalence relation on elements of $IL(\eta)$ defined by $\Psi_1 \sim \Psi_2$ if and only $\Psi_1 = i_y \Psi_2 i_y^{-1}$ for some $y \in F_n$. Since $\mathcal{H}$ is abelian, conjugation by any automorphism Θ representing $\mu \in \mathcal{H}$ preserves $IL(\eta)$. It is immediate from the definitions that conjugation by Θ permutes the equivalence classes of $IL(\eta)$ and that this permutation is independent of the choice of Θ representing μ. In other words, $\mathcal{H}$ acts by conjugation on the set of equivalence classes in $IL(\eta)$.

Lemma 2.4 implies that there are only finitely many equivalence classes in $IL(\eta)$. We may therefore assume, after passing to a finite index subgroup, that conjugation defines a trivial action of $\mathcal{H}$ on the equivalence classes in $IL(\eta)$. Applying this to the equivalence class containing Φ proves that each μ is represented by $\Theta \in Aut(F_n)$ such that $\Theta\Phi\Theta^{-1} = \Phi$.

(Step 2: Choosing Θ when x is given) Suppose now that x is given and that Θ represents μ and fixes x. Since both Φ and $\Phi_1 = \Theta\Phi\Theta^{-1}$ represent η and fix x, $\Phi_1 = i_x^k \Phi$ for some $k \in \mathbb{Z}$. It suffices to show that $k = 0$. Suppose not. Arguing by induction on m, we see that $\Theta^m \Phi \Theta^{-m} = i_x^{mk}\Phi$ for all m. Thus $\text{Fix}(\hat{T}_x^{mk}\hat{\Phi}) = \text{Fix}(\hat{\Theta}^m \hat{\Phi} \hat{\Theta}^{-m}) = \hat{\Theta}^m(\text{Fix}(\hat{\Phi}))$ contains a point, say R_m, that is neither the backward endpoint P nor the forward endpoint Q of the axis of T_x. Since $\hat{T}_x$ acts cocompactly on $C_\infty \setminus \{P, Q\}$ and since $\text{Fix}(\hat{T}_x^{mk}\hat{\Phi})$ is $\hat{T}_x$-invariant, we may assume, after passing to a subsequence, that the R_m's converge to some point $R \in C_\infty \setminus \{P, Q\}$. Choose a neighborhood V of Q in C_∞ that is disjoint from R and hence disjoint from R_m for all sufficiently large m. There is a neighborhood U of $\hat{\Phi}(R)$ in C_∞ such that $\hat{T}_x^l(U) \subset V$ for all sufficiently large l. This contradicts the fact that $T_x^{mk}(\hat{\Phi}(R_m)) = R_m$ for all m and we conclude that $k = 0$ as desired.

Denote $\text{Fix}(\Phi)$ by $\mathbb{F}$, $\text{Fix}(\hat{\Phi})$ by C and the group of covering translations that preserve C by $\mathcal{T}(C)$. Recall from Lemma 2.1 that $\mathcal{T}(C)$ is also the set of covering translations that commute with Φ. Since each $\mu \in \mathcal{H}$ is represented by an automorphism Θ that commutes with Φ and fixes x if x is given, $\mathbb{F}$ is $\mathcal{H}$-invariant, up to conjugacy, and the conjugacy class of x in $\mathbb{F}$ is $\mathcal{H}|\mathbb{F}$-invariant, if x is given.

(Step 3: The case that $\mathcal{T}(C)$ is the trivial group). If $\mathcal{T}(C)$ is the trivial group, then (Lemma 2.1) each μ is represented by a unique Θ that commutes with Φ and the assignment $\mu \mapsto \Theta$ defines $\tilde{\mathcal{H}}$.

(Step 4: The case that $\mathcal{T}(C)$ has rank one). Suppose next that $\mathcal{T}(C)$ has rank one. Since $\mathbb{F}$ is $\mathcal{H}$-invariant, up to conjugacy, we may assume that x is given and that T_x generates $\mathcal{T}(C)$.

By the main theorem of [Coo87], C contains only finitely many $\hat{T}_x$-orbits. The argument in this special case is short so we include it for completeness. Since T_x acts co-compactly on $C_\infty \setminus \{P, Q\}$, it suffices to show that if S is an accumulation point of C, then $S = P$ or Q. Arguing as in the proof of Lemma 2.4, using triples of points in C limiting on S, we find lifts $\tilde{v}_i \in \Gamma$ of a vertex $v \in G$ so that if T_i is the covering translation that carries $\tilde{v}_1$ to $\tilde{v}_i$, then $\hat{T}_i$ commutes with $\hat{\Phi}$. But then $T_i = T_x^{k_i}$ with $k_i \to \infty$ (or $-\infty$). Since $T_i(\tilde{v}_1) \to S$, S must be an endpoint of the axis of T_x.

Given $\mu \in H$, choose Θ representing μ that commutes with Φ and that fixes x. Then $\hat{\Theta}$ preserves C and commutes with $\hat{T}_x$ so conjugation by $\hat{\Theta}$ permutes the $\hat{T}_x$ orbits of C. Since different choices of $\hat{\Theta}$ differ only by composition with some iterate of $\hat{T}_x$, this permutation of the $\hat{T}_x$ orbits of C depends only on μ. In other words, conjugation by Θ defines an action of $\mathcal{H}$ on the $\hat{T}_x$ orbits of C. After passing to a finite index subgroup, we may assume that the action is trivial. Choose a point $R \in C \setminus \{P, Q\}$. After replacing $\hat{\Theta}$ with some $\hat{T}_x^k \hat{\Theta}$, we may assume that $\hat{\Theta}$ fixes R. The assignment $\mu \mapsto \Theta$ defines $\tilde{\mathcal{H}}$.

(Step 5: The case that $\mathcal{T}(C)$ has rank at least two). Finally suppose that $\mathcal{T}(C)$ has rank at least two. Denote $\mathcal{H}|\mathbb{F}$ by $\mathcal{H}^*$. Since $\mathbb{F}$ is the fixed subgroup of an automorphism Φ representing η, the image of η in $\mathcal{H}^*$ is trivial and the rank of $\mathcal{H}^*$ is less than that of $\mathcal{H}$. By the inductive hypothesis, there is a lift $\widetilde{\mathcal{H}^*} \subset Aut(\mathbb{F})$ to elements that fix x if x is given. By Lemma 2.2, $\widetilde{\mathcal{H}^*}$ extends to the desired lift $\tilde{\mathcal{H}}$. $\square$

6 Proof of Theorem 1.1

Let us first dispose of the index bound in the statement of Theorem 1.1. Suppose that $\mathcal{H}$ is a finitely generated, virtually free abelian subgroup of $Out(F_n)$. By passing to a subgroup of index at most $D(n)$ where $D(n) := |GL(n, \mathbb{Z}/3\mathbb{Z})| < 3^{n^2}$, we may assume that the image of $\mathcal{H}$ in $GL(n, \mathbb{Z}/3\mathbb{Z})$ is trivial. Thus, $\mathcal{H}$ has a subgroup of index at most this number that is a torsion free Bieberbach group of vcd at most $vcd(Out(F_n)) = 2n - 3$ (see [CV86]). Also, a Bieberbach group of vcd at most k has a subgroup of index at most $D(k)$ that is free abelian (see, for example, [Cha86]). Thus, $\mathcal{H}$ has a free abelian subgroup of index at most $D(n)D(2n-3) < 3^{5n^2}$.

It remains now to show that a solvable subgroup $\mathcal{H}$ of $Out(F_n)$ is finitely generated and virtually free abelian. After passing to a subgroup of finite index if necessary, there is (Theorem 7.0.1 of [BFH96a]) an exact sequence

$$1 \to \mathcal{H}_0 \to \mathcal{H} \stackrel{\Omega}{\to} \mathbb{Z}^b \to 1$$

where $\mathcal{H}_0$ is UPG. We may assume without loss that $\mathcal{H}_0$ is non-trivial and, by Lemma 4.11, that $\mathcal{H}$ acts trivially on both $\mathcal{H}_0$ and on $A(\mathcal{H}_0)$.

Let $\mathcal{K}$ be an abelian Kolchin representative for $\mathcal{H}_0$. Throughout this section $[x]$ is an element of $A(\mathcal{H}_0)$, α is the corresponding element of $A(\mathcal{K})$, $s = s_i$ or s_α and $T^\alpha = T_x$. We will take advangate of the isomorphism between $\mathcal{H}_0$ and $\mathcal{K}$ to view $\hat{s}$ as a lift of $\mathcal{H}_0$ and to write $\hat{s}(\eta)$ for $\eta \in \mathcal{H}_0$. We will also take advantage of the isomorphism between the group of lifts of elements of $\mathcal{K}$ and the group of automorphisms of F_n representing elements of $\mathcal{H}_0$ to view $s(\eta)$ as an element of $Aut(F_n)$; in this context we will write $\tilde{s}(\eta) \in Aut(F_n)$.

Lemma 6.1. *The lift $\tilde{s}(\mathcal{H}_0) \subset Aut(F_n)$ virtually extends to a lift $\tilde{S}(\mathcal{H}) \subset Aut(F_n)$ such that $\tilde{S}(\mu)$ fixes x for each $\mu \in \mathcal{H}_0$.*

Proof of Lemma 6.1. Denote $\bigcap_{\eta\in\mathcal{H}_0} Fix(\tilde{s}(\eta))$ by $\mathbb{F}$, $\bigcap_{\eta\in\mathcal{H}_0} Fix(\hat{s}(\eta))$ by C, and the group of covering translations that preserve C by $\mathcal{T}(C)$. Since T^α commutes with each $s(\eta)$, $\mathcal{T}(C)$ contains T^α and so has rank at least one.

Since $\mathcal{H}$ acts trivially on $A(\mathcal{H}_0)$, each $\mu \in \mathcal{H}$ is represented by an automorphism Θ that fixes x. Since $\mathcal{H}$ acts trivially on $\mathcal{H}_0$, $\Theta\tilde{s}(\eta)\Theta^{-1}$ represents η for each $\eta \in \mathcal{H}_0$. Both $\Theta\tilde{s}(\eta)\Theta^{-1}$ and $\tilde{s}(\eta)$ fix x, so $\Theta\tilde{s}(\eta)\Theta^{-1} = i_x^k\tilde{s}(\eta)$ for some $k \in \mathbb{Z}$. Arguing exactly as in Step 2 of the proof of Proposition 1.2, we see that $k = 0$ and hence that Θ commutes with $\tilde{s}(\eta)$. Thus C is $\hat{\Theta}$-invariant, $\mathbb{F}$ is $\mathcal{H}$-invariant and the conjugacy class of x in $\mathbb{F}$ is $\mathcal{H}|\mathbb{F}$-invariant.

The proof now divides into cases, depending on the rank of $\mathcal{T}(C)$. Suppose that $\mathcal{T}(C)$ is generated by T^α. Lemma 2.3 implies that if $Q \in C_\infty$ is an accumulation point of C, then there is a vertex $w \in G$ and lifts $\tilde{w}_i \in \Gamma$ of w that are fixed by each $s(\eta)$ and that limit on Q. As in Step 4 of the proof of Proposition 1.2, we conclude that Q is an endpoint of the axis of T^α and that C contains only finitely many $\hat{T}_x$-orbits. Continuing as in Step 4 of the proof of Proposition 1.2, we see that there is a well defined action of $\mathcal{H}$ on the finitely many $\hat{T}_x$-orbits of C. After passing to a finite index subgroup, this action is trivial. By Lemma 4.5 there exists $R \in C_\infty$ that is not an endpoint of the axis of T^α. Composing $\hat{\Theta}$ with an iterate of $\hat{T}^\alpha$, we may assume that $\hat{\Theta}$ fixes R. The assignment $\eta \mapsto \Theta$ defines $\tilde{S}$.

We may now assume that $\mathcal{T}(C)$ has rank at least two. By Lemma 2.2, the restriction $\mathcal{H}^* = \mathcal{H}|\mathbb{F}$ is well defined. Since $\tilde{s}(\eta)|\mathbb{F}$ is the identity for each $\eta \in \mathcal{H}_0$, $\mathcal{H}_0|\mathbb{F}$ is trivial. The restriction homomorphism $R : \mathcal{H} \to \mathcal{H}^*$ therefore factors through the abelian group $\mathcal{H}/\mathcal{H}_0$ proving that $\mathcal{H}^*$ is abelian. By Propositon 1.2, there is a virtual lift $\widetilde{S^*} \subset Aut(\mathbb{F})$ of $\mathcal{H}^*$ such that each element of $\widetilde{S^*}$ fixes x. Let $\tilde{S}$ be the unique extension of $\widetilde{S^*}$ to a virtual lift of $\mathcal{H}$ given by Lemma 2.2. Since $\mathcal{H}_0|\mathbb{F}$ is trivial, $\tilde{S}(\eta)|\mathbb{F}$ is the identity for each $\eta \in \mathcal{H}_0$. Thus $\tilde{S}(\eta)$ and $\tilde{s}(\eta)$ agree on $\mathbb{F}$ and so must agree everywhere. $\square$

It is now a simple matter to complete the proof of Theorem 1.1. We may assume that the extensions $\tilde{S}_i$ and $\tilde{S}_\alpha$ of $\tilde{s}_i$ and $\tilde{s}_\alpha$ produced by Lemma 6.1 are defined on all of $\mathcal{H}$. Define a homomorphism $\phi_i' : \mathcal{H} \to \mathbb{Z}$ by $\tilde{S}_i(\psi) = i_x^{\phi_i'(\psi)} \tilde{S}_\alpha(\psi)$ and note that the product $\Phi' : \mathcal{H} \to \mathbb{Z}^r$ of the ϕ_i''s extends the embedding $\Phi : \mathcal{H}_0 \to \mathbb{Z}^r$ of Lemma 3.6. Define $\Psi : \mathcal{H} \to \mathbb{Z}^{b+r}$ to be the product of Ω and Φ'. Since $\mathcal{H}_0$ is the kernel of Ω and $\Phi'|\mathcal{H}_0 = \Phi$ is an embedding, Ψ is an embedding. $\square$

References

[BFH96a] M. Bestvina, M. Feighn, and M. Handel, *The Tits Alternative for Out(F_n) I: Dynamics of Exponentially Growing Automorphisms*, preprint, 1996.

[BFH96b] M. Bestvina, M. Feighn, and M. Handel, *The Tits Alternative for Out(F_n) II: a Kolchin Theorem*, preprint, 1996.

[BH92] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, Annals of Math., vol 135, 1992, pp. 1–51

[BL94] H. Bass and A. Lubotzky, *Linear-central filtrations on groups*, The mathematical legacy of Wilhelm Magnus: groups, geometry and special functions (W. Abikoff, J. Birman, and K. Kuiken, eds.), Contemp. Math., vol. 169, American Mathematical Society, 1994, pp. 45–98.

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

[Cha86] L. S. Charlap, *Bieberbach groups and flat manifolds*, Springer-Verlag, 1986.

[Coo87] D. Cooper, *Automorphisms of free groups have finitely generated fixed point sets*, J. Algebra **111** (1987), 453–456.

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

[LL99] G. Levitt and M. Lustig, *Periodic ends, growth rates, Hölder dynamics for automorphisms of free groups*, preprint

[HT85] M. Handel and W.P. Thurston, *New proofs of some results of Nielsen*, Adv. in Math. **56** (1985), 173–191.