5-2020

Hyperbolic Endomorphisms of Free Groups

Jean Pierre Mutanguha
University of Arkansas, Fayetteville

Hyperbolic Endomorphisms of Free Groups

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy in Mathematics

by

Jean Pierre Mutanguha
Oklahoma Christian University
Bachelor of Science in Mathematics, 2014
University of Arkansas
Master of Science in Mathematics, 2018

May 2020
University of Arkansas

This dissertation is approved for recommendation to the Graduate Council.

Matt Clay, Ph.D.
Dissertation Director

Matt Day, Ph.D.
Committee Member

Yo'av Rieck, Ph.D.
Committee Member

Abstract

We prove that ascending HNN extensions of free groups are word-hyperbolic if and only if they have no Baumslag-Solitar subgroups. This extends Brinkmann's theorem that free-by-cyclic groups are word-hyperbolic if and only if they have no $\mathbb{Z}^2$ subgroups. To get started on our main theorem, we first prove a structure theorem for injective but nonsurjective endomorphisms of free groups. With the decomposition of the free group given by this structure theorem, we (more or less) construct representatives for nonsurjective endomorphisms that are expanding immersions relative to a homotopy equivalence. This structure theorem initializes the development of (relative) train track theory for nonsurjective endomorphisms.

Acknowledgements

I would like to thank my family and friends for the love and support they have always showered me with, especially in the last few years of graduate school. Mom and dad, thank you for believing in me and my dream to be a mathematician.

This dissertation would not have been possible without the guidance of my advisor Matt Clay. Matter of fact, the main sources I used in this dissertation (the appendices) were the first three readings we did together! There have been so many moments when I have doubted myself and he has always been understanding and encouraging. I would also like to thank the rest of my committee: Matt Day and Yo'av Rieck. They have all introduced me to so much interesting mathematics that will fuel my research for the rest of my career. In addition to my committee, I am grateful to Ilya Kapovich and Derrick Wigglesworth for the long conversations we had on the subject of my dissertation; I had a lot of outlandish ideas and their patient probing helped me hone in on the not-wrong ones.

Finally, I thank my department — it goes without saying that I would not be where I am today without their support.

Table of Contents

1 Introduction

One of the main themes in *Geometric Group Theory* is understanding how algebraic properties of groups determine the geometric properties of spaces they act nicely on and vice-versa. For example, any group acting geometrically (properly and cocompactly) on hyperbolic space $\mathbb{H}^n$ with $n \geq 2$ will have a finite presentation, a solvable word-problem, both Hopfian/co-Hopfian properties, and, if $n \geq 3$, finite outer automorphism groups. Since acting geometrically on hyperbolic space guarantees various useful algebraic properties, an obvious (but hard) question to ask is: what are the necessary and sufficient conditions for a group to act geometrically on hyperbolic space?

This question becomes tractable with more hypotheses. For instance, the fundamental group of a surface will act geometrically on $\mathbb{H}^2$ if and only if the surface is closed and has negative Euler characteristic. Going up a dimension to $\mathbb{H}^3$ leads to Thurston's geometrization conjecture but we only highlight one of its components: Thurston's hyperbolization theorem for 3-manifolds that fiber over the circle [15]. Succinctly, one instance of the theorem states that the mapping torus (defined below) of a closed surface's homeomorphism will have a fundamental group that acts geometrically on $\mathbb{H}^3$ if and only if the homeomorphism has no periodic homotopy classes of essential simple closed curves, or equivalently, the fundamental group has no $\mathbb{Z}^2$ subgroups.

This dissertation proves a *combinatorial* analogue of Thurston's thoerem, a dimension lower. δ-Hyperbolic spaces are geodesic spaces whose geodesic triangles are all *δ-thin*, a feature of $\mathbb{H}^n$ that accounts for many of the algebraic properties groups inherit from acting geometrically on $\mathbb{H}^n$. Groups that act geometrically on proper δ-hyperbolic spaces are called **word-hyperbolic groups**. We find necessary and sufficient conditions for a mapping torus of a π_1-injective graph map to have a word-hyperbolic fundamental group.

Main Theorem. *If $f : \Gamma \to \Gamma$ is a π_1-injective graph map and M_f is its mapping torus, then $\pi_1(M_f)$ is word-hyperbolic if and only if $\pi_1(M_f)$ has no $BS(1,d)$ subgroups for $d \geq 1$.*

The complete theorem has more equivalent statements but we postpone giving the full result until the relevant definitions and context is given. The **mapping torus** of a map $g : X \to X$ is defined to be the quotient space $M_g = (X \times [0,1]) /{\sim_g}$ where identification is given by $(x,1) {\sim_g} (g(x),0)$ for all $x \in X$. For $d \geq 1$, the **(metabelian) Baumslag-Solitar** subgroups are $BS(1,d) = \pi_1(M_d)$ where $d : S^1 \to S^1$ is a degree d map. It is a well-known fact that groups with $BS(1,d)$ subgroups are not word-hyperbolic and so the main theorem is the converse of this fact for mapping tori of graph maps.

In our setting, $F = \pi_1(\Gamma)$ is a free group of finite rank and the map $f : \Gamma \to \Gamma$ induces an injective endomorphism $\phi : F \to F$. The fundamental group of M_f is denoted by $F{*_\phi}$ and known as an **ascending HNN extension of** $\boldsymbol{F}$. The main theorem can now be restated as *"$F{*_\phi}$ is word-hyperbolic if and only if it has no $BS(1,d)$ subgroups for $d \geq 1$."*

Peter Brinkmann [4] proved the case of this theorem when ϕ is an automorphism. He uses the Bestvina-Feighn combination theorem [1] and improved relative train tracks to show word-hyperbolicity of the mapping tori when their fundamental groups have no $BS(1,1) = \mathbb{Z}^2$ subgroups.

Extending Brinkmann's theorem to all injective endomorphisms of F remained diffucult for two reasons. When the injective endomorphism $\phi : F \to F$ is not surjective, then

1. we do not know if *improved relative train tracks* exist; and

2. we have to somehow deal with $BS(1,d)$ subgroups where $d \geq 2$.

To complicate matters further, even if we assume improved relative train tracks exist, it is unclear how their properties could help us deal with nonabelian $BS(1,d)$ subgroups. For example, knowing an endomorphism has an irreducible train track (the best possible improved relative train track) is still not enough to address the second part.

It turns out, assuming the endomorphism is induced by a graph immersion (a seemingly strong assumption a priori) is enough to address both difficulties. In the context of this theorem, being an immersion is stronger than being a train track and, hence, the first difficulty is averted. Furthermore, the fact that edges in loops persist under iteration of immersions, i.e., no cancellation occurs, becomes crucial to accounting for nonabelian $BS(1,d)$ subgroups.

Ilya Kapovich proved the special case of the theorem under the extra assumption that the graph map is an immersion on the rose [6]. In previous work [8], we generalized Kapovich's argument to work for all graph immersions. Thus prior to this dissertation, the theorem was known for two extreme cases: ϕ is an automorphism or ϕ is induced by an immersion. Since the persistence of expanding edges in loops under iteration of an immersion plays a crucial role in how we deal with nonabelian $BS(1,d)$ subgroups, we are naturally led to three interesting questions:

1. which nonsurjective endomorphisms cannot be induced by expanding graph immersions?
2. can they instead be induced by graph maps that share the same persistence property with immersions, i.e., graph maps that are very close to being immersions themselves?
3. will such graph maps be enough to address the two difficulities listed above?

The first two questions are interesting in their own right and, in a sense, the original contribution of this dissertation is answering them.

Patrick Reynolds [12] made some progress in answering the first question. He proved that irreducible nonsurjective endomorphisms were induced by expanding graph immersions. Being irreducible means having no invariant proper free factor system and the precise definition will be given in the next section. Irreducible endomorphisms ought to be generic (intuitively speaking), so our previous generalization of Kapovich's theorem applies to "most" nonsurjective endomorphisms and the assumption that our graph maps are immersions is not as strong as it seemed! Furthermore, our generalization was a definite improvement on Kapovich's theorem since "most" nonsurjective endomorphisms are induced by expanding immersions *not* on roses but on trivalent graphs. Reynolds' proof involved tools like graph of actions and Q-maps for $\mathbb{R}$-tree with dense orbits. We later gave a different proof [9] of his result using *mostly* the semi-action of endomorphisms on the *spine of Culler-Vogtmann's outer space*. Part of the work done here was improving this idea to *only* use the action on the spine, which sets us up for *relativizing* the argument.

We will now summarize our complete answers to the first two questions above. The order we answer the questions in this introduction follows the progression ideas from the paper [9]. However, the actual proof given in the later sections answers them backwards (or rather simultaneously). Reynolds' theorem is a partial answer to the first question: if $\phi : F \to F$ is injective, nonsurjective, and cannot be induced by an expanding immersion, then it is *reducible*, i.e., it has an invariant free factor system. But it is easy to construct reducible endomorphisms that are induced by expanding immersions. On the other hand, obvious examples of endomorphisms not induced by expanding immersions are automorphisms. By passing to covers, it is not hard to convince ourselves that ϕ cannot be induced by an expanding immersion if it has a fixed free factor on which it acts as an automorphism. So the naive answer to the first question is the converse: *"If ϕ cannot be induced by an expanding immersion, then it has a fixed free factor system."*

Indeed, we give two equivalent characterizations of when an injective endomorphism is induced by an expanding immersion:

Corollary 6.6. *Let $\phi : F \to F$ be an injective endomorphism. Then the following conditions are equivalent:*

1. *ϕ is induced by an expanding graph immersion;*
2. *for all nontrivial $g \in F$, there is a $k = k(g)$ such that $[\phi^k(x)] \neq [g]$ for all $x \in F$;*
3. *ϕ has no nontrivial fixed free factor system.*

Condition (2) is equivalent to Reynolds' definition of ϕ being *expansive*. If some nontrivial conjugacy class $[g]$ witnesses the failure of (2), we will say $[g]$ *has an infinite tail.*

Implications (1) $\implies$ (2) $\implies$ (3) are easy. We sketch the original idea for (3) $\implies$ (1). Fix a maximal filtration and as the base case of our induction, we apply Reynolds' theorem to get an expanding immersion for the bottom stratum. At each induction step, we collapse the lower stratum and, using nonsurjectivity again, produce an *relative immersion* of the current stratrum, then we patch this relative immersion with the expanding immersion of the lower stratum to get an expanding immersion of the current graph. Once we are at the top of the filtration, we have an expanding immersion that induces the endomorphism. Of course, more needs to be said as nonsurjectivity is not enough to complete the induction step. Nevertheless, this strategy suggests the answer for the second question:

Theorem 6.5. *If $\phi : F \to F$ is injective and nonsurjective, then it is induced by an expanding relative immersion.*

Loosely, a relative immersion is a map of a *free splitting* of F that induces an immersion on the corresponding covering *Bass-Serre* trees. Again, precise definitions appear later. The proof of the answer is essentially the same induction described for the first answer. We first collapse a maximal fixed factor system and use nonsurjectivity to prove a variant of Reynolds' theorem for the base case: that there is an expanding relative immersion for the bottom stratum. Presumably, we know how to construct relative immersions since we needed them in the induction step in the previous argument. So once the base case is covered, we can invoke the previous induction step and be done.

The difficulty in proving the above theorem lies in collapsing the "right" fixed free factor system at the base and determining its properties that will be needed for the construction of relative immersions. Conveniently, there is a canonical choice of a maximal fixed free factor system to collapse first and it has the properties neded to complete the base case and induction step of Theorem 6.5.

Proposition 3.3. *Let $\phi : F \to F$ be an injective endomorphism. Then ϕ has a unique* *maximal fixed free factor system $\mathcal{A}$. Precisely, $\mathcal{A}$ supports every conjugacy class with an infinite tail as well as every finitely generated periodic subgroup of F.*

This is the one-line summary of the proof of the proposition: repeatedly apply the *bounded cancellation lemma* to growing Stallings graphs and find a stable/nonexpanding subgraph that represents a maximal fixed free factor system. The property that the chosen fixed free factor system supports all conjugacy classes with infinite tails is what allows us to complete the proofs of Collorary 6.6 and Theorem 6.5. So, for the corollary, we actually prove (3) $\Longrightarrow$ (2) and then (3)+(2) $\Longrightarrow$ (1).

To be a bit more precise, we collapse more than the fixed free factor system. The collapsed invariant systems must also be invariant under backward iteration. Thus we take preimages to get what we call the *elliptic free factor system*.

By now, we have completely answered the first two questions and the third question can be rephrased as: will expanding relative immersions allow us to prove the main theorem? The answer is: *"Yes*. (*Terms and conditions apply.)"* Simply having expanding relative immersions is not enough to understand all the dynamics because we are missing what happens in the collapsed free factor system; however, if they are expanding immersions relative to the elliptic free factor system, then the restriction of the endomorphism to the collapsed system is (almost) an automorphism and we have a much better understanding of those dynamics. In summary, we get to state a more detailed version of the main theorem:

Theorem 9.6. *Let $\phi : F \to F$ be an injective endomorphism. Then the following statements are equivalent:*

1. *$F*_\phi$ is word-hyperbolic;*
2. *$F*_\phi$ contains no $BS(1, d)$ subgroups, $d \geq 1$;*
3. *No iterate ϕ^k has an invariant cyclic subgroup (with index $d \geq 1$) for all $k \geq 1$.*

We have omitted a fourth equivalent condition as it is rather technical for the introduction.

Once again, the content of the theorem is the implication (3) $\implies$ (1) and we will briefly give the idea behind it. Start with an expanding immersion relative to the elliptic free factor system and on the elliptic system choose an improved relative train track. Together, the hybrid of an expanding immersion and an improved relative train track addresses the first difficulty raised earlier (Proposition 9.5). To address the second difficulty, i.e., to deal with invariant cyclic subgroups with index $d \geq 2$, we need only concern ourselves with the dynamics of the relative immersion as such subgroups do not exist in the elliptic free factor system. Fortunately, relative immersions do have the persistence property that we needed to generalize Kapovich's theorem (Proposition 8.5).

We believe the relative immersions we construct in this dissertation will prove to be extremely useful in understanding the dynamics of free group endomorphisms. For example, it seems like our techniques can be improved to give a trichotomy of the *Dehn functions* of ascending HNN extensions and characterising exactly when each happen:

1. linear $\iff$ no $BS(1,d)$ subgroups for $d \geq 1$ (Theorem 9.6);
2. quadratic $\iff$ has $\mathbb{Z}^2$ but no $BS(1,d)$ subgroups for $d \geq 2$;
3. exponential $\iff$ has $BS(1,d)$ subgroup for some $d \geq 2$;

The first equivalence is a reformulation of Theorem 9.6 since groups are word-hyperbolic if and only if they have linear Dehn functions. We conclude the introduction with an outline of the dissertation.

Outline. The next section introduces most of the definitions and preliminary results, most importantly the bounded cancellation lemma. Section 3 proves the existence of a unique maximal fixed free factor system. Section 5 develops the theory of relative representatives. This is done to set-up the construction of relative immersion in Section 6. Section 8 introduces pullbacks of an endomorphism and their relation to nonabelian $BS(1,d)$ subgroups and the final Section 9 proves the main theorem. There are interludes following Sections 3 and 6 that remind us of the big picture and one provides some useful examples.

2 Preliminaries

Unless otherwise stated, F will be a free group with finite rank at least 2. A **nontrivial free factor system** of F is a collection of nontrivial free factors $\mathcal{A} = \{A_1, \ldots, A_l\}$ of F such that $A_i \cap A_j$ is trivial if $i \neq j$ and $\langle A_1, \ldots, A_l \rangle$ is free factor of F. We define the **trivial free factor system** to be the collection consisting of the trivial subgroup; this will allow us to treat *relative/absolute* cases simultaneously in our proofs. A free factor system $\mathcal{A}$ is **proper** if $\mathcal{A} \neq \{F\}$, i.e., some free factor $A \in \mathcal{A}$ is proper. We shall say a free factor system $\mathcal{A}$ supports a collection of subgroups $\mathcal{B}$ if for every $B \in \mathcal{B}$, there is an $A \in \mathcal{A}$ and $x \in F$ such that $B \leq xAx^{-1}$ and, when $\mathcal{B}$ is a free factor system, we shall denoted it by $\mathcal{B} \preceq \mathcal{A}$. The $\preceq$-relation on free factor systems is a preorder, i.e., it is reflexive and transitive; it determines an equivalence relation on the set of free factor systems and a partial order on the set of equivalence classes. Free factor systems will always be considered up to this equivalence relation. In particular, we are free to replace free factors in a system with conjugates whenever convenient.

Let $\phi : F \to F$ be an endomorphism. A free factor system $\mathcal{A} = \{A_1, \ldots, A_l\}$ is **ϕ-invariant** if $\mathcal{A}$ supports $\phi(\mathcal{A}) = \{\phi(A_1), \ldots, \phi(A_l)\}$, i.e, there exists a set of elements $\{x_1, \ldots, x_l\} \subset F$ and a function $\sigma : \{1, \ldots, l\} \to \{1, \ldots, l\}$ such that $\phi(A_i) \leq x_i A_{\sigma(i)} x_i^{-1}$ for all i. A ϕ-invariant free factor system $\mathcal{A}$ is **ϕ-fixed** if its free factors are permuted up to conjugacy, i.e., σ is a permutation and $\phi(A_i) = x_i A_{\sigma(i)} x_i^{-1}$ for all i. When ϕ is an automorphism, all ϕ-invariant free factor systems are ϕ-fixed. When ϕ is injective, then, for any $k \geq 1$ and free factor system $\mathcal{A}$ of F, $\phi^k(\mathcal{A})$ is a free factor system of $\phi^k(F)$. A subgroup $A \leq F$ is **eventually ϕ-periodic** if $\phi^m(A) = x\phi^n(A)x^{-1}$ for some $m > n \geq 1$ and $x \in F$, and it is **ϕ-periodic** if $\phi^n(A) = xAx^{-1}$ for some $n \geq 1$ and $x \in F$.

The endomorphism ϕ is **reducible** if it has a nontrivial proper invariant free factor system and **irreducible** otherwise. One immediate consequence of Stallings folds [14] is the injectivity of irreducible endomorphisms (Lemma 2.1 below). So we will usually drop the injective hypothesis when specializing results to the irreducible cases.

We preface the lemma with a summary of Stallings' *folding theorem*. Let R be a rose whose edges are labelled by a basis $\{a_1, \ldots, a_n\}$ of F and let $f : R \to R$ be the map where $f(a_i)$ is the immersed edge-path in R labelled by $\phi(a_i)$. Stallings showed that f factors as $\iota \circ f_l \circ \cdots \circ f_1$ where each f_i is a fold and ι is an immersion [14]. We will use this again to prove bounded cancellation (Lemma 2.5 below).

Lemma 2.1. *If $\phi : F \to F$ is irreducible, then it is injective.*

Proof. If ϕ is not injective, then at least one of the folds in Stallings' factorization of ϕ *collapses* a subgraph of the domain. In particular, the kernel of ϕ contains a proper free factor $A \leq F$; therefore, $\phi(A) = \{1\} \leq A$ and ϕ is reducible. □

Two group homomorphisms $h_1, h_2 : A \to B$ are said to be **equivalent**, denoted $[h_1] = [h_2]$, if there is an inner automorphism $i_b : B \to B$ such that $h_2 = i_b \circ h_1$. **Outer endomorphisms** of F will refer to equivalence classes on the set of endomorphisms of F.

For the topological view, Γ will be a connected finite graph with the same rank as F, i.e., connected, finite 1-dimensional CW-complex such that $\pi_1(\Gamma) \cong F$. A graph map $f : \Gamma \to \Gamma'$ will be a continuous map of graphs that sends vertices to vertices and any edge to a vertex or immersed path. Since a graph map $f : \Gamma \to \Gamma'$ is allowed to send edges to vertices, an edge e of Γ **pretrivial** if $f(e)$ is a vertex of Γ'. Let K be the maximum of the (combinatorial) length of the edge-path $f(e)$ as e varies over all the edges of Γ. Then f is K-Lipschitz, a fact that will be used throughout the dissertation. Generally, $K(f)$ will denote some (convenient) Lipschitz constant for a graph map f.

A **direction at a vertex** $v \in \Gamma$ is a half-edge attached to the vertex. The set of directions at vertex v is denoted by $T_v\Gamma$. If the graph map $f : \Gamma \to \Gamma$ has no pretrivial edges, then the restriction to initial segments induces the derivative map at $\boldsymbol{v}$, $df_v : T_v\Gamma \to T_{f(v)}\Gamma$. The graph map f is an **immersion** if it is locally injective, i.e., it has no pretrivial edges and the derivative maps df_v are injective for all vertices v. An immersion f is **expanding** if it has no invariant subforest (subgraph with contractible components) and the length of $f^n(e)$ is unbounded as $n \to \infty$ for every edge e of Γ.

Let $\mathcal{A} = \{A_1, \ldots, A_l\}$ be a nontrivial free factor system of F, then an **$\mathcal{A}$-marked graph** (Γ_*, α_*) is a collection of graphs $\Gamma_* = \{\Gamma_1, \ldots, \Gamma_l\}$ where the finite graphs Γ_i have vertices that are at least trivalent and are indexed by *markings*, i.e., isomorphisms, $\alpha_i : A_i \to \pi_1(\Gamma_i)$. A **marked graph** (Γ, α) is an $\{F\}$-marked graph. Suppose $\mathcal{A}$ is ϕ-invariant for some injective endomorphism $\phi : F \to F$; so there is a function $\sigma : \{1, \ldots, l\} \to \{1, \ldots, l\}$ and inner automorphisms $i_{x_i} : F \to F$ such that the restrictions $\phi_i = (i_{x_i} \circ \phi)|_{A_i}$ are homomorphisms $\phi_i : A_i \to A_{\sigma(i)}$. The collection $\{\phi_i\}$ is a **restriction of ϕ to $\mathcal{A}$**, denoted by $\phi|_{\mathcal{A}}$. A **topological representative for $\phi|_{\mathcal{A}}$** is a graph map $f_* : \Gamma_* \to \Gamma_*$ of an $\mathcal{A}$-marked graph (Γ_*, α_*) whose restrictions to graphs Γ_i are graph maps $f_i : \Gamma_i \to \Gamma_{\sigma(i)}$ with no pretrivial edges such that $[\pi_1(f_i) \circ \alpha_i] = [\alpha_{\sigma(i)} \circ \phi_i]$ for all $i \in \{1, \ldots, l\}$.

Choosing a marked graph (R, α) where R is a rose is equivalent to choosing a basis $\mathcal{B}$ for F determined by the pre-images of the (oriented) petals over the marking α. Fix a marked graph (Γ, α); for any conjugacy class of a nontrivial element $g \in F$, denoted by $[g]$, we define its **length with respect to α**, $\|g\|_\alpha$, to be the (combinatorial) length of the immersed loop in Γ representing $[g]$. When (R, α) is a marked rose corresponding to a basis $\mathcal{B}$, we shall denote the respective length of $[g]$ by $\|g\|_{\mathcal{B}}$.

For any nontrivial subgroup $H \leq F$, the **Stallings (subgroup) graph** for H with respect to (Γ, α) is the smallest graph $S(H)$ with a marking $\beta : H \to \pi_1(S(H))$ and an immersion $\iota : S(H) \to \Gamma$ such that $[\pi_1(\iota) \circ \beta] = [\alpha|_H]$. Alternatively, $S(H)$ is the **core** of the cover $\hat{\Gamma}_H$ of Γ corresponding to $\alpha(H)$, i.e., the smallest deformation retract of $\hat{\Gamma}_H$, and ι is the restriction to $S(H)$ of the covering map $\hat{\Gamma}_H \to \Gamma$. If H and H' are in the same conjugacy class, $[H]$, then there is a homeomorphism $h : S(H) \to S(H')$ such that $\iota = \iota' \circ h$. The converse holds as well. So the Stallings graph is well-defined for the conjugacy class $[H]$ and vice-versa. Suppose $\phi : F \to F$ is injective, $\mathcal{A}$ is a nontrivial free factor system of

F, and $k \geq 1$. We define $S(\phi^k(\mathcal{A}))$ to be the disjoint union of the Stallings graphs of the free factors of $\phi^k(F)$ in $\phi^k(\mathcal{A})$. By definition, $S(\phi^k(\mathcal{A}))$ is a $\phi^k(\mathcal{A})$-marked graph with the markings $\beta_* = \{\beta_i : A_i \to \pi_1(S_i) \mid S_i = S(\phi^k(A_i))$ is a component of $S(\phi^k(\mathcal{A}))\,\}$.

Unlike the *base graph* Γ, we allow a Stallings graph $S = S(H)$ or $S(\phi^k(\mathcal{A}))$ to have bivalent vertices. More precisely, we assume S is subdivided so that $\iota : S \to \Gamma$ is simplicial/isometric, i.e., ι maps edges to edges. With this subdivision, we get a combinatorial metric on (S, β) that is compatible with (Γ, α), i.e., for any nontrivial element g in H or $\phi^k(\mathcal{A})$, $\|g\|_\alpha = \|g\|_\beta$. For graphs with bivalent vertices, **branch points** are vertices that are not bivalent and **natural edges** are maximal edge-paths whose interior vertices are bivalent.

Lemma 2.2. *Let $\phi : F \to F$ be injective and H be a finitely generated nontrivial subgroup of F. If H is not eventually ϕ-periodic, then the length of the longest natural edge in $S(\phi^k(H))$ is unbounded as $k \to \infty$.*

Proof. Suppose the length of the longest natural edge in $S(\phi^k(H))$ with respect to some marked graph (Γ, α) was uniformly bounded as $k \to \infty$. We want to show that H is eventually ϕ-periodic. The number of natural edges in $S(\phi^k(H))$ is bounded by $3 \cdot \mathrm{rank}(H) - 3$. So our assumption implies there is a bound on the volume of (number of edges in) the graphs $S(\phi^k(H))$ as $k \to \infty$. So the sequence $S(\phi^k(H))$ is eventually periodic, i.e., there are integers $m > n \geq 1$ and an isometry $h : S(\phi^m(H)) \to S(\phi^n(H))$ such that $\iota_m = \iota_n \circ h$. Since a Stallings graph determines the conjugacy class of its defining subgroup, we have $[\phi^m(H)] = [\phi^n(H)]$, i.e., H is eventually ϕ-periodic. □

Conversely, the next lemma handles the case when an invariant free factor system consists entirely of eventually periodic free factors:

Lemma 2.3. *Let $\phi : F \to F$ be injective and $\mathcal{A}$ be a nontrivial ϕ-invariant free factor system. If all free factors in $\mathcal{A}$ are eventually ϕ-periodic, then some nonempty subset $\mathcal{B} \subset \mathcal{A}$ is a ϕ-fixed free factor system and $\phi^k(\mathcal{A})$ is supported by $\mathcal{B}$ for some $k \geq 0$.*

Proof. Let $\sigma : \{1, \ldots, l\} \to \{1, \ldots, l\}$ be the function used to define the ϕ-invariance of $\mathcal{A} = \{A_1, \ldots, A_l\}$. Then there is a nonempty subset $J \subset \{1, \ldots, l\}$ on which σ acts as a bijection and $\sigma^l(\{1, \ldots, l\}) = J$. Let $\mathcal{B} \preceq \mathcal{A}$ by the nontrivial ϕ-invariant free factor system corresponding to J. Then $\phi^l(\mathcal{A})$ is supported by $\mathcal{B}$ since the image of $\sigma^{\,l}$ is J. It remains to show that $\mathcal{B}$ is ϕ-fixed. Set j to be the order of $\sigma|_J$, fix $B \in \mathcal{B}$, and let $i_x : F \to F$ be the inner automorphism such that $i_x \circ \phi^j(B) \leq B$. Define $\psi = i_x \circ \phi^j$. As B is eventually ϕ-periodic and hence eventually ϕ^j-periodic, there are integers $m > n \geq 1$ such that $[\psi^m(B)] = [\phi^{jm}(B)] = [\phi^{jn}(B)] = [\psi^n(B)]$. Therefore, there is an element $y \in F$ such that $y\psi^n(B)y^{-1} = \psi^m(B) \leq \psi^n(B)$. But no finitely generated subgroup of F is conjugate to a proper subgroup of itself (Lemma 2.4 below). So $\psi^m(B) = \psi^n(B)$ and, by injectivity of ϕ, $\psi(B) = B$. Since $B \in \mathcal{B}$ was arbitrary, ϕ^j fixes the free factors of $\mathcal{B}$ up to conjugation; as $\mathcal{B}$ is ϕ-invariant, it must be ϕ-fixed. □

The following fact will be used again in the proof of Lemma 3.5.

Lemma 2.4. *No finitely generated subgroup of F is conjugate to a proper subgroup of itself.*

Proof. By Marshall Hall's theorem, free groups are *subgroup separable/locally extended residual finiteness (LERF)*, i.e., for any finitely generated subgroup $H \leq F$ and element $g \in F \setminus H$, there is a finite group G and a surjective homomorphism $\varphi : F \to G$ such that $\varphi(g) \notin \varphi(H)$. We sketch a proof of this due to Stallings in Appendix A.

For a contradiction, suppose there is an element $y \in F$ such that $yHy^{-1} \leq H$ and $g \in H \setminus yHy^{-1}$. By subgroup separability, there is a finite group G and homomorphism $\varphi : F \to G$ such that $\varphi(g) \notin \varphi(yHy^{-1})$. But $g \in H$ implies $\varphi(g) \in \varphi(H)$ and, by finiteness of G, $yHy^{-1} \leq H$ implies $\varphi(yHy^{-1}) = \varphi(H)$ — a contradiction. □

The next lemma, also known as the **Bounded Cancellation Lemma**, will be used extensively in this dissertation. At the risk of overloading notation, for an edge-path p in a graph Γ, $[p]$ denotes the immersed edge-path that is homotopic to p rel. endpoints; for a loop ρ in Γ, $[\rho]$ will be the immersed edge-path that is freely homotopic to ρ.

Lemma 2.5 (Bounded Cancellation)**.** *Let $g : \Gamma \to \Gamma'$ be a π_1-injective graph map. Then there is a constant $C(g)$ such that, for any natural-edge-path decomposition $p_1 \cdot p_2$ of an immersed path in the universal cover $\tilde{\Gamma}$, the edge-path $[\tilde{g}(p_1)] \cdot [\tilde{g}(p_2)]$ is contained in the $C(g)$-neighborhood of $[\tilde{g}(p_1) \cdot \tilde{g}(p_2)]$.*

The following proof is due to Bestvina-Feighn-Handel [2, Lemma 3.1].

Proof. Any graph map $g : \Gamma \to \Gamma'$ factors as a pretrevial edge collapse and edge subdivision g_0, a composition of $r \geq 0$ Stallings folds $g_1 \circ \cdots \circ g_r$, and an isometric immersion g_{r+1}. The collapse, subdivision, and immersion obviously have cancellation constants 0 while each fold has cancellation constant 1 by π_1-injectivity. Thus we may choose $C(g) = r$. □

If $f : \Gamma \to \Gamma$ is a topological representative for an injective endomorphism $\phi : F \to F$, $\mathcal{A}$ is a nontrivial ϕ-invariant free factor system of F, and $\hat{\Gamma}_k$ is the disjoint union of covers of Γ corresponding to $\phi^k(\mathcal{A})$ for some $k \geq 1$, then f lifts to a map $\hat{f}_k : \hat{\Gamma}_k \to \hat{\Gamma}_k$ and the deformation retraction $\hat{\Gamma}_k \to S(\phi^k(\mathcal{A}))$ induces a map $\bar{f}_k : S(\phi^k(\mathcal{A})) \to S(\phi^k(\mathcal{A}))$ with $K(\bar{f}_k) = K(f)$ and $C(\bar{f}_k) = C(f)$. We shall call $\bar{f}_k$ the **(k-th) homotopy lift of f**.

The maps $\bar{f}_k$ need not be topological representatives for $\phi|_{\mathcal{A}}$ with respect to natural edges: they might map branch points to bivalent vertices (interior of natural edges) or fail to be locally injective on natural edges. We hope to replace $\bar{f}_k$ with homotopic maps that are as close to topological representatives as possible for our purposes while still maintaining uniform control on the Lipschitz and cancellation constants.

As the bounded cancellation lemma only considers natural-edge-paths p_1, p_2, homotopies that are supported in the interior of natural edges will not affect the cancellation constant. Using a (tightening) homotopy supported in the interior of natural edges of $S(\phi^k(\mathcal{A}))$, we may replace $\bar{f}_k$ with a homotopic map that maps any natural edge to either a vertex or an immersed path and has the same Lipschitz and cancellation constants. One would usually collapse the pretrivial edges of $S(\phi^k(\mathcal{A}))$ to get a map that is locally injective on natural edges but we will not since we want to preserve compatibility:

$\|\cdot\|_\beta$ is the restriction of $\|\cdot\|_\alpha$ to $\phi^k(\mathcal{A})$. The next lemma uses bounded cancellation to measure how close the homotopy lift $\bar{f}_k$ is to mapping branch points to branch points.

Lemma 2.6. *Let $f : \Gamma \to \Gamma$ be a topological representative for an injective endomorphism $\phi : F \to F$. For any $k \geq 1$, if $\bar{f}_k : S(\phi^k(\mathcal{A})) \to S(\phi^k(\mathcal{A}))$ is a homotopy lift of f with $C(\bar{f}_k) = C(f)$, then $\bar{f}_k$ maps branch points to the $C(f)$-neighborhood of branch points.*

Proof. Set $C = C(f) = C(\bar{f}_k)$. If $S(\phi^k(\mathcal{A}))$ is the C-neighborhood of its branch points, then there is nothing to prove. Suppose $\nu \in S(\phi^k(\mathcal{A}))$ is a bivalent vertex whose distance to the nearest branch point is $> C$. We need to show that ν is not the $\bar{f}_k$-image of any branch point. Set ϵ_1 and ϵ_2 to be the distinct oriented directions originating from ν and $\bar{\epsilon}_1, \bar{\epsilon}_2$ are the same directions with opposite orientation.

Suppose $v \in S(\phi^k(\mathcal{A}))$ is a branch point and $\bar{f}_k(v) = \nu$. As v is a branch point, there are at least three distinct oriented directions originating from v: e_1, e_2, and e_3. Let p_{12} be an immersed path that starts and ends with e_1 and $\bar{e}_2$ respectively and define p_{23} similarly. Set $p_{13} = [p_{12} \cdot p_{23}]$ and $p'_{13} = [p_{12} \cdot \bar{p}_{23}]$, where $\bar{p}_{23}$ is the reversal of the path p_{23} (See Figure 1). Although the paths are loops, we still treat them as paths, i.e., reduction is done rel. the endpoints (v or ν). Without loss of generality, assume $[\bar{f}_k(p_{12})]$ starts with ϵ_1.

If $[\bar{f}_k(p_{12})]$ ends with $\bar{\epsilon}_1$, then $[\bar{f}_k(p_{12})] = \mu_1 \cdot \rho \cdot \bar{\mu}_1$, where μ_1 is an extension of ϵ_1 to an embedded path from ν to a branch point and ρ is an immersed loop. By hypothesis, μ_1 is longer than C. Since p_{12} starts and ends with e_1 and $\bar{e}_2$ respectively, the concatenation $p_{12} \cdot p_{12}$ is an immersed path such that $[\bar{f}_k(p_{12})] \cdot [\bar{f}_k(p_{12})]$ has $\bar{\mu}_1 \cdot \mu_1$ as a subpath, violating bounded cancellation. So we may assume $[\bar{f}_k(p_{12})]$ starts and ends with ϵ_1 and $\bar{\epsilon}_2$.

If $[\bar{f}_k(p_{23})]$ starts and ends with ϵ_2 and $\bar{\epsilon}_1$, then $[\bar{f}_k(p_{13})] = [\bar{f}_k(p_{12}) \cdot \bar{f}_k(p_{23})]$ starts and ends with ϵ_1 and $\bar{\epsilon}_1$ respectively, which violates bounded cancellation for the same reason given in the previous paragraph. Similarly, if $[\bar{f}_k(p_{23})]$ starts and ends with ϵ_1 and $\bar{\epsilon}_2$, we rule out this possibility by considering $[\bar{f}_k(p'_{13})]$. We have ruled out all cases, and therefore, no branch point v of $S(\phi^k(\mathcal{A}))$ is mapped to ν. □

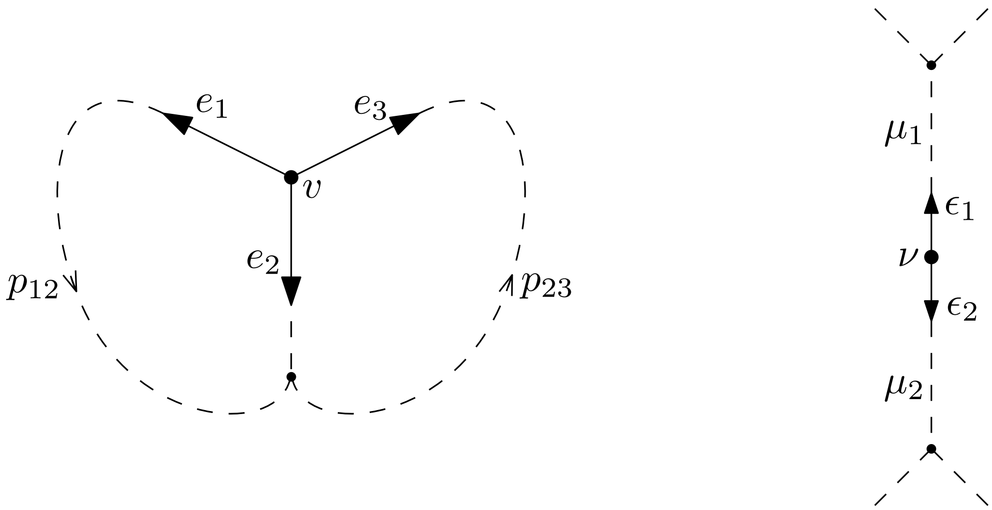

Figure 1: Schematic for paths p_{12}, p_{23}, p_{13}, and p'_{13}. The path p_{13} starts with e_1 follows the dashed path and ends with $\bar{e}_3$. The path p'_{13} is the "figure 8" path traced by p_{12} then $\bar{p}_{23}$.

As $\bar{f}_k$ maps branch points to the $C(f)$-neighborhood of branch points, we replace it with a homotopic map that maps branch points to branch points. The homotopy can potentially increase both the Lipschitz and cancellation constants. However, since the homotopy is moving the image of a branch point a distance at most $C(f)$, we can use $K(\bar{f}_k) = K(f) + C(f)$ and $C(\bar{f}_k) = 2C(f)$. We will now summarize the properties of homotopy lifts that we will be using in the next sections.

Proposition 2.7. *Suppose $f : \Gamma \to \Gamma$ is a topological representative for an injective endomorphism ϕ and $\mathcal{A}$ is a nontrivial ϕ-invariant free factor system of F. For any $k \geq 1$, there is a map $\bar{f}_k : S(\phi^k(\mathcal{A})) \to S(\phi^k(\mathcal{A}))$ on the Stallings graph that has these properties:*

1. *it is a homotopy lift of f, i.e., $\iota \circ \bar{f}_k \simeq f \circ \iota$;*
2. *it maps branch points to branch points and any natural edge to a branch point or an immersed path;*
3. *$K(\bar{f}_k) = K(f) + C(f)$ and $C(\bar{f}_k) = 2C(f)$.*

The point being that the Lipschitz and cancellation constants are independent of k.

3 Elliptic free factor systems

In this section, we will construct a canonical invariant free factor system for any given injective endomorphism of F. This free factor system, called the *elliptic free factor system*, is crucial for the construction of *(expanding) relative immersions* later in the dissertation.

Suppose $\phi : F \to F$ is an injective endomorphism with an invariant free factor system $\mathcal{A}$. We shall say $\mathcal{A}$ **supports a nontrivial conjugacy class** $[g]$ in F if there is a nontrivial element $g' \in A, A \in \mathcal{A}$ such that $[g'] = [g]$. Similarly, $\mathcal{A}$ **supports an infinite ϕ-tail of a nontrivial conjugacy class** $[g]$ in F if for every $n \geq 1$, there is a nontrivial element $g_n \in A, A \in \mathcal{A}$ such that $[\phi^n(g_n)] = [g]$. We now state the main technical result of this section.

Theorem 3.1. *If $\phi : F \to F$ is an injective endomorphism, then ϕ has a nontrivial fixed free factor system if and only if some nontrivial conjugacy class in F has an infinite ϕ-tail.*

Remark. Patrick Reynolds defined *expansive* endomorphisms [12, Definition 3.8] and ϕ is expansive in this sense if and only if some conjugacy class has an infinite ϕ-tail. Under this equivalence, Reynolds' Remark 3.12 in [12] is a weaker form of this theorem.

Proof. The forward direction is obvious: if ϕ has a nontrivial fixed free factor system, then it has a periodic free factor $A \leq F$ and this implies all nontrivial conjugacy classes $[g]$ in A have an infinite ϕ-tail. The main content of the theorem is in the reverse direction.

Suppose the ϕ-invariant free factor system $\{F\}$ supports an infinite ϕ-tail of some nontrivial conjugacy class $[g]$ in F. If ϕ is an automorphism, then $\{F\}$ is a nontrivial ϕ-fixed free factor system and we are done. So we assume ϕ is nonsurjective and proceed by descending down the poset of free factor systems. The following claim is the key idea:

Claim (Descent). *If $\mathcal{D}$ is a ϕ-invariant free factor system that supports an infinite ϕ-tail of $[g]$ but contains a free factor that is not eventually ϕ-periodic, then some ϕ-invariant free factor system $\mathcal{D}' \prec \mathcal{D}$ supports an infinite ϕ-tail of $[g]$.*

Since the poset of free factor systems has no infinite chains, the descent starts with the free factor system $\{F\}$ and nonsurjectivity of ϕ and then finds a ϕ-invariant free factor sytem $\mathcal{A}$ that supports an infinite ϕ-tail of $[g]$ and whose free factors are eventually ϕ-periodic; such a free factor system contains a ϕ-fixed nontrivial free factor system by Lemma 2.3 and we are done. It remains to show the proof of the descent claim.

Note that the assumption that $\mathcal{D}$ contains a free factor that is not eventually ϕ-periodic is more general than we really need for our conclusion. In fact, we could have worked with a more natural assumption, $\mathcal{D}$ has no periodic free factors, and gotten the same conclusion. However, we give the general argument instead since we plan on invoking a variation of the claim again in the next proposition.

Proof of descent. Let (Γ, α) be a marked graph, $f : \Gamma \to \Gamma$ be a topological representative for ϕ, and set $K = K(f) + C(f)$ and $C = 2C(f)$. Suppose $\mathcal{D}$ is a ϕ-invariant free factor system that supports an infinite ϕ-tail $(g_n)_{n \geq 1}$ of $[g]$ and contains a free factor that is not eventually ϕ-periodic. Then, for all $k \geq 1$, there is an immersed loop $\rho_k(g_k)$ in $\Delta_k = S(\phi^k(\mathcal{D}))$ of length $\|g\|_\alpha$ corresponding to $[\phi^k(g_k)] = [g]$, where $\Delta_k = S(\phi^k(\mathcal{D}))$ is the Stallings graph with respect to (Γ, α). Set $L = \max\{\|g\|_\alpha, C\}$.

For all $k \geq 1$, let $\bar{f}_k : \Delta_k \to \Delta_k$ be the homotopy lifts of f given by Proposition 2.7. In particular, these lifts map branch points to branch points, map any natural edge to either a branch point or an immersed path, and have Lipschitz and cancellation constants K and C respectively. As $(g_n)_{n \geq 1}$ is an infinite ϕ-tail of $[g]$ supported in $\mathcal{D}$, we get, for any fixed $k \geq 1$, an infinite sequence of immersed loops $(\rho_k(g_n))_{n \geq k}$ in Δ_k such that $\bar{f}_k^{n-k}(\rho_k(g_n))$ have length L for all $n > k$ but they need not be homotopic to $\rho_k(g_k)$ in Δ_k as $\phi^n(g_n)$ need not be conjugate to $\phi^k(g_k)$ in $\phi^k(F)$.

Form a directed graph $\mathbb{G}_k$ whose vertices are the natural edges of Δ_k and there is an edge $E_i \to E_j$ if $\bar{f}_k$ maps E_i over E_j. Note that the number of natural edges of Δ_k is at most $N = 3 \cdot \operatorname{rank}(F) - 3$ and so $\mathbb{G}_k$ has at most N vertices.

Since $\mathcal{D}$ contains a free factor that is not eventually ϕ-periodic, the length of natural edges in Δ_k is unbounded as $k \to \infty$ by Lemma 2.2. Fix $k \gg 0$ such that the longest natural edge in Δ_k is longer than $L \cdot K^{N-1}$. Let $\mathcal{L}_0$ be the natural edges of Δ_k longer than $L \cdot K^{N-1}$ and $\mathcal{L}$ be the union of $\mathcal{L}_0$ and all natural edges that have a directed path to $\mathcal{L}_0$ in $\mathbb{G}_k$. Since $\bar{f}_k$ is K-Lipschitz and the shortest directed path in $\mathbb{G}_k$ from a natural edge in $\mathcal{L}$ to $\mathcal{L}_0$ has at most N natural edges on it, every natural edge in $\mathcal{L}$ is longer than L. The natural edges in $\mathcal{L}$ will be referred to as the *long* natural edges and the remaining natural edges as *short*.

Set $\Delta' \subset \Delta_k$ to be the union of short natural edges, which is a proper subgraph since long natural edges exist. The subgraph is automatically $\bar{f}_k$-invariant by the construction of $\mathcal{L}$. Since $\rho_k(g_k)$ is an immersed loop in Δ_k with length $\|g\|_\alpha \leq L$, its natural edges are short and Δ' is a nonempty, noncontractible proper subgraph of Δ_k. Therefore, Δ' determines a nontrivial ϕ-invariant proper free factor system $\mathcal{D}' \prec \mathcal{D}$. Technically, it determines a $\phi|_{\phi^k(F)}$-invariant free factor system of $\phi^k(F)$ but, as ϕ is injective, this corresponds to a ϕ-invariant free factor system of F. It remains to show that $\mathcal{D}'$ supports an infinite ϕ-tail of $[g]$.

Let $\mathbb{L} \subset \mathbb{G}_k$ be the full subgraph of generated by the long natural edges $\mathcal{L}$. If there are no directed cycles in $\mathbb{L}$, then $\bar{f}_k^N(\Delta_k) \subset \Delta'$; in this case, the sequence of nontrivial loops $(\bar{f}_k^N(\rho_k(g_n)))_{n \geq k+N}$ in Δ' naturally determines an infinite ϕ-tail of $[g]$ supported in $\mathcal{D}'$ and we are done. Now suppose there are directed cycles in $\mathbb{L}$ and ρ be an immersed loop in Δ_k that contains a long natural edge in such a cycle. Then, by bounded cancellation and the fact long natural edges are longer than $L \geq C$, $[\bar{f}_k^m(\rho)]$ contains a long natural edge in the same directed cycle in $\mathbb{L}$ for all $m \geq 1$. Consequently, none of the immersed loops $\rho_k(g_n)$ in Δ_k contain a long natural edge that is in a directed cycle of $\mathbb{L}$. Therefore, as far as the sequence of loops $(\rho_k(g_n))_{n \geq k}$ is concerned, we may assume there are no directed cycles in $\mathbb{L}$ and, as before, the sequence $(\bar{f}_k^N(\rho_k(g_n)))_{n \geq k+N}$ determines an infinite ϕ-tail of $[g]$ supported in $\mathcal{D}'$. □

The following dichotomy is (equivalent to) a result in Reynolds' thesis.

Corollary 3.2 ([12, Proposition 3.11])**.** *If $\phi : F \to F$ is irreducible, then either ϕ is an automorphism or no nontrivial conjugacy class in F has an infinite ϕ-tail.*

Proof. Suppose ϕ is irreducible and there is a nontrivial conjugacy class with an infinite ϕ-tail. By Theorem 3.1, there is a nontrivial ϕ-fixed free factor system $\mathcal{A}$. Since ϕ is irreducible, $\mathcal{A} = \{F\}$ and ϕ is an automorphism. □

The fixed free factor system given by Theorem 3.1 may depend on the chosen conjugacy class in F that has an infinite tail. The following proposition constructs a canonical fixed free factor system for injective endomorphisms that supports all conjugacy classes with an infinite tail as well as all finitely generated periodic subgroups. The proof will use both ascent and descent in the poset of free factor systems!

Proposition 3.3. *If $\phi : F \to F$ is an injective endomorphism, then there is a unique maximal ϕ-fixed free factor system $\mathcal{A}$. Precisely, $\mathcal{A}$ supports every a nontrivial conjugacy class in F with an infinite ϕ-tail and every finitely generated ϕ-periodic subgroup of F.*

Proof. If ϕ is an automorphism, then let $\mathcal{A} = \{F\}$ and we are automatically done. If ϕ has no fixed nontrivial free factor system, then the trivial free factor system is automatically the unique maximal ϕ-fixed free factor system. By Theorem 3.1, there are no conjugacy classes with an infinite ϕ-tail and, consequently, there are no ϕ-periodic subgroups. So the trivial free factor system vacuously supports infinite ϕ-tails and ϕ-periodic subgraphs. So we may assume ϕ is nonsurjective and has a fixed nontrivial free factor system $\mathcal{D}_0$. We proceed by ascending up the poset of free factor systems:

Claim (Ascent)**.** *Let $A \leq F$ be a finitely generated ϕ-periodic subgroup and $[g]$ be a nontrivial conjugacy class in F with an infinite ϕ-tail. If $\mathcal{D}$ is a ϕ-fixed nontrivial free factor system that does not support both $\{A\}$ and $[g]$, then some ϕ-fixed proper free factor system $\mathcal{D}' \succ \mathcal{D}$ supports both $\{A\}$ and $[g]$.*

Once again, as there are no infinite chains in the poset of free factor systems, the ascent starts with the ϕ-fixed nontrivial proper free factor $\mathcal{D}_0$ and will stop with the necessarily unique maximal ϕ-fixed proper free factor system $\mathcal{A}$ that supports all finitely generated ϕ-periodic subgroups of F and all conjugacy classes in F with an infinite ϕ-tail. It remains to prove the ascent claim.

Proof of ascent. Let $A \leq F$ be a finitely generated ϕ-periodic subgroup, $[g]$ be a nontrivial conjugacy class in F with an infinite ϕ-tail, and $\mathcal{D}$ be a ϕ-fixed free factor system that does not support both $\{A\}$ and $[g]$. We now describe the descent portion:

Claim (Descent)**.** *If $\mathcal{D}''$ is ϕ-invariant free factor system that supports $\{A\}$, an infinite ϕ-tail of $[g]$, and $\mathcal{D}$ but contains a free factor that is not eventually ϕ-periodic, then some ϕ-invariant free factor system $\mathcal{D}''' \prec \mathcal{D}''$ supports $\mathcal{D}$, $\{A\}$, and an infinite ϕ-tail of $[g]$.*

Starting with $\{F\}$ and nonsurjectivity of ϕ, the descent will find a ϕ-invariant proper free factor system $\mathcal{D}^*$ that supports $\mathcal{D}$, $\{A\}$, and an infinite ϕ-tail of $[g]$ and whose free factors are eventually ϕ-periodic. By Lemma 2.3, $\mathcal{D}^*$ contains a ϕ-fixed free factor subsystem $\mathcal{D}' \subset \mathcal{D}^*$ such that $\phi^k(\mathcal{D}^*)$ is supported by $\mathcal{D}'$ for some $k \geq 0$. As A is ϕ-periodic, it is supported by $\mathcal{D}'$. Similarly, $\mathcal{D}'$ supports $[g]$ since $\mathcal{D}^*$ supports an infinite ϕ-tail of $[g]$. So $\mathcal{D}'$ is a ϕ-fixed free factor system that supports $\mathcal{D}$, $\{A\}$, and $[g]$ as needed for ascent. It remains to prove the descent claim. □

Proof of descent. Let (Γ, α) be a marked graph, $f : \Gamma \to \Gamma$ be a topological representative for ϕ, and set $K = K(f) + C(f)$ and $C = 2C(f)$. Suppose $\mathcal{D}''$ is a ϕ-invariant free factor system that supports $\mathcal{D}$, $\{A\}$, and an infinite ϕ-tail of $[g]$. Let $S(\phi^k(\mathcal{D}))$ and $S(\phi^k(A))$ be the Stallings graphs of $\phi^k(\mathcal{D})$ and A with respect to (Γ, α). Since A and the free factors of $\mathcal{D}$ are finitely generated and ϕ-periodic, the length of the longest immersed loop in $S(\phi^k(\mathcal{D}))$ or $S(\phi^k(A))$ that covers any edge at most twice is uniformly bounded by some L_0 for all $k \geq 1$. Set $L = \max\{L_0, \|g\|_\alpha, C\}$. The proof now mimics that of descent in Theorem 3.1 so we only give a sketch.

For all $k \geq 1$, let $\Delta''_k = S(\phi^k(\mathcal{D}''))$ and, by Proposition 2.7, there is a K-Lipschitz homotopy lift $\bar{f}_k : \Delta''_k \to \Delta''_k$ that maps branch points to branch points and has cancellation constant C. As some free factor in $\mathcal{D}''$ is not eventually ϕ-periodic, we can fix $k \gg 0$ so that the longest natural edge in Δ''_k is longer than $L \cdot K^{N-1}$ by Lemma 2.2. Define the long and short natural edges as before and deduce long natural edges are longer than L. Set $\Delta''' \subset \Delta''_k$ to be the union of the short natural edges, which is necessarily proper and $\bar{f}_k$-invariant. Recall that immersed loops of $S(\phi^k(\mathcal{D}))$ and $S(\phi^k(A))$ that cover any edge at most twice have length bounded by $L_0 \leq L$ and these Stallings graphs have isometric natural immersions into Δ''_k. Hence the images of these natural immersions lie in the subgraph of short natural edges Δ'''. So Δ''' is nonempty and noncontractible. The subgraph Δ''' determines a ϕ-invariant proper free factor system $\mathcal{D}''' \prec \mathcal{D}''$ that supports both $\mathcal{D}$ and $\{A\}$ since both $S(\phi^k(\mathcal{D}))$ and $S(\phi^k(A))$ have immersions into Δ'''. From the proof of descent in Theorem 3.1, $L \geq C$ implies $\mathcal{D}'''$ supports an infinite ϕ-tail of $[g]$. $\square$

Although this proposition produces a canonical ϕ-fixed free factor system for an injective endomorphism, we shall enlarge the system again to get a better ϕ-invariant free factor system that gives us some control of the *relative dynamics* of ϕ. We do this by taking iterated preimages of the free factor system.

Proposition 3.4. *If $\phi : F \to F$ is an injective endomorphism and $\mathcal{A}$ is the maximal ϕ-fixed free factor system, then there is a unique maximal ϕ-invariant free factor system $\mathcal{A}^* \succeq \mathcal{A}$ such that $\phi^k(\mathcal{A}^*)$ is supported by $\mathcal{A}$ for some $k \geq 0$.*

We shall call the free factor system given by this proposition the **(canonical) ϕ-elliptic free factor system**. The *elliptic* nature will be apparent from the proof.

Proof. Let $\phi : F \to F$ be an injective endomorphism and $\mathcal{A}$ be the maximal ϕ-fixed free factor system of F given by Proposition 3.3. If ϕ surjective or has no nontrivial fixed free factor systems, then $\mathcal{A}^* = \mathcal{A}$ is $\{F\}$ or trivial respectively and we are done. Thus, we assume that ϕ is nonsurjective and $\mathcal{A}$ is a nontrivial proper free factor system. Let T be a simplicial tree with a minimal F-action whose edge stabilizers are trivial and nontrivial point stabilizers are conjugates of free factors in $\mathcal{A}$. The quotient $F \backslash T$ is a *graph of groups* decomposition of F with trivial edge groups and $\mathcal{A}$ as the nontrivial vertex groups; this is also known as a *free splitting* of F. We give a quick review of *Bass-Serre theory* in Appendix B. For any $k \geq 1$, we get a $\phi^k(F)$-action on the minimal subtree $T_k \subset T$ of $\phi^k(F)$. Under this action, the quotient $\phi^k(F) \backslash T_k$ is a free splitting of $\phi^k(F)$. By the injectivity of ϕ, we get a free splitting of F whose vertex groups form a free factor system of F denoted by $\phi^{-k} \cdot \mathcal{A} = \mathcal{A}_k$. Since $\mathcal{A}$ is ϕ-invariant, we get that $\mathcal{A} \preceq \mathcal{A}_k \preceq \mathcal{A}_{k+1}$ for all $k \geq 1$ and, consequently, all $\mathcal{A}_k$ are ϕ-invariant. As there are no infinite chains in the poset of free factor systems of F, we get that the chain of ϕ-invariant free factor systems $\mathcal{A}_k$ $(k \geq 1)$ stabilizes and we can set $\mathcal{A}^*$ to be the maximal free factor system in the chain. By construction, $\mathcal{A}^*$ supports any subgroup $A \leq F$ such that $\phi^k(A)$ is supported by $\mathcal{A}$ for some $k \geq 0$ and this implies the uniqueness of $\mathcal{A}^*$. □

The next lemma states that the elliptic free factor system is simply a disjoint union of the maximal fixed free factor system with a free factor system that eventually gets mapped into the maximal fixed free factor system.

Lemma 3.5. *If $\phi : F \to F$ is an injective endomorphism, $\mathcal{A}$ is the maximal ϕ-fixed free factor system, and $\mathcal{A}^*$ is the ϕ-elliptic free factor system, then $\mathcal{A} \subset \mathcal{A}^*$ after replacing the free factors of $\mathcal{A}$ with conjugates if necessary.*

Proof. Suppose $\sigma : \{1, \ldots l\} \to \{1, \ldots, l\}$ is the function associated to the ϕ-invariance of the ϕ-elliptic free factor system $\mathcal{A}^* = \{A_1, \ldots A_l\}$. Then there is a maximal nonempty subset $J \subset \{1, \ldots, l\}$ on which σ is a bijection. Let $\mathcal{A}_J = \{\, A_j \in \mathcal{A}^* \,:\, j \in J \,\}$. Since $\mathcal{A}^*$ supports the maximal ϕ-fixed free factor system $\mathcal{A}$, it follows that $\mathcal{A} \preceq \mathcal{A}_J$. Replace free factors of $\mathcal{A}$ with conjugates if necessary and assume each $A \in \mathcal{A}$ is a subgroup of some $A_j \in \mathcal{A}_J$. We want to show that $\mathcal{A} \subset \mathcal{A}_J$. Choose a free factor $A \in \mathcal{A}$ and let $A_j \in \mathcal{A}_J$ be the free factor such that $A \leq A_j$. Furthermore, fix an inner automorphism $i_x : F \to F$ such that $i_x \circ \phi^s(A_j) \leq A_j$ for some $s \geq 0$. Set $\psi = i_x \circ \phi^s$ and note that ψ and ϕ must have the same maximal fixed and elliptic free factor systems by construction. Consequently, $A_j \in \mathcal{A}^*$ implies $\psi^k(A_j)$ is conjugate to a subgroup of a ψ-periodic free factor $A' \in \mathcal{A}$ for some $k \geq 0$. We must have $A' \leq A_j$ since $A' \leq A_{j'} \in \mathcal{A}_J$, $\{A'\}$ supports $\{\psi^k(A_j)\}$, $\psi(A_j) \leq A_j$, and $\sigma|_J$ is a bijection. So $i_y \circ \psi^k(A_j) \leq A' \leq A_j$ for some inner automorphism $i_y : F \to F$. The ψ-periodicity of A' implies $(i_y \circ \psi^k)^m(A') \leq A'$ is conjugate to A' for some $m \geq 1$. But Lemma 2.4 says no finitely generated subgroup of F is conjugate to a proper subgroup of itself. Therefore, $(i_y \circ \psi^k)^m(A') = A'$ and, by injectivity of ψ, $i_y \circ \psi^k(A_j) = A' = A_j$. In particular, $A = A' = A_j$. As this holds for arbitrary free factors $A \in \mathcal{A}$, we get $\mathcal{A} \subset \mathcal{A}_J$. $\square$

It is obvious that the maximal ϕ-fixed free factor system is proper if and only if ϕ is nonsurjective. The same holds for the elliptic free factor system:

Lemma 3.6. *Let $\phi : F \to F$ be an injective endomorphism. The ϕ-elliptic free factor system is*

1. *nontrivial if and only if there is a nontrivial conjugacy class with an infinite ϕ-tail;*

2. *proper if and only if ϕ is nonsurjective.*

Proof. Let $\mathcal{A}^*$ be the ϕ-elliptic free factor system. Then, by construction and injectivity of ϕ, $\mathcal{A}^*$ is nontrivial if and only if ϕ has a nontrivial fixed free factor system. The first equivalence then follows from Theorem 3.1.

The forward direction in the second equivalence is the fact that the elliptic free factor system of an automorphism is $\{F\}$. For the backward direction. Suppose $\mathcal{A}^* = \{F\}$. By construction of $\mathcal{A}^*$ and injectivity of ϕ, the maximal ϕ-fixed free factor system $\mathcal{A}$ is nontrivial. So $\mathcal{A} = \{F\}$ by Lemma 3.5 and ϕ is an automorphism. $\square$

4 — Interlude —

Let us take a short moment to summarize the previous two sections and how they will relate to the next two sections. We will end the interlude with three useful examples.

In the previous section, we considered pairs of free factor systems $\mathcal{A} \prec \{F\}$ where we only varied $\mathcal{A}$ until we found an invariant free factor system with unique properties with respect to an injective endomorphism $\phi : F \to F$, e.g., the unique maximal ϕ-fixed free factor system. The main tools we used were bounded cancellation and the fact Stallings graphs $S(\phi^k(F))$ have arbitrarily long natural edges (Lemma 2.2).

The whole discussion was done in terms of marked graphs (Γ, α) and Stallings graphs $S(\phi^k(\mathcal{A}))$ that come with immersions into the base graph Γ. On the other hand, we could have equivalently done everything in terms of F-trees T, i.e., simplicial trees T with a free minimal F-action $\alpha : F \to \mathrm{Isom}(T)$, and $\phi^k(\mathcal{A})$-forests that come with equivariant embeddings into the ambient tree T. Taking universal covers gives the translation from the graph-setting to the tree-setting while taking quotients by deck transformations gives the translation back.

In the next section, we will consider more general pairs of invariant free factor systems $\mathcal{A} \prec \mathcal{B}$ where $\mathcal{A}$ is more or less fixed/understood and there is no free factor system strictly between the two in the poset of invariant free factor systems. This discussion will be done in terms of forests T_* with minimal $\mathcal{B}$-actions whose edge-stabilizers are trivial. We shall distinguish $\mathcal{A}$ by having its conjugates in $\mathcal{B}$ be the nontrivial point stabilizers under the $\mathcal{B}$-actions and we call these $(\mathcal{B}, \mathcal{A})$-actions. In lieu of Stallings graphs $S(\phi^k(\mathcal{B}))$, we will be looking at minimal subforests $T_*(\phi^k(\mathcal{B}))$ in an ambient forest T_*. The main objective of the next section is to define *relative representatives* and give a relative version of bounded cancellation and Lemma 2.2. These relativizations will be the main tools for understanding relative dynamics.

We now give examples to illustrate the constructions from the previous section and motivate the relative setting in the next sections. Let $F = F(a, b)$ be the free group on two generators and $\phi, \varphi, \psi : F \to F$ be injective endomorphisms given by

$$\phi : (a, b) \mapsto (ab, ba), \quad \varphi : (a, b) \mapsto (a, bab^{-1}), \quad \text{and} \quad \psi : (a, b) \mapsto (a, abab).$$

The **standard rose** is the marked rose (R, α) such that $\|a\|_\alpha = \|b\|_\alpha = 1$; it will be graphically represented by a rose with two oriented petals labelled by a and b respectively. A Stallings graph S with respect to the standard rose along with its corresponding immersion $\iota : S \to R$ are graphically represented as R-digraphs, i.e., S will be an oriented graph whose oriented edges are labelled by a or b and no two a- or b-edges share an origin as shown in the following figure.

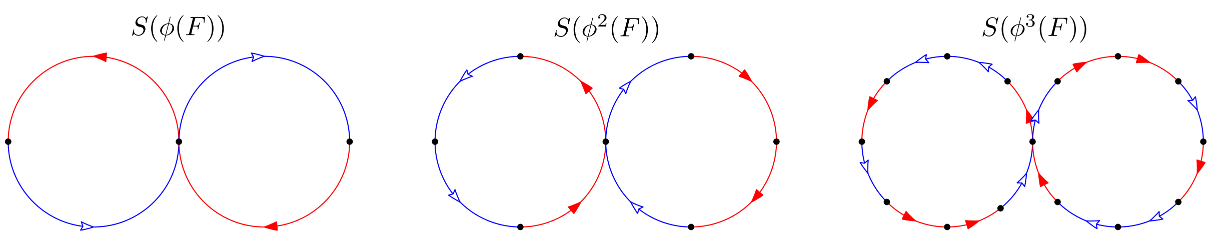

Figure 2: Stallings graphs $S(\phi^k(F))$ with respect to the standard rose for $k = 1, 2, 3$.

The Stallings graphs $S(\phi^k(F))$ are roses for all $k \geq 1$ and each petal doubles in size with each iteration. Since ϕ is injective, we may consider $S(\phi^k(F))$ as marked roses (Γ_k, α_k) for $k \geq 1$, and forgetting the bivalent vertices in Γ_k, we recover the standard rose. So topologically, the marked roses (Γ_k, α_k) are (*equivalent to*) the standard rose and this is equivalent to ϕ having a topological representative defined on the standard rose that is an immersion [9, Lemma 3.2]. Indeed, the obvious map on the standard rose is an immersion.

Remembering the bivalent vertices in Γ_k again, the fact that the petals of Γ_k double in size with each iteration implies ϕ is in fact induced by an expanding immersion, which in turn implies ϕ has no fixed nontrivial free factor system. We shall see later (Corollary 6.6) that the converse holds as well: if an injective endomorphism has no fixed nontrivial free factor system then it is induced by an expanding graph immersion.

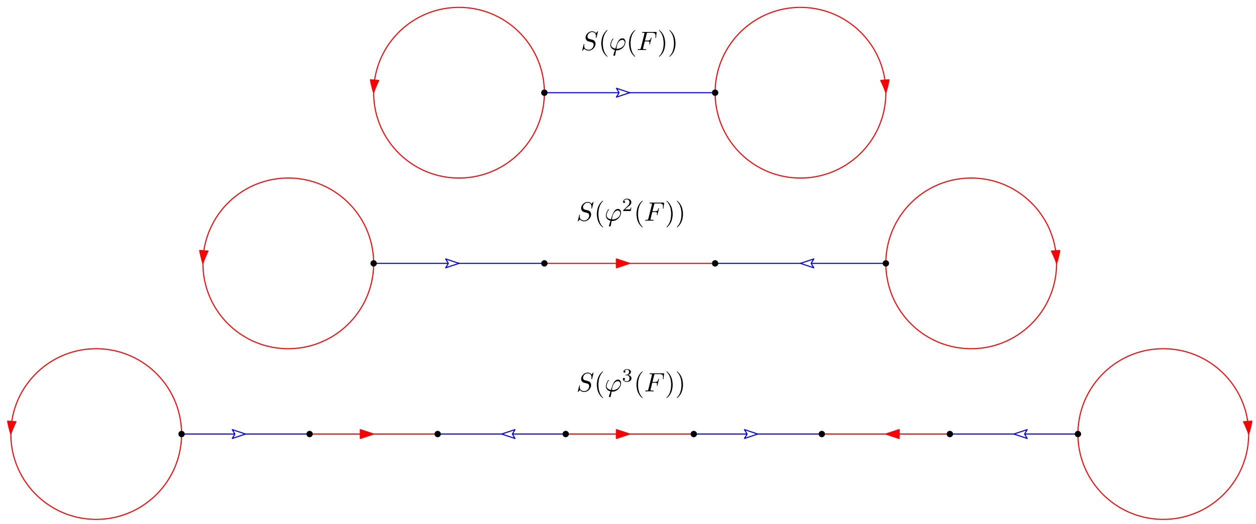

Figure 3: Stallings graphs $S(\varphi^k(F))$ for $k = 1, 2, 3$.

The Stallings graphs $S(\varphi^k(F))$ are barbells for all $k \geq 1$, the middle bars roughly double in length with each iteration, and the embedded loops are labelled by a. Considering the graphs as marked barbells (B_k, β_k) for $k \geq 1$ and forgetting the bivalent vertices in B_k, we get the *standard barbell* (B, β), i.e., the barbell formed by expanding the vertex of the standard rose to a separating edge. So topologically, the marked barbells (B_k, β_k) are the standard barbell and ϕ is induced by an immersion $g : B \to B$ of the standard barbell.

Remembering the bivalent vertices in B_k, the fact that all the barbells B_k have a loop labelled by a implies $[a]$ has an infinite φ-tail. Hence, the immersion g is not expanding, the maximal φ-fixed free factor system is $\mathcal{A}_\varphi = \{\langle a \rangle\}$, and the φ-elliptic free factor system is $\mathcal{A}^*_\varphi = \{\langle a \rangle, \langle b \rangle\}$. However, if we lift g to the universal cover $\tilde{g} : \tilde{B} \to \tilde{B}$ and then collapse all translates of axes of a and b in $\tilde{B}$, then we get an expanding immersion $\bar{g} : T \to T$ on an $(F, \mathcal{A}^*_\varphi)$-tree T where every edge doubles in length.

Finally, the Stallings graphs $S(\psi^k(F))$ are roses for all $k \geq 1$ where one petal roughly doubles in length with each iteration and another is labelled by a. Unlike the previous two examples, the corresponding marked roses are all distinct for $k \geq 1$ due to the folding occuring at each iteration. In particular, ψ cannot be induced by a graph immersion! But the fact the marked roses all have a petal labelled by a implies $[a]$ has an infinite ψ-tail. In

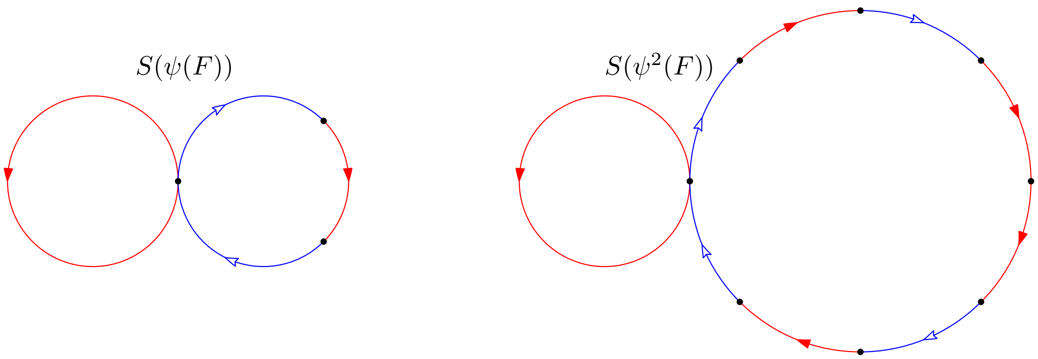

Figure 4: Stallings graphs $S(\psi^k(F))$ for $k = 1, 2$.

this case, the maximal ψ-fixed free factor system and the ψ-elliptic free factor system are $\mathcal{A}_\psi = \mathcal{A}^*_\psi = \{\langle a \rangle\}$. Despite ψ not being induced by an immersion, there is a topological representative $h : R \to R$ defined on the standard rose such that collapsing the translates of the axis of a in the universal cover $\tilde{R}$ induces an expanding immersion $\bar{h} : Y \to Y$ on an $(F, \mathcal{A}^*_\psi)$-tree Y where every edge doubles in length.

The main result of the next two sections is that this construction will always produce expanding immersions: any nonsurjective injective endomorphism of a free group has a topological representative $f : \Gamma \to \Gamma$ such that collapsing the translates of axes in $\tilde{\Gamma}$ of the elliptic free factor system will induce an expanding immersion $\bar{f} : T \to T$ on a tree T.

5 Relative representatives

This section defines relative representatives that will be the basis for inductively studying dynamical properties of free group endomorphisms. For the whole section, we suppose $\phi : F \to F$ is an injective endomorphism and $\phi^{-1} \cdot \mathcal{A} = \mathcal{A}$, i.e, $\mathcal{A}$ is a ϕ-invariant proper free factor system such that there is no free factor system $\mathcal{A}' \succ \mathcal{A}$ for which $\mathcal{A}$ supports $\phi(\mathcal{A}')$; e.g., consider an injective-nonsurjective ϕ and its elliptic free factor system.

An $(F, \mathcal{A})$-tree is a simplicial tree T with no bivalent vertices along with a minimal F-action by isometries whose edge stabilizers are trivial and point stabilizers are trivial or conjugates of the free factors in $\mathcal{A}$. Any given $(F, \mathcal{A})$-tree endowed with the combinatorial metric has an associated length function $l_T : F \to \mathbb{R}$. Precisely, any isometry of a simplicial tree is either *elliptic* (fixes a point) or *loxodromic* (preserves an *axis* of least displacement). If $g \in F$ is elliptic, then $l_T(g) = 0$; otherwise, $l_T(g)$ is the translation distance of g acting on its axis. As this will be relevant in the next section, note that an $(F, \{F\})$-tree is a point.

For any $(F, \mathcal{A})$-tree T, the condition $\phi^{-1} \cdot \mathcal{A} = \mathcal{A}$ implies that *loxodromic* elements in F have loxodromic ϕ-images. Let T and T' be $(F, \mathcal{A})$- and $(F', \mathcal{A}')$-trees respectively and $\psi : F \to F'$ be an injective homomorphism such that $\psi^{-1} \cdot \mathcal{A}' = \mathcal{A}$. A graph map $f : T \to T'$ is **ψ-equivariant** if $f(g \cdot x) = \psi(g) \cdot f(x)$ for all $g \in F$ and $x \in T$. The first step is to establish the bounded cancellation lemma for trees.

Lemma 5.1 (Bounded Cancellation)**.** *Let T and T' be $(F, \mathcal{A})$- and $(F', \mathcal{A}')$-trees respectively and $\psi : F \to F'$ be an injective homomorphism such that $\psi^{-1} \cdot \mathcal{A}' = \mathcal{A}$. If $g : T \to T'$ is a ψ-equivariant graph map, then there is a constant $C(g)$ such that for every natural-edge-path decomposition $p_1 \cdot p_2$ of an immersed path in T, the edge-path $[g(p_1)] \cdot [g(p_2)]$ is contained in the $C(g)$-neighborhood of $[g(p_1) \cdot g(p_2)]$.*

Proof. The proof is exactly the same as before. The map g factors as an equivariant pretrivial edge collapse and subdivision, a composition of $r \geq 0$ equivariant folds, and an equivariant isometric embedding. We may choose $C(g) = r$. □

More generally, we want to consider pairs $\mathcal{A} \prec \mathcal{B}$ of ϕ-invariant free factor systems. A $(\mathcal{B}, \mathcal{A})$-forest T_* is a forest whose components are $(B_i, \mathcal{A}_i)$-trees T_i for each free factor $B_i \in \mathcal{B}$, where $\mathcal{A}_i$ is either the trivial free factor system or the nonempty maximal subset of $\mathcal{A}$ supported by $\{B_i\}$; this is essentially the relative analogue of $\mathcal{B}$-marked graphs. Let $\sigma : \{1, \dots l\} \to \{1, \dots l\}$ be the function corresponding to the ϕ-invariance of $\mathcal{B}$ and $\phi|_{\mathcal{B}} = \{\phi_i : B_i \to B_{\sigma(i)}\}$ be a restriction of ϕ. An **$\boldsymbol{\mathcal{A}}$-relative representative for $\boldsymbol{\phi|_{\mathcal{B}}}$** is a map $f_* : T_* \to T_*$ whose restrictions to components are ϕ_i-equivariant graph maps $f_i : T_i \to T_{\sigma(i)}$ with no pretrivial edges. Additionally, we shall say the relative representative is **minimal** if it has no *orbit-closed* invariant subforests with bounded components.

Let T_* be a $(\mathcal{B}, \mathcal{A})$-forest and $f_* : T_* \to T_*$ be a $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$. For all $k \geq 1$, set $T_*(\phi^k(\mathcal{B})) \subset T_*$ to be the minimal subforest for $\phi^k(\mathcal{B})$; these subforests are the relative analogues of Stallings graphs. We will assume the minimal subforests $T_*(\phi^k(\mathcal{B}))$ inherit their simplicial structure from the *ambient forest* T_* and so they might have bivalent vertices unlike T_*. For a graph of groups decomposition with bivalent vertices, **branch points** are vertices that are images of branch points of the *Bass-Serre* tree and natural edges are images of natural edges of the tree.

For any $k \geq 1$, let $f_{*,k} : T_*(\phi^k(\mathcal{B})) \to T_*$ be the restriction of f_* to $T_{*,k}$ and then replace it with an equivariantly homotopic map $T_*(\phi^k(\mathcal{B})) \to T_*(\phi^k(\mathcal{B}))$ that is induced by the deformation retraction of $f_{*,k}(T_*(\phi^k(\mathcal{B})))$ to $T_*(\phi^k(\mathcal{B}))$, which we suggestively call the **($\boldsymbol{k}$-th) homotopy restriction of $\boldsymbol{f_*}$**. Note that if $X \subset T_*(\phi^k(\mathcal{B}))$ is an axis such that $f_*|_X$ is an immersion, then $f_{*,k}|_X$ is still an immersion.

These homotopy restrictions are the relative analogues of homotopy lifts. As before, using an equivariant (tightening) homotopy supported in the interior of natural edges, we replace $f_{*,k} : T_*(\phi^k(\mathcal{B})) \to T_*(\phi^k(\mathcal{B}))$ with a homotopic map that maps any natural edge to either a vertex or an immersed path and has Lipschitz and cancellation constants $K(f_*)$ and $C(f_*)$. Once again, if $X \subset T_*(\phi^k(\mathcal{B}))$ is an axis such that $f_{*,k}|_X$ is an immersion before tightening, then it is still an immersion after tightening. The proof of the following lemma is also almost the same as that of Lemma 2.6.

Lemma 5.2. *Let $f_* : T_* \to T_*$ be an $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$. Then for any $k \geq 1$, if $f_{*,k} : T_*(\phi^k(\mathcal{B})) \to T_*(\phi^k(\mathcal{B}))$ is a homotopy restriction of f_* with $C(f_{*,k}) = C(f_*)$, then $f_{*,k}$ maps branch points to the $C(f_*)$-neighborhood of branch points.*

Proof. Set $C = C(f_{*,k}) = C(f_*)$ and let ν be a bivalent vertex in $T_*(\phi^k(\mathcal{B}))$ whose distance to the nearest branch point is greater than C. In particular, ν has a trivial stabilizer. We denote by ϵ_1, ϵ_2 the 2 distinct directions at ν oriented away from the vertex. Suppose v is a branch point of $T_*(\phi^k(\mathcal{B}))$ such that $f_{*,k}(v) = \nu$. So v has a trivial stabilizer under the action of $\phi^k(\mathcal{B})$. Choose 3 distinct directions at v: e_1, e_2, and e_3. Let p_{12} be an embedded path in $T_*(\phi^k(\mathcal{B}))$ that starts with e_1 and ends with a translate $\bar{e}_2$. Since v has a trivial stabilizer, the path determines a unique loxodromic element g_{12} in $\phi^k(\mathcal{B})$ with axis a_{12} such that p_{12} is a fundamental domain of the axis under the translation action of g_{12}. Without loss of generality, $[f_{*,k}(p_{12})]$ starts with ϵ_1.

If $[f_{*,k}(p_{12})]$ ends with the translate $\phi(g_{12})\bar{\epsilon}_1$, then $[f_*(p_{12})] = \mu \cdot \rho \cdot (\phi(g_{12})\bar{\mu})$, where μ is an extension of ϵ_1 to an embedded path from ν to the axis of $\phi(g_{12})$ and ρ is a fundamental domain of the axis of $\phi(g_{12})$. By assumption, μ is longer than C. Decompose the axis $a_{12} = a_- \cdot a_+$ at v, then $[f_{*,k}(a_-)] \cdot [f_{*,k}(a_+)]$ has $\bar{\mu} \cdot \mu$ as a subpath, violating bounded cancellation. The remaining cases are handled similarly. Upon ruling out all cases, we conclude that no branch point v of $T_*(\phi^k(\mathcal{B}))$ is mapped to ν. □

This lemma allows us to replace the current homotopy restriction with an equivariantly homotopic map $f_{*,k}$ that maps branch points to branch points, maps any natural edge to a branch point or an immersed path, and has Lipschitz and cancellation constants $K(f_{*,k}) = K(f_*) + C(f_*)$ and $C(f_{*,k}) = 2C(f_*)$ respectively. If $X \subset T_*(\phi^k(\mathcal{B}))$ is an axis and $f_{*,k}|_X$ is an immersion before the homotopy, then $f_{*,k}(X)$ is an immersed path after the homotopy; the restriction $f_{*,k}|_X$ may fail to be an immersion due to pretrivial edges.

Collapsing a maximal (orbit-closed) $f_{*,k}$-invariant subforest of $T_*(\phi^k(\mathcal{B}))$ with bounded components and forgetting the bivalent vertices induces a minimal $\phi^k(\mathcal{A})$-relative representative $g_{*,k} : Y_{*,k} \to Y_{*,k}$ for $\phi|_{\phi^k(\mathcal{B})}$ defined on a $(\phi^k(\mathcal{B}), \phi^k(\mathcal{A}))$-forest $Y_{*,k}$. Note that the collapsed maximal subforest contains the pretrivial edges as $\phi^{-1} \cdot \mathcal{A} = \mathcal{A}$. Since $g_{*,k}$ is induced by equivariantly collapsing a forest and forgetting bivalent vertices, we may use $K(g_{*,k}) = K(f_*) + C(f_*)$, $C(g_{*,k}) = 2C(f_*)$, and $l_{Y_{*,k}} \leq l_{T_*}|_{\phi^k(\mathcal{B})}$. If X is an axis of b in $T_*(\phi^k(\mathcal{B}))$ and $f_{*,k}|_X$ is an immersion modulo pretrivial edges, then $g_{*,k}|_{X'}$ is an immersion, where X' is the axis of b in $Y_{*,k}$. The following summary is an analogue of Proposition 2.7.

Lemma 5.3. *Let $f_* : T_* \to T_*$ be an $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$ and $k \geq 1$. There is a $(\phi^k(\mathcal{B}), \phi^k(\mathcal{A}))$-forest $Y_{*,k}$ (with no bivalent vertices) and a minimal $\phi^k(\mathcal{A})$-relative representative $g_{*,k} : Y_{*,k} \to Y_{*,k}$ for $\phi|_{\phi^k(\mathcal{B})}$ such that:*

1. *$l_{Y_{*,k}} : \phi^k(\mathcal{B}) \to \mathbb{R}$ is dominated by the restriction $l_{T_*}|_{\phi^k(\mathcal{B})} = l_{T_*(\phi^k(\mathcal{B}))}$;*

2. *$K(g_{*,k}) = K(f_*) + C(f_*)$ and $C(g_{*,k}) = 2C(f_*)$.*

3. *If f_* restricted to the axes of some conjugacy class $[b]$ in $\phi^k(\mathcal{B})$ is an immersion, then $g_{*,k}$ restricted to the axes of $[b]$ is also an immersion.*

For an $\mathcal{A}$-relative representative f_* for $\phi|_{\mathcal{B}}$, we define the **transition matrix** $A(f_*)$. Let $A(f_*)$ be a nonnegative integer square matrix whose rows and columns are indexed by the orbits of edges in T_*; the entry of $A(f_*)$ in row-i and column-j, $A(f_*)(i,j)$, is given by the number of translates of e_i that are contained in the immersed edge-path $f_*(e_j)$, where e_i is a orbit representative for the i-th orbit of edges. An $\mathcal{A}$-relative representative f_* is **irreducible** if the matrix $A(f_*)$ is irreducible, i.e., for any pair (i,j), there is a positive integer n_{ij} such that $A(f_*)^{n_{ij}}(i,j) > 0$. In this case, the **stretch factor** of f_*, $\lambda(f_*) \geq 1$, is the *Perron-Frobenius eigenvalue* of $A(f_*)$. An irreducible $\mathcal{A}$-relative representative is **expanding** if $\lambda(f_*) > 1$. Note that irreducible $\mathcal{A}$-relative representatives are minimal.

We say $\phi|_{\mathcal{B}}$ is **irreducible relative to $\mathcal{A}$** if there is no ϕ-invariant free factor system $\mathcal{C}$ such that $\mathcal{A} \prec \mathcal{C} \prec \mathcal{B}$. If $\mathcal{A}$ is trivial and $\mathcal{B} = \{F\}$, then we recover the definition of ϕ's irreducibility. The next lemma is the most useful property of an irreducible restriction for our purposes.

Lemma 5.4. *If $\phi|_{\mathcal{B}}$ is irreducible relative to $\mathcal{A}$, then every minimal $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$ is irreducible.*

Proof. Suppose some minimal $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$ has a reducible transition matrix; in particular, it has an invariant $\mathcal{B}$-equivariant proper subforest (with unbounded components) that determines a ϕ-invariant free factor system $\mathcal{C}$ such that $\mathcal{A} \prec \mathcal{C} \prec \mathcal{B}$. So $\phi|_{\mathcal{B}}$ is not irreducible relative to $\mathcal{A}$. □

Remark. Bestvina-Handel give the absolute version of this property as the definition of irreducibility and then prove that it is equivalent to the definition of irreducibility given in this dissertation [3, Lemma 1.16]. The relative version of this equivalence holds as well but we will not prove it as it is not needed.

Bestvina-Handel used the next proposition to construct *train tracks* [3, Theorem 1.7].

Proposition 5.5. *If $\phi|_{\mathcal{B}}$ is irreducible relative to $\mathcal{A}$, then there is an irreducible $\mathcal{A}$-relative representative $f_* : T_* \to T_*$ for $\phi|_{\mathcal{B}}$ with the minimal stretch factor, i.e., if $f'_* : T'_* \to T'_*$ is an irreducible $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$, then $\lambda(f'_*) \geq \lambda(f_*)$.*

The minimal stretch factor will be denoted by $\lambda(\phi, \mathcal{B}, \mathcal{A})$.

Proof. Let $g_* : Y_* \to Y_*$ be a minimal $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$ and suppose $\phi|_{\mathcal{B}}$ is irreducible relative to $\mathcal{A}$. Then g_* is an irreducible $\mathcal{A}$-relative representative by Lemma 5.4 with stretch factor $\lambda(g_*)$. By the lack on bivalent vertices, any irreducible $\mathcal{A}$-relative representative has a transition matrix of size $\leq N = 3 \cdot \operatorname{rank}(F) - 3$. Suppose B is an irreducible integer square matrix with Perron-Frobenius eigenvalue $\lambda(B) \leq \lambda(g_*)$. Then B

has a positive right eigenvector $\vec{v}$ associated with $\lambda(B)$. So for all $k \geq 1$, B^k has right eigenvector $\vec{v}$ associated with eigenvalue $\lambda(B)^k$. Assuming the smallest entry of $\vec{v}$ is 1 (rescale if necessary), we get that the minimum row-sum of B^k is at most $\lambda(B)^k$ for any $k \geq 1$. If B has no more than N rows, then the largest entry of B is at most the minimum row-sum of $B^{N!}$, which we know is at most $\lambda(B)^{N!} \leq \lambda(g_*)^{N!}$. So there are finitely many irreducible integer square matrices with size $\leq N$ and Perron-Frobenius eigenvalue $\leq \lambda(g_*)$. Thus, there is a finite set of stretch factors $\leq \lambda(g_*)$ for irreducible $\mathcal{A}$-relative representatives for $\phi|_{\mathcal{B}}$. In particular, there is an irreducible $\mathcal{A}$-relative representative $f_* : T_* \to T_*$ for $\phi|_{\mathcal{B}}$ with the minimal stretch factor. □

Bestvina-Handel's work [3] can be adapted to show that an irreducible $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$ with the minimal stretch factor is in fact an *$\mathcal{A}$-relative train track* (See Appendix C) and, conversely, bounded cancellation implies all irreducible $\mathcal{A}$-relative train tracks for $\phi|_{\mathcal{B}}$ have the minimal stretch factor. We do not prove this converse as it is not needed. The next lemma is an application of train track theory that will be invoked once, in the second half of the proof of Proposition 5.7.

Lemma 5.6. *If $\phi|_{\mathcal{B}}$ is irreducible relative to $\mathcal{A}$ and $f_* : T_* \to T_*$ is an irreducible $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$ with the minimal stretch factor, then there is an element g in $\mathcal{B}$ with an axis a_g such that the restriction of f_*^k to a_g is an immersion for all $k \geq 1$.*

Such an axis will be known as an **f_*-legal axis**.

Proof. If $\lambda(f_*) = 1$, then f_* is a simplicial embedding and we are done. So we may assume $\lambda(f_*) > 1$. By minimality of its stretch factor, f_* is an expanding irreducible $\mathcal{A}$-relative *train track* for $\phi|_{\mathcal{B}}$ (Theorem C.1), i.e., for any edge e in T_*, $f_*^k(e)$ is an expanding immersed path for all $k \geq 1$. A 2-edge path $e_1 \cdot e_2$ is *f_*-legal* if it is a translate of a subpath of $f_*^k(e)$ for some edge e and integer $k \geq 1$. By irreducibility of f_*, every edge e is contained in a 3-edge path $e_- \cdot e \cdot e_+$ whose 2-edge subpaths are both f_*-legal. This means we can form an axis a_g whose 2-edge subpaths are all f_*-legal. By the train track property, the restriction of f_*^k to a_g is an immersion for all $k \geq 1$. $\square$

The main tools from Section 2 that were used in the previous section were Lemma 2.2, bounded cancellation, and Proposition 2.7. The relative analogues of the latter two have been established in this section. We now state the main technical result of this section, an analogue of Lemma 2.2 — analogous in the sense that both give sufficient conditions for iterated subgroup graphs to have arbitrarily long natural edges.

Proposition 5.7. *Let $\mathcal{A} \prec \mathcal{B}$ be a chain of ϕ-invariant free factor systems that support the ϕ-elliptic free factor system. If $\phi|_\mathcal{B}$ is irreducible relative to $\mathcal{A}$ and $\lambda(\phi, \mathcal{B}, \mathcal{A}) > 1$, then the length of the longest natural edge in $T_*(\phi^k(\mathcal{B}))$ is unbounded as $k \to \infty$.*

Before starting the proof, we will first describe the *(absolute) vertex blow-up* construction. Let $f_* : T_* \to T_*$ be a $\mathcal{A}$-relative representative for $\phi|_\mathcal{B}$ defined on some $(\mathcal{B}, \mathcal{A})$-forest T_*. Replace the free factors of $\mathcal{A}_i \subset \mathcal{A}$ with conjugates if necessary and assume they are subgroups of $B_i \in \mathcal{B}$. Fix some $\mathcal{A}$-marked roses $(R_\mathcal{A}, \alpha_\mathcal{A})$ and a topological representative $f_\mathcal{A} : R_\mathcal{A} \to R_\mathcal{A}$ for $\phi|_\mathcal{A}$. Define $\Gamma_\mathcal{B}$ to be the graph formed by identifying the appropriate vertices of the graph of groups $\mathcal{B}\backslash T_*$ with the basepoints of roses $(R_\mathcal{A}, \alpha_\mathcal{A})$. If $c : R_\mathcal{A} \to \Gamma_\mathcal{B}$ is the inclusion map, then *Bass-Serre theory* gives markings $\alpha_\mathcal{B} = \{\alpha_i : B_i \to \pi_1(\Gamma_i)\}$ such that $\pi_1(c) \circ \alpha_\mathcal{A} = (\alpha_\mathcal{B})|_\mathcal{A}$. Thus, $(\Gamma_\mathcal{B}, \alpha_\mathcal{B})$ is a $\mathcal{B}$-marked graph. This construction and, in general any pair of graphs $\Gamma'_\mathcal{A} \subset \Gamma'_\mathcal{B}$ with collections of markings $\alpha'_\mathcal{A}, \alpha'_\mathcal{B}$ such that $\pi_1(c') \circ \alpha'_\mathcal{A} = (\alpha'_\mathcal{B})|_\mathcal{A}$ will be referred to as **(vertex) blow-up**.

We note that the Stallings graph $S(\phi^k(\mathcal{B}))$ with respect to $(\Gamma_\mathcal{B}, \alpha_\mathcal{B})$, as a $\phi^k(\mathcal{B})$-marked graph, is a blow-up of $\phi^k(\mathcal{B})\backslash T_*(\phi^k(\mathcal{B}))$: let $\iota : S(\phi^k(\mathcal{B})) \to \Gamma_\mathcal{B}$ be the Stallings graph's immersion and $S_\mathcal{A} \subset S(\phi^k(\mathcal{B}))$ be the core of the subgraph $\iota_\mathcal{B}^{-1}(R_\mathcal{A})$. Since $\phi^{-1} \cdot \mathcal{A} = \mathcal{A}$, $S_\mathcal{A} = S(\phi^k(\mathcal{A}))$ is marked by an isomorphism $\alpha'_\mathcal{A} : \phi^k(\mathcal{A}) \to \pi_1(S_\mathcal{A})$ and $\alpha'_\mathcal{A}$ is the restriction of the marking $\alpha'_\mathcal{B} : \phi^k(\mathcal{B}) \to \pi_1(S(\phi^k(\mathcal{B})))$ to $\phi^k(\mathcal{A})$ with respect to the inclusion $S_\mathcal{A} \subset S(\phi^k(\mathcal{B}))$. Therefore, $S(\phi^k(\mathcal{B}))$ is also a vertex blow-up of $\phi^k(\mathcal{B})\backslash Y_{*,k}$. The noncontractible components of the subgraph $\iota_{\mathcal{B}}^{-1}(R_{\mathcal{A}})$ will be known as the *lower stratum* and the rest of the graph as the *top stratum*.

Construct a topological representative $f_{\mathcal{B}} : \Gamma_{\mathcal{B}} \to \Gamma_{\mathcal{B}}$ for $\phi|_{\mathcal{B}}$ that agrees with $f_{\mathcal{A}}$ on the $\mathcal{A}$-marked roses $R_{\mathcal{A}}$ and induces f_* on the Bass-Serre forest T_* upon collapsing the roses $R_{\mathcal{A}}$. For any $k \geq 1$, we let $g_{*,k} : Y_{*,k} \to Y_{*,k}$ be the minimal $\phi^k(\mathcal{A})$-relative representatives for $\phi|_{\phi^k(\mathcal{B})}$ given by Lemma 5.3 and $\bar{f}_k : S(\phi^k(\mathcal{B})) \to S(\phi^k(\mathcal{B}))$ be the homotopy lift of $f_{\mathcal{B}}$ given by Proposition 2.7. By Lemma 5.3(3), if an element b in $\mathcal{B}$ has an f_*-legal axis, then $\phi^k(b)$ has a $g_{*,k}$-legal axis. It can be arranged for $S(\phi^k(\mathcal{A})) \subset S(\phi^k(\mathcal{B}))$ to be $\bar{f}_k$-invariant and $\bar{f}_k$ to induce $g_{*,k}$ on the $(\phi^k(\mathcal{B}), \phi^k(\mathcal{A}))$-forest $Y_{*,k}$ upon collapsing a maximal invariant proper subgraph containing $S(\phi^k(\mathcal{A}))$ and forgetting bivalent vertices.

Now for the idea behind the proof. By irreducibility of the restriction $\phi|_{\mathcal{B}}$, we may assume the map $g_{*,k}$ is an expanding irreducible representative for $\phi|_{\phi^k(\mathcal{B})}$. Suppose the forests $T_*(\phi^k(\mathcal{B}))$ had uniformly bounded natural edges. Then there is a sequence of loxodromic elements b_k in $\phi^k(\mathcal{B})$ with uniformly bounded translation length $l_{T_*}(g_k)$. Now suppose that the vertex blow-up $S(\phi^k(\mathcal{B}))$ had natural edges with aribtrarily long top stratum subpaths. Bounded cancellation, the fact $\bar{f}_k$ induces $g_{*,k}$, and the irreducibility of $g_{*,k}$ imply $g_{*,k}$ is an expanding irreducible *immersion*. However, this contradicts the first assumption since $l_{Y_{*,k}} \leq l_{T_*(\phi^k(\mathcal{B}))}$. So the second supposition is false and the natural edges of $S(\phi^k(\mathcal{B}))$ have top stratum subpaths with uniformly bounded length. Using the Lipschitz property, expanding irreducibility of $g_{*,k}$, and existence of a $g_{*,k}$-legal axis (train track theory), we find uniformly bounded lower stratum paths in $S(\phi^k(\mathcal{B}))$ connecting the origin of any oriented top stratum subpath of a natural edge to another top stratum subpath of a natural edge. Consequently, we are able to build uniformly bounded immersed loops in $S(\phi^k(\mathcal{B}))$ that contain top stratum subpaths. This implies some loxodromic conjugacy class in $\mathcal{B}$ has an infinite ϕ-tail, which contradicts Proposition 3.3: any conjugacy class with an infinite ϕ-tail is elliptic. So the first supposition is false too and the forests $T_*(\phi^k(\mathcal{B}))$ have arbitrarily long natural edges for $k \geq 1$.

Proof of Proposition 5.7. Suppose $\mathcal{A} \prec \mathcal{B}$ are ϕ-invariant free factor systems that supports the ϕ-elliptic free factor system, $\phi|_{\mathcal{B}}$ is irreducible relative to $\mathcal{A}$, $f_{\mathcal{A}} : R_{\mathcal{A}} \to R_{\mathcal{A}}$ is a topological representative for $\phi|_{\mathcal{A}}$ defined on $\mathcal{A}$-marked roses $(R_{\mathcal{A}}, \alpha_{\mathcal{A}})$, and $f_* : T_* \to T_*$ is an expanding irreducible $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$ with the minimal stretch factor $\lambda(f_*) > 1$ (Proposition 5.5). By Lemma 5.6, there is an element b in $\mathcal{B}$ with an f_*-legal axis. Set $(\Gamma_{\mathcal{B}}, \alpha_{\mathcal{B}})$ to be the vertex blow-up of $\mathcal{B}\backslash T_*$ with respect to the $\mathcal{A}$-marked roses $(R_{\mathcal{A}}, \alpha_{\mathcal{A}})$. The discussion preceding the proof gives minimal $\phi^k(\mathcal{A})$-relative representatives $g_{*,k} : Y_{*,k} \to Y_{*,k}$ for $\phi|_{\phi^k(\mathcal{B})}$ and homotopy lifts $\bar{f}_k : S(\phi^k(\mathcal{B})) \to S(\phi^k(\mathcal{B}))$ of a topological representative $f_{\mathcal{B}} : \Gamma_{\mathcal{B}} \to \Gamma_{\mathcal{B}}$ for $\phi|_{\mathcal{B}}$ which have these properties: for all $k \geq 1$,

1. $l_{Y_{*,k}} : \phi^k(\mathcal{B}) \to \mathbb{R}$ is dominated by the restrictions $l_{T_*}|_{\phi^k(\mathcal{B})} = l_{T_*(\phi^k(\mathcal{B}))}$;
2. $\phi^k(b)$ has a $g_{*,k}$-legal axis;
3. $\bar{f}_k$ maps branch points to branch points and any natural edge to a branch point or an immersed path;
4. $K = K(\bar{f}_k) = K(f_{\mathcal{B}}) + C(f_{\mathcal{B}})$ and $C = C(\bar{f}_k) = 2C(f_{\mathcal{B}})$;
5. $\bar{f}_k$ induces $g_{*,k}$ on $Y_{*,k}$ upon collapsing $R_{\mathcal{A}} \subset \Gamma_{\mathcal{B}}$.

The collection $\phi|_{\phi^k(\mathcal{B})}$ is conjugate to $\phi|_{\mathcal{B}}$ by injectivity of ϕ. So $\phi|_{\phi^k(\mathcal{B})}$ is irreducible relative to $\phi^k(\mathcal{A})$ and $\lambda(f_*)$ is the minimal stretch factor for $\phi|_{\phi^k(\mathcal{B})}$ relative to $\phi^k(\mathcal{A})$. Furthermore, the minimal $\phi^k(\mathcal{A})$-relative representatives $g_{*,k}$ are irreducible (Lemma 5.4) and $\lambda(g_{*,k}) \geq \lambda(f_*) > 1$ by the minimality of $\lambda(f_*)$.

Suppose, for a contradiction, there is a bound $L \geq 1$ such that all natural edges in $T_*(\phi^k(\mathcal{B}))$ are shorter than L for all $k \geq 1$. Then, for all $k \geq 1$, there is a loxodromic element b_k in $\phi^k(\mathcal{B})$ such that $l_{T_*}(b_k) \leq (3N-3)L$, where $N = 3 \cdot \mathrm{rank}(F) - 3$. Every edge E in $\Gamma_{*,k} = \phi^k(\mathcal{B})\backslash Y_{*,k}$ lifts to a $\phi^k(\mathcal{B})$-orbit of a natural edge E' in $\phi^k(\mathcal{B})\backslash T_*(\phi^k(\mathcal{B}))$, which corresponds to a top stratum subpath $\bar{E}$ of a natural edge in $S(\phi^k(\mathcal{B}))$.

Claim. *The subpath $\bar{E}$ in $S(\phi^k(\mathcal{B}))$ has length $\leq C \cdot K^{N-1}$ for all edges E in $\Gamma_{*,k}$ and $k \geq 1$.*

Suppose, the graph $\Gamma_{*,k}$ has an edge E_0 whose corresponding subpath $\bar{E}_0$ in $S(\phi^k(\mathcal{B}))$ is longer than $C \cdot K^{N-1}$ for some $k \geq 1$. As we did in the proof of Theorem 3.1, we construct the set of *long* edges $\mathcal{E}$ by looking at all the edges of $\Gamma_{*,k}$ that are eventually mapped over E_0. Here, an edge E_1 in $\Gamma_{*,k}$ mapped over E_0 if there are lifts E_1' and E_0' in $Y_{*,k}$ such that $g_{*,k}$ maps E_1' over E_0'. Since $\bar{f}_k : S(\phi^k(\mathcal{B})) \to S(\phi^k(\mathcal{B}))$ is K-Lipschitz and it induces $g_{*,k}$ on $Y_{*,k}$, each long edge in $\Gamma_{*,k}$ corresponds to a top statrum subpath in $S(\phi^k(\mathcal{B}))$ longer than C. Since $g_{*,k}$ is an irreducible $\phi^k(\mathcal{A})$-relative representative, all edges eventually map over E_0 and hence are long. The long natural edges of $S(\phi^k(\mathcal{B}))$ will be the natural edges in $S(\phi^k(\mathcal{B}))$ containing top stratum subpaths.

Suppose an edge E of $\Gamma_{*,k}$ had a lift E' in $Y_{*,k}$ that is the initial segment of the $g_{*,k}$-image of two edges that share an initial vertex. Then the top stratum subpath $\bar{E}$ is in a long natural edge of $S(\phi^k(\mathcal{B}))$ that is the initial segment of $\bar{f}_k$-images of natural edges that share an initial vertex; this violates bounded cancellation since long natural edges of $S(\phi^k(\mathcal{B}))$ longer than $C = C(\bar{f}_k)$. Hence, there is no folding in $g_{*,k}$, i.e., $g_{*,k}$ is an expanding irreducible $\phi^k(\mathcal{A})$-relative *immersion*. We may now find an $m \geq 1$ such that all loxodromic elements b in $\phi^k(\mathcal{B})$ have $l_{Y_{*,k}}(\phi^m(b)) > (3N-3)L$. Since $l_{Y_{*,k}}$ is dominated by $l_{T_*}|_{\phi^k(\mathcal{B})}$, we get that $l_{T_*}(b') > (3N-3)L$ for all loxodromic elements b' in $\phi^{k+m}(\mathcal{B})$. This contradicts the assumption that $l_{T_*}(b_{k+m}) \leq (3N-3)L$ for some loxodromic b_{k+m} in $\phi^{k+m}(\mathcal{B})$. So the top stratum subpath $\bar{E}$ in $S(\phi^k(\mathcal{B}))$ has length $\leq C \cdot K^{N-1}$ for all natural edges E of $\Gamma_{*,k}$ and $k \geq 1$. This ends the proof of the claim.

Next, we prove the existence of lower stratum paths in $S(\phi^k(\mathcal{B}))$ with uniformly bounded lengths connecting top stratum paths. Suppose E_0, E_1, and E_2 are edges of $\Gamma_{*,k}$ with lifts E_0', E_1', and E_2' in $Y_{*,k}$ such that $E_1' \cdot E_2'$ is a subpath of the immersed path $g_{*,k}(E_0')$. Then $\bar{E}_1 \cdot P_{12} \cdot \bar{E}_2$ is a subpath of immersed path $\bar{f}_k(\bar{E}_0)$ for some lower stratum path P_{12} in $S(\phi^k(\mathcal{B}))$. Since $\bar{E}_0$ has length bounded by $C \cdot K^{N-1}$ and $\bar{f}_k$ is K-Lipschitz, the path P_{12} has length bounded by $C \cdot K^N$. We say the 2-edge path $E_1 \cdot E_2$ in $\Gamma_{*,k}$ has a *nondegenerate turn bounded by* $C \cdot K^N$.

As $g_{*,k}$ is an expanding irreducible relative representative that has a legal axis (this is part of the argument that invokes train track theory), every edge E' in $Y_{*,k}$ can be extended to an immersed 3-edge path $E'_- \cdot E' \cdot E'_+$ that is a translate of a subpath of $g^n_{*,k}(E')$ and $n \leq 2{\cdot}N!$. In particular, any edge in $\Gamma_{*,k}$ can be extended to a 3-edge path whose 2-edge subpaths both have nondegenerate turns bounded by $C \cdot K^{N-1} \cdot K^{2 \cdot N!}$, i.e., every top stratum subpath $\bar{E}$ can be extended to an immersed path $\bar{E}_- {\cdot} P_- {\cdot} \bar{E} {\cdot} P_+ {\cdot} \bar{E}_+$ with top stratum subpaths $\bar{E}_-$, $\bar{E}_+$ and lower stratum paths P_-, P_+ with length bounded by $C \cdot K^{N-1} \cdot K^{2 \cdot N!}$.

Using this bound on lower stratum paths and the bound on top stratum subpaths given by the claim, we can now form an immersed loop ρ_k in $\Gamma_{*,k}$ with the properties:

1. ρ_k lifts to an axis in $Y_{*,k}$ for some loxodromic conjugacy class $[b'_k]$;

2. ρ_k passes any edge of $\Gamma_{*,k}$ at most twice and only takes short turns (including the turn at the endpoint), which implies it has at most $2N$ edges and (short) turns; and

3. ρ_k corresponds to a loop in $S(\phi^k(\mathcal{B}))$ with length bounded by $2N{\cdot}C(1 + K^{2 \cdot N!})K^{N-1}$.

In summary, for each $k \geq 1$, we find a loxodromic conjugacy class $[b'_k]$ in $\phi^k(\mathcal{B})$ with $\alpha_{\mathcal{B}}$-length bounded by a constant independent of k. As there are finitely many conjugacy classes with $\alpha_{\mathcal{B}}$-length bounded by any given constant, the sequence of conjugacy classes $([b'_k])_{k=1}^{\infty}$ has a constant infinite subsequence. Thus, some loxodromic conjugacy class $[b']$ has an infinite ϕ-tail supported in $\mathcal{B}$. This is the contradiction that ends the proof — recall that the free factor system $\mathcal{A}$ supports the maximal ϕ-fixed free factor system by hypothesis and the latter supports all conjugacy classes with an infinite ϕ-tail (Proposition 3.3); on the other hand, $\mathcal{A}$ does not support loxodromic conjugacy classes in $\mathcal{B}$ by definition. □

6 Relative immersions

The main result of this section is the existence of expanding immersions for nonsurjective endomorphisms relative to their elliptic free factor systems.

Let $\phi : F \to F$ be an injective endomorphism, $\mathcal{A} \prec \mathcal{B}$ be a pair of ϕ-invariant free factor systems such that $\phi^{-1} \cdot \mathcal{A} = \mathcal{A}$, and $\phi|_{\mathcal{B}} : \mathcal{B} \to \mathcal{B}$ be a restriction of ϕ to $\mathcal{B}$. An **$\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}}$** is an $\mathcal{A}$-relative representative $f_* : T_* \to T_*$ for $\phi|_{\mathcal{B}}$ that is also an immersion. An $\mathcal{A}$-relative immersion f_* for $\phi|_{\mathcal{B}}$ is **expanding** if it is minimal and every edge expands under f_*-iteration. Recall that a relative representative is minimal if it has no orbit-closed invariant subforests with bounded components.

There will be two possible ways of obtaining a relative immersion for a relatively irreducible restriction with the minimal stretch factor λ. If $\lambda = 1$, then an irreducible representative with stretch factor λ is automatically an isometric immersion. The next proposition shows how to construct an immersion when $\lambda > 1$. This construction is unique to nonsurjective endomorphisms because we require that the restriction be irreducible relative to a free factor system that supports the ϕ-elliptic free factor system — when ϕ is an automorphism, the ϕ-elliptic free factor system is $\{F\}$ and no such restriction exists.

Proposition 6.1. *Let $\phi : F \to F$ be injective and $\mathcal{A} \prec \mathcal{B}$ be a chain of ϕ-invariant free factor systems that support the ϕ-elliptic free factor system. If $\phi|_{\mathcal{B}}$ is irreducible relative to $\mathcal{A}$ and $\lambda(\phi, \mathcal{B}, \mathcal{A}) > 1$, then there is an expanding irreducible $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}}$.*

Proof. Suppose $\phi : F \to F$ is injective, $\mathcal{A} \prec \mathcal{B}$ are ϕ-invariant free factor systems that supports the ϕ-elliptic free factor system, $\phi|_{\mathcal{B}}$ is irreducible relative to $\mathcal{A}$, and $f_* : T_* \to T_*$ is an expanding irreducible $(\mathcal{B}, \mathcal{A})$-relative representative for $\phi|_{\mathcal{B}}$ with minimal stretch factor $\lambda(f_*) > 1$. Set $K = K(f_*) + C(f_*)$ and $C = 2C(f_*)$. Recall that, for all $k \geq 1$, there is a homotopy restriction $f_{*,k} : T_*(\phi^k(\mathcal{B})) \to T_*(\phi^k(\mathcal{B}))$ that maps branch points to branch points, maps any natural edge to a branch point or an immersed path, and has Lipschitz and cancellation constants $K(f_{*,k}) = K$ and $C(f_{*,k}) = C$.

The first part of the proof proceeds as a relativized version of the proof of Theorem 3.1. By Proposition 5.7, we may fix $k \gg 0$ such that the set of natural edges $\mathcal{L}_0$ in $T_*(\phi^k(\mathcal{B}))$ longer than $C \cdot K^{N-1}$ is nonempty, where $N = 3 \cdot \text{rank}(F) - 3$. Choose $\mathcal{L}$ to be the set of all natural edges that eventually get mapped over those in $\mathcal{L}_0$ by $f_{*,k}$ and call $\mathcal{L}$ the long natural edges. As $f_{*,k}$ is K-Lipschitz and there are at most N orbits of natural edges in $T_*(\phi^k(\mathcal{B}))$, the long natural edges are longer that C. Injectivity of ϕ implies $\phi|_{\phi^k(\mathcal{B})}$ is conjugate to $\phi|_{\mathcal{B}}$. So $\phi|_{\phi^k(\mathcal{B})}$ is irreducible relative to $\phi^k(\mathcal{A})$, $\lambda(f_*)$ is the minimal stretch factor for $\phi|_{\phi^k(\mathcal{B})}$ relative to $\phi^k(\mathcal{A})$, and the short natural edges of $T_*(\phi^k(\mathcal{B}))$ form an orbit-closed $f_{*,k}$-invariant subforest with bounded components.

Collapse a maximal $f_{*,k}$-invariant subforest of $T_*(\phi^k(\mathcal{B}))$ that has bounded components and contains the short natural edges then forget the bivalent vertices; this induces a minimal $\phi^k(\mathcal{A})$-relative representative $g_{*,k} : Y_{*,k} \to Y_{*,k}$ for $\phi|_{\phi^k(\mathcal{B})}$. The map $g_{*,k}$ is irreducible $\phi^k(\mathcal{A})$-relative representatives for $\phi|_{\phi^k(\mathcal{B})}$ (Lemma 5.4) and $\lambda(g_{*,k}) \geq \lambda(f_*)$ by the minimality of $\lambda(f_*)$. So $g_{*,k}$ is an expanding irreducible $\phi^k(\mathcal{A})$-relative representative.

Since the lifts in $T_*(\phi^k(\mathcal{B}))$ of all edges in $Y_{*,k}$ are longer than the cancellation constant C, there is no folding in $g_{*,k}$ — otherwise, there would be folding in $f_{*,k}$ identifying paths longer than its cancellation constant, absurd. Thus, $g_{*,k}$ is an expanding irreducible $\phi^k(\mathcal{A})$-relative immersion for $\phi|_{\phi^k(\mathcal{B})}$. By injectivity of ϕ, we can view $Y_{*,k}$ as a $(\mathcal{B}, \mathcal{A})$-forest and $g_{*,k}$ as an expanding irreducible $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}}$. □

We are now ready to state and prove our base case for the construction. In light of the previous proposition, the point is that a restriction $\phi|_{\mathcal{B}}$ that is irreducible relative to the ϕ-elliptic free factor system $\mathcal{A}$ will satisfy $\lambda(\phi, \mathcal{B}, \mathcal{A}) > 1$.

Proposition 6.2. *Let $\phi : F \to F$ be injective and $\mathcal{A} \prec \mathcal{B}$ be a chain of ϕ-invariant free factor systems where $\mathcal{A}$ is the ϕ-elliptic free factor system. If $\phi|_{\mathcal{B}}$ is irreducible relative to $\mathcal{A}$, then there is an expanding irreducible $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}}$.*

Proof. Let $\phi : F \to F$ be injective and $\phi|_{\mathcal{B}}$ be irreducible relative to the ϕ-elliptic free factor system $\mathcal{A}$. Then there is an irreducible $\mathcal{A}$-relative representative $f_* : T_* \to T_*$ for $\phi|_{\mathcal{B}}$ with stretch factor $\lambda(f_*) = \lambda(\phi, \mathcal{B}, \mathcal{A}) \geq 1$ (Proposition 5.5). We say $B_i \in \mathcal{B}$ is loxodromic if $T_i \in T_*$ is not a point, i.e., B_i contains a loxodromic element; similarly, the component $B_i \backslash T_i$ of the graph of groups $\mathcal{B} \backslash T_*$ is loxodromic if B_i is loxodromic. If $\lambda(f_*) = 1$, then the induced map on the loxodromic components of $\mathcal{B} \backslash T_*$ is a graph isomorphism. So for some $k \geq 1$, if $A \in \mathcal{A}$ is supported by a loxodromic $B \in \mathcal{B}$, then A is ϕ^k-invariant. By Proposition 3.4 and Lemma 3.5, the set of all $A \in \mathcal{A}$ supported by loxodromic free factors of $\mathcal{B}$ form a ϕ-fixed free factor subsystem. As f_* induces a graph isomorphism on the loxodromic components of $\mathcal{B} \backslash T_*$ and these components' set of vertex groups is a ϕ-fixed free factor system, we get that f_* is surjective when restricted to the unbounded components of the forest T_* and the loxodromic free factors of $\mathcal{B}$ form a ϕ-fixed free factor system. This is a contradiction since ϕ-periodic free factors are elliptic (Proposition 3.3). Therefore, $\lambda(f_*) > 1$ and the result follows from Proposition 6.1. □

Specializing this proposition to the case where ϕ is irreducible and nonsurjective gives an alternate proof to a result due to Reynolds.

Corollary 6.3 ([12, Corollary 3.23])**.** *If $\phi : F \to F$ is irreducible and nonsurjective, then ϕ is induced by an expanding irreducible graph immersion.*

The proof of Proposition 6.2 given here can be thought of as a relativization of our previous proof [9, Theorem 4.5] of Reynolds' result but it does have one crucial difference: there is no mention here of *limit trees* in the compactification of outer space.

The next proposition is the induction step for our construction.

Proposition 6.4. *Let $\phi : F \to F$ be injective, $\mathcal{A}$ be the ϕ-elliptic free factor system, and $\mathcal{A} \prec \mathcal{B} \prec \mathcal{C}$ be a chain of ϕ-invariant free factor systems. If there is an expanding $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}}$ and a $\mathcal{B}$-relative immersion for $\phi|_{\mathcal{C}}$, then there is an expanding $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{C}}$.*

Although the proof gets a bit technical, the idea is rather simple: a $\mathcal{B}$-relative immersion for $\phi|_{\mathcal{C}}$ (top stratum) and an expanding $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}}$ (lower stratum) can be patched together via a *(relative) vertex blow-up* to get a minimal $\mathcal{A}$-relative representative $g_* : Y_* \to Y_*$ whose only possible folds would have to happen between a top and lower stratum edge of Y_*. As the restriction of g_* to the lower stratum is an expanding immersion, we may assume the edges in the lower stratum are longer than the cancellation constant. This means no lower statum edge is identified by a fold and so no folding in g_* is possible. Thus, g_* is a minimal $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{C}}$, which will be expanding if $\mathcal{A}$ is the ϕ-elliptic free factor system.

Let us now describe the vertex blow-up of a $(\mathcal{C}, \mathcal{B})$-forest $T_{\mathcal{C}}$ with respect to a $(\mathcal{B}, \mathcal{A})$-forest $T_{\mathcal{B}}$. For any free factor $C_i \in \mathcal{C}$, let $\mathcal{B}_i$ the maximal subset of $\mathcal{B}$ that is supported by C_i. Replace the free factors of $\mathcal{B}$ with conjugates if necessary and assume the free factors $B \in \mathcal{B}_i$ are subgroups of the free factors $C_i \in \mathcal{C}$. Identifying the appropriate vertices of the graph of groups $\mathcal{C} \backslash T_{\mathcal{C}}$ with basepoints on the graph of groups $\mathcal{B} \backslash T_{\mathcal{B}}$ results in a graph of groups decomposition for $\mathcal{C}$ whose Bass-Serre forest T_* is a $(\mathcal{C}, \mathcal{A})$-forest that contains $T_{\mathcal{B}}$ as a subforest.

Proof of Proposition 6.4. Let $\phi : F \to F$ be injective, $\mathcal{B} \prec \mathcal{C}$ be a chain of ϕ-invariant free factor systems that support the ϕ-elliptic free factor system $\mathcal{A}$. Suppose $f_{\mathcal{B}} : T_{\mathcal{B}} \to T_{\mathcal{B}}$ is an expanding $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}}$ and $f_{\mathcal{C}} : T_{\mathcal{C}} \to T_{\mathcal{C}}$ is a $\mathcal{B}$-relative immersion for $\phi|_{\mathcal{C}}$ then define T_* to be the vertex blow-up of $T_{\mathcal{C}}$ with respect to $T_{\mathcal{B}}$. The edges of T_* are of two types: the lower stratum, which are edges that are contained in the $\mathcal{C}$-orbit of $T_{\mathcal{B}}$, and the top stratum, which are the remaining edges.

Let $f_* : T_* \to T_*$ be a minimal $\mathcal{A}$-relative representative for $\phi|_{\mathcal{C}}$ that agrees with $f_{\mathcal{B}}$ on $T_{\mathcal{B}}$ and induces $f_{\mathcal{C}}$ upon collapsing the lower stratum. For all $k \geq 1$, set $T_*(\phi^k(\mathcal{C}))$ and $T_*(\phi^k(\mathcal{B}))$ be the minimal subforests of T_* for $\phi^k(\mathcal{C})$ and $\phi^k(\mathcal{B})$ respectively. Similarly, define the minimal subforest $T_{\mathcal{B}}(\phi^k(\mathcal{B})) \subset T_{\mathcal{B}}$. By the inclusion of $T_{\mathcal{B}}$ in T_*, we get an isometric identification of $T_*(\phi^k(\mathcal{B}))$ with $T_{\mathcal{B}}(\phi^k(\mathcal{B}))$. However, we want to consider these two forests differently with respect to the simplicial structure of branch points and natural edges. In particular, there may be branch points of $T_*(\phi^k(\mathcal{C}))$ that are bivalent when considered as points on the subforest $T_*(\phi^k(\mathcal{B}))$. So by *"natural edges of* $T_*(\phi^k(\mathcal{B}))$*"*, we mean those inherited from the parent forest $T_*(\phi^k(\mathcal{C}))$; on the other hand, by *"natural edges of* $T_{\mathcal{B}}(\phi^k(\mathcal{B}))$*"*, we do mean exactly that. Under the identification of the two forests, the natural edges of $T_*(\phi^k(\mathcal{B}))$ partition any natural edge of $T_{\mathcal{B}}(\phi^k(\mathcal{B}))$ into at most $2\mathcal{X}$ segments, where $\mathcal{X} = \text{rank}(F) - 1$.

Since $f_{\mathcal{B}} : T_{\mathcal{B}} \to T_{\mathcal{B}}$ is an $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}}$, the restrictions of $f_{\mathcal{B}}$ to $T_{\mathcal{B}}(\phi^k(\mathcal{B}))$ are $\phi^k(\mathcal{A})$-relative immersions $f_{\mathcal{B},k} : T_{\mathcal{B}}(\phi^k(\mathcal{B})) \to T_{\mathcal{B}}(\phi^k(\mathcal{B}))$ for $\phi|_{\phi^k(\mathcal{B})}$ that are conjugate to $f_{\mathcal{B}}$. As f_* agrees with $f_{\mathcal{B}}$ on $T_{\mathcal{B}}$, the restriction of f_* to $T_*(\phi^k(\mathcal{B}))$ is exactly $f_{\mathcal{B},k}$ for all $k \geq 1$. Similarly define $\phi^k(\mathcal{B})$-relative immersions $f_{\mathcal{C},k} : T_{\mathcal{C}}(\phi^k(\mathcal{C})) \to T_{\mathcal{C}}(\phi^k(\mathcal{C}))$ for $\phi|_{\phi^k(\mathcal{C})}$ that are conjugate to $f_{\mathcal{C}}$. As f_* induces $f_{\mathcal{C}}$ upon collapsing the lower stratum, any edges in $f_*(T_*(\phi^k(\mathcal{C})))$ but not $T_*(\phi^k(\mathcal{C}))$ must be in the lower stratum and the restriction of f_* to $T_*(\phi^k(\mathcal{C}))$ induces $f_{\mathcal{C},k}$ for all $k \geq 1$ upon collapsing the lower stratum. Applying the deformation retraction $f_*(T_*(\phi^k(\mathcal{C}))) \to T_*(\phi^k(\mathcal{C}))$ produces a homotopy restriction $f_{*,k}$ that still agrees with $f_{\mathcal{B},k}$ on $T_*(\phi^k(\mathcal{B}))$ and still induces $f_{\mathcal{C},k}$. By Lemma 5.2, we know that $f_{*,k}$ maps branch points of $T_*(\phi^k(\mathcal{C}))$ to $C(f_*)$-neighborhoods of branch points.

That the map $f_{*,k}$ induces the immersion $f_{\mathcal{C},k}$ upon collapsing the lower stratum means any branch point of $T_*(\phi^k(\mathcal{C}))$ that is not mapped to a branch point must be in the lower stratum. By the same token, $f_{*,k}$ agreeing with the immersion $f_{\mathcal{B},k}$ on $T_*(\phi^k(\mathcal{B}))$ means any branch point of $T_*(\phi^k(\mathcal{B}))$ that is not mapped to a branch point must be a bivalent point in $T_{\mathcal{B}}(\phi^k(\mathcal{B}))$. So we can apply a *bounded* homotopy to get a homotopy restriction $f_{*,k} : T_*(\phi^k(\mathcal{C})) \to T_*(\phi^k(\mathcal{C}))$ that maps branch points to branch points, maps any natural edge to a branch point or immersed path, still induces $f_{\mathcal{C},k}$ upon collapsing $T_*(\phi^k(\mathcal{B}))$, but differs from $f_{\mathcal{B},k}$ on $T_*(\phi^k(\mathcal{B}))$ by a homotopy supported in the natural edges of $T_{\mathcal{B}}(\phi^k(\mathcal{B}))$.

By the bound on the homotopy, $K(f_{*,k}) = K(f_*) + C(f_*)$ and $C(f_{*,k}) = 2C(f_*)$ can be taken as the new Lipschitz and cancellation constants of $f_{*,k}$.

Since $f_\mathcal{B} : T_\mathcal{B} \to T_\mathcal{B}$ is an expanding $\mathcal{A}$-relative immersion for $\phi|_\mathcal{B}$, the minimal subforest $T_\mathcal{B}(\phi^k(\mathcal{B}))$ is an $f_\mathcal{B}$-invariant subforest whose lengths of natural edges grows exponentially with k for all $k \geq 1$. So there is a $k_0 \gg 0$ such that all natural edges of $T_\mathcal{B}(\phi^k(\mathcal{B}))$ are longer than $2\mathcal{X} \cdot C \cdot K^{3\mathcal{X}-1}$ for all $k \geq k_0$. By the pigeonhole principle, each natural edge of $T_\mathcal{B}(\phi^k(\mathcal{B}))$ contains a natural edge of $T_*(\phi^k(\mathcal{B}))$ longer than $C \cdot K^{3\mathcal{X}}$ for all $k \geq k_0$. Let $\mathbb{G}_k$ be the directed graph of natural edges of $T_*(\phi^k(\mathcal{B}))$ where a directed edge $E_i \to E_j$ corresponds to $f_{*,k}$ mapping E_i over E_j. Set $\mathcal{S}_0$ to be those natural eges with length $\leq C$ and $\mathcal{S}$ to be those natural edges with directed path from $\mathcal{S}_0$ in $\mathbb{G}_k$; these will be the *short* natural edges. Since $f_{*,k}$ is K-Lipschitz and the shortest path between any natural edges in $\mathbb{G}_k$ has $3\mathcal{X}$ natural edges, all the short natural edges have length $\leq C \cdot K^{3\mathcal{X}-1}$. So the short natural edges $\mathcal{S}$ form an orbit-closed $f_{*,k}$-invariant subforest of $T_*(\phi^k(\mathcal{B})) \subset T_*(\phi^k(\mathcal{C}))$ with bounded components as $\mathcal{S}$ does not cover any natural edge of $T_\mathcal{B}(\phi^k(\mathcal{B}))$.

Collapsing all translates in $T_*(\phi^k(\mathcal{C}))$ of the short natural edges in $T_*(\phi^k(\mathcal{B}))$ induces a map $g'_{*,k} : Y'_{*,k} \to Y'_{*,k}$ with the same cancellation constant $C(g'_{*,k}) = C$. Now collapse the pretrivial edges until the induced map $g_{*,k} : Y_{*,k} \to Y_{*,k}$ has none. The new map $g_{*,k}$ still induce the immersion $f_{\mathcal{C},k}$ upon collapsing the rest of the lower stratum and, as a result, folding in $g_{*,k}$ may only occur between initial segments of natural edges of $Y_{*,k}$ whose $g_{*,k}$-images are lower stratum natural edges. However, all natural edges in the lower stratum of $Y_{*,k}$ are longer than C by construction and so no folding in $g_{*,k}$ is possible by bounded cancellation, i.e., $g_{*,k}$ is an immersion.

Collapsing a maximal invariant subforest with bounded components and forgetting bivalent vertices if necessary, we may assume $g_{*,k} : Y_{*,k} \to Y_{*,k}$ is a minimal $\phi^k(\mathcal{A})$-relative immersion for $\phi|_{\phi^k(\mathcal{C})}$. By the injectivity of ϕ, the restrictions $\phi|_{\phi^k(\mathcal{C})}$ and $\phi|_\mathcal{C}$ are conjugate and we can view $g_{*,k}$ as a minimal $\mathcal{A}$-relative immersion for $\phi|_\mathcal{C}$. It remains to show that every natural edge of $Y_{*,k}$ expands under $g_{*,k}$-iteration.

Suppose not, i.e., suppose there is a natural edge of $Y_{*,k}$ whose $g_{*,k}$-iterates have uniformly bounded length. Since $g_{*,k}$ is minimal, the non-expanding edges in the graph of groups $\mathcal{C}\backslash Y_{*,k}$ contain fixed subgraph $\mathcal{F}$ that supports a loxodromic element. The subgraph $\mathcal{F}$ is a free splitting of a ϕ-invariant free factor. Recall that the ϕ-elliptic free factor system $\mathcal{A}$ decomposes as a union of the maximal ϕ-fixed free factor system and free factors that eventually get mapped into this fixed system (Proposition 3.4 and Lemma 3.5). By construction, the point stabilizers of $Y_{*,k}$ are conjugates of $\mathcal{A}$. So any vertex of the fixed graph $\mathcal{F}$ is labelled by either the trivial group or a free factor of the maximal ϕ-fixed free factor system. Thus $\mathcal{F}$ is a free splitting of a ϕ-fixed free factor system that supports some loxodromic element. However, this contradicts Proposition 3.3 which states that all ϕ-periodic free factors are supported by the maximal ϕ-fixed free factor system and hence will be elliptic. Therefore, $g_{*,k}$ is an expanding $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{C}}$. □

We are now ready to inductively construct expanding relative immersions.

Theorem 6.5. *If $\phi : F \to F$ is injective and nonsurjective, then there is an expanding $\mathcal{A}$-relative immersion for ϕ, where $\mathcal{A}$ is the ϕ-elliptic free factor system.*

Proof. Suppose $\phi : F \to F$ is injective but not surjective. By Lemma 3.6, the ϕ-elliptic free factor system $\mathcal{A}$ is proper. The naive approach is to assume there exists a chain $\mathcal{A} = \mathcal{B}_0 \prec \cdots \prec \mathcal{B}_n = \{F\}$ in the poset of ϕ-invariant free factor system such that the restrictions $\phi|_{\mathcal{B}_{m+1}}$ are irreducible relative to $\mathcal{B}_m$ for all $m \geq 1$. This assumption is typical when working with automorphisms. For each restriction $\phi|_{\mathcal{B}_{m+1}}$, if the minimal stretch factor $\lambda(\phi, \mathcal{B}_{m+1}, \mathcal{B}_m) = 1$, then there is automatically a $\mathcal{B}_m$-relative immersion for $\phi|_{\mathcal{B}_{m+1}}$; and if $\lambda(\phi, \mathcal{B}_{m+1}, \mathcal{B}_m) > 1$, then there is an expanding $\mathcal{B}_m$-relative immersion for $\phi|_{\mathcal{B}_{m+1}}$ by Proposition 6.1. In either case, there is a $\mathcal{B}_m$-relative immersion for $\phi|_{\mathcal{B}_{m+1}}$. By Proposition 6.2, $\lambda(\phi, \mathcal{B}_1, \mathcal{B}_0) > 1$ and there is an expanding $\mathcal{B}_0$-relative immersion for $\phi|_{\mathcal{B}_1}$. By inductively patching these immersions together using Proposition 6.4, we get an expanding $\mathcal{B}_0$-relative immersion for ϕ and we are done. Unfortunately, since ϕ is nonsurjective, it could be that no chain $\mathcal{A} = \mathcal{B}_0 \prec \cdots \prec \mathcal{B}_n = \{F\}$ satisfies the naive assumption we made at the start. Recall that our definition of $\phi|_{\mathcal{B}_{m+1}}$ being irreducible relative to $\mathcal{B}_m$ presupposed $\phi^{-1} \cdot \mathcal{B}_m = \mathcal{B}_m$. Fortunately, this is a minor complication that can be easily addressed. The proof follows the approach described above closely but uses a chain with slightly weaker conditions on it.

We first construct a chain $\mathcal{A} = \mathcal{B}_0 \prec \cdots \prec \mathcal{B}_n = \{F\}$ in the poset of ϕ-invariant free factor system that we will induct on. Let $\mathcal{A} \prec \mathcal{B}_1$ be a chain of ϕ-invariant free factor systems such that $\phi|_{\mathcal{B}_1}$ is irreducible relative to $\mathcal{A}$. Suppose $\mathcal{B}_m$ has been constructed for some $m \geq 1$ and let $\mathcal{C} \succeq \mathcal{B}_m$ be the maximal free factor system in the chain $\mathcal{B}_m \preceq \phi^{-1} \cdot \mathcal{B}_m \preceq \phi^{-2} \cdot \mathcal{B}_m \preceq \cdots$ of ϕ-invariant free factor systems. If $\phi^{-1} \cdot \mathcal{B}_m = \mathcal{B}_m = \mathcal{C}$, then let $\mathcal{B}_m \prec \mathcal{B}_{m+1}$ be a chain of ϕ-invariant free factor systems such that $\phi|_{\mathcal{B}_{m+1}}$ is irreducible relative to $\mathcal{B}_m$. If $\mathcal{B}_m \prec \mathcal{C}$, then let $\mathcal{B}_m \prec \cdots \prec \mathcal{B}_{m+k} = \mathcal{C}$ be the chain of ϕ-invariant free factor systems such that $\mathcal{B}_{m+i} = \phi^{-i} \cdot \mathcal{B}_m$ for $1 \leq i \leq k$.

We proceed by inducting on the resulting chain between $\mathcal{A}$ and $\{F\}$. For the base case, $\phi|_{\mathcal{B}_1}$ is irreducible relative to $\mathcal{A}$; therefore, there is an expanding $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}_1}$ by Proposition 6.2. For our induction hypothesis, suppose that there is an expanding $\mathcal{A}$-relative immersion $f_{\mathcal{B}_m} : T_{\mathcal{B}_m} \to T_{\mathcal{B}_m}$ for $\phi|_{\mathcal{B}_m}$ for some $m \geq 1$. By our construction of the chain, either $\phi|_{\mathcal{B}_{m+1}}$ is irreducible relative to $\mathcal{B}$ or $\phi(\mathcal{B}_{m+1})$ is supported by $\mathcal{B}_m$. We deal with these two cases separately.

Case 1. Suppose $\phi|_{\mathcal{B}_{m+1}}$ is irreducible relative to $\mathcal{B}_m$. By Proposition 5.5, there is an irreducible $\mathcal{B}_m$-relative representative $f_* : T_* \to T_*$ for $\phi|_{\mathcal{B}_{m+1}}$ with minimal stretch factor. If $\lambda(f_*) = 1$, then $f_{\mathcal{B}_{m+1}} = f_*$ is a $\mathcal{B}_m$-relative isometric immersion for $\phi|_{\mathcal{B}_{m+1}}$. If $\lambda(f_*) > 1$, then there is an expanding $\mathcal{B}$-relative immersion $f_{\mathcal{B}_{m+1}}$ for $\phi|_{\mathcal{B}_{m+1}}$ by Proposition 6.1. In either case, we get a $\mathcal{B}_m$-relative immersion $f_{\mathcal{B}_{m+1}} : T_{\mathcal{B}_{m+1}} \to T_{\mathcal{B}_{m+1}}$ for $\phi|_{\mathcal{B}_{m+1}}$ defined on a $(\mathcal{B}_{m+1}, \mathcal{B}_m)$-forest $T_{\mathcal{B}_{m+1}}$. Thus, there is an expanding $\mathcal{A}$-relative immersion for $\phi|_{\mathcal{B}_{m+1}}$ by the induction hypothesis and Proposition 6.4.

Case 2. Now suppose $\phi(\mathcal{B}_{m+1})$ is supported by $\mathcal{B}_m$. Let $T_{\mathcal{B}_m}(\phi(\mathcal{B}_{m+1})) \subset T_{\mathcal{B}_m}$ be the minimal subforest of $\phi(\mathcal{B}_{m+1})$ and $T_{\mathcal{B}_{m+1}} = T_{\mathcal{B}_m}(\phi(\mathcal{B}_{m+1}))$ be the same forest after forgetting bivalent vertices. By injectivity of ϕ, we may consider $T_{\mathcal{B}_{m+1}}$ as a $(\mathcal{B}_{m+1}, \mathcal{A})$-forest that comes with a natural $\phi|_{\mathcal{B}_{m+1}}$-equivariant immersion $g : T_{\mathcal{B}_{m+1}} \to T_{\mathcal{B}_m}$. Since $f_{\mathcal{B}_m} : T_{\mathcal{B}_m} \to T_{\mathcal{B}_m}$ is an immersion, we can identify a subdivision of $T_{\mathcal{B}_m}$ with the minimal subforest $T_{\mathcal{B}_{m+1}}(\mathcal{B}_m) \subset T_{\mathcal{B}_{m+1}}$ of $\mathcal{B}_m$. Composing g with the subdivision and inclusion $T_{\mathcal{B}_{m+1}}(\mathcal{B}_m) \subset T_{\mathcal{B}_{m+1}}$ gives an $\mathcal{A}$-relative immersion $f_{\mathcal{B}_{m+1}} : T_{\mathcal{B}_{m+1}} \to T_{\mathcal{B}_{m+1}}$ for $\phi|_{\mathcal{B}_{m+1}}$, which is expanding since its image lies in $T_{\mathcal{B}_{m+1}}(\mathcal{B}_m)$ and its restriction to $T_{\mathcal{B}_{m+1}}(\mathcal{B}_m)$ is the expanding $\mathcal{A}$-relative immersion $f_{\mathcal{B}_m}$ after forgetting bivalent vertices. □

This gives us a complete characterization of when an injective endomorphism is induced by an expanding graph immersion.

Corollary 6.6. *Let $\phi : F \to F$ be an injective endomorphism. Then the following conditions are equivalent:*

1. *ϕ is induced by an expanding graph immersion;*
2. *no nontrivial conjugacy class in F has an infinite ϕ-tail;*
3. *ϕ has no nontrivial fixed free factor system.*

Proof. The implications (1) $\implies$ (2) $\implies$ (3) are obvious. Suppose ϕ has no nontrivial fixed free factor system. Then ϕ is nonsurjective and the ϕ-elliptic free factor system is trivial. By Theorem 6.5, there is an expanding ϕ-equivariant immersion $\tilde{f} : T \to T$ defined on an F-tree T. So ϕ is induced by an expanding immersion f of the graph $\Gamma = F\backslash T$. □

Remark. Unpacking the implication (3) $\implies$ (1) given by Theorem 6.5 in the corollary's proof, we use the fact (3) $\implies$ (2) (Theorem 3.1) and then show (3) + (2) $\implies$ (1).

7 — Interlude —

Let us now contextualize the previous section and, especially, Corollary 6.6. In our previous work [8], we determined exactly when the mapping torus of an expanding graph immersion has a word-hyperbolic fundamental group.

Theorem 7.1 ([8, Theorem 6.3])**.** *Let $\phi : F \to F$ be induced by an expanding graph immersion. $F*_\phi$ is word-hyperbolic if and only if it has no $BS(1,d)$ subgroups for $d \geq 2$.*

By Reynolds' result (Corollary 6.3), we knew this theorem applied to all irreducible and nonsurjective endomorphisms but it was unknown at the time what the general class of endomorphisms induced by expanding graph immersions was. It is clear that if ϕ has a periodic free factor, then it cannot be induced by an expanding graph immersion. Corollary 6.6 shows that the existence of periodic free factors is the only obstruction to being induced by an expanding graph immersion. So Theorem 7.1 can be restated as:

Theorem 7.2. *Let $\phi : F \to F$ be an injective endomorphism with no fixed free factor system. $F*_\phi$ is word-hyperbolic if and only if it has no $BS(1,d)$ subgroups for $d \geq 2$.*

The interesting thing about the restatement is that it is purely algebraic, i.e., there is no mention of topological graph immersions. Since fixed free factor systems correspond to *free-by-cyclic* subgroups $F' \rtimes \mathbb{Z}$ of $F*_\phi$, the restatement suggests that our proof of Theorem 7.1 can be adapted to work for all nonsurjective injective endomorphisms.

In what follows, we will use expanding relative immersions to relativize the proof of Theorem 7.1. For instance, the sequence of lemmas/propositions in the next section is essentially identical to the sequence in [8, Section 3]. However, this will *not* constitute an alternate proof of Brinkmann's theorem: $F \rtimes \mathbb{Z}$ is word-hyperbolic if and only if it has no $\mathbb{Z}^2$ subgroups [4]. We assume Brinkmann's theorem as the base case of our generalization.

8 Pullbacks

In our previous work [8, Sections 3], topological pullbacks for a graph immersion $f : \Gamma \to \Gamma$ were used to give sufficient conditions for $\pi_1(f) = \phi : F \to F$ to have an invariant nonfixed *cyclic subgroup system*. The goal of this section is to drop the immersion hypothesis and give sufficient conditions that apply to all injective endomorphisms of F.

Suppose immersions $f_1 : \Gamma_1 \to \Gamma$ and $f_2 : \Gamma_2 \to \Gamma$ induce inclusions of free groups $H_1, H_2 \leq F$ respectively. Then components of the core *(topological) pullback* of (f_1, f_2) are in one-to-one correspondence with nontrivial intersection $H_1 \cap gH_2g^{-1}$ as g ranges over (H_1, H_2)-double coset representatives of $H_1 \backslash F / H_2$. For a graph immersion $f : \Gamma \to \Gamma$ that induces an endomorphism $\phi : F \to F$, we get a one-to-one correspondence between components of the core pullback of (f^k, f^k) and nontrivial conjugacy classes $[\phi^k(F) \cap g\phi^k(F)g^{-1}]$ as $[[g]]$ ranges over $\phi^k(F)$-double cosets for all $k \geq 1$. We will not define topological pullbacks since we will be working algebraically in this section.

Given subgroups $H_1, H_2 \leq F$, we define the **(algebraic) pullback** of (H_1, H_2), denoted by $H_1 \wedge H_2$, to be the set of all nontrivial **components** $[H_1 \cap gH_2g^{-1}]$ as $[[g]]$ ranges over the (H_1, H_2)-double cosets in F. When H_1 and H_2 are finitely generated, then their pullback is a finite set.

For an endomorphism $\phi : F \to F$ and $k \geq 1$, define the algebraic pullbacks $\Lambda_k = \phi^k(F) \wedge \phi^k(F)$. There is a natural restriction $\phi|_{\Lambda_k} : \Lambda_k \to \Lambda_{k+1}$ of ϕ to Λ_k given by

$$[\phi^k(F) \cap g\phi^k(F)g^{-1}] \quad \mapsto \quad [\phi^{k+1}(F) \cap \phi(g)\phi^{k+1}(F)\phi(g)^{-1}].$$

Lemma 8.1. *If $\phi : F \to F$ is injective and $k \geq 1$, then $\phi|_{\Lambda_k} : \Lambda_k \to \Lambda_{k+1}$ as a function on the set of components is injective.*

Proof. Fix $k \geq 1$ and let g_1 and g_2 be $\phi^k(F)$-double coset representatives such that $\phi(g_1)$ and $\phi(g_2)$ were in the same ϕ^{k+1}-double coset. So there are $x, y \in \phi^{k+1}(F)$ such that

$x\phi(g_1)y = \phi(g_2)$. Equivalently, there are $x', y' \in \phi^k(F)$ such that $\phi(x')\phi(g_1)\phi(y') = \phi(g_2)$. Since ϕ is injective, we get $x'g_1y' = g_2$. So g_1 and g_2 are in the same $\phi^k(F)$-double coset. $\square$

For the rest of the section, assume $\phi : F \to F$ is an injective endomorphism. Then we get a chain of injections $\{F\} = \Lambda_0 \to \Lambda_1 \to \Lambda_2 \to \cdots$. Furthermore, injectivity of ϕ implies the restriction to components are isomorphisms. We will be mainly interested in $\hat{\Lambda}_{k+1} = \Lambda_{k+1} \setminus \phi(\Lambda_k)$, i.e., the components of Λ_{k+1} that are not part of $\phi(\Lambda_k)$. Equivalently, $\hat{\Lambda}_{k+1}$ has the following description:

$$\hat{\Lambda}_{k+1} = \{\, [\phi^{k+1}(F) \cap g\phi^{k+1}(F)g^{-1}] \in \Lambda_{k+1} \ : \ g \notin \phi(F)\,\}.$$

We might say *pullbacks stabilize* if $\hat{\Lambda}_k = \emptyset$ for some k. The image $\phi(F)$ is malnormal in F if and only if $\hat{\Lambda}_1 = \emptyset$. Pullback stability is a sort of generalization of *malnormality* for $\phi^k(F)$ with respect to ϕ.

Lemma 8.2. *Suppose $\phi : F \to F$ is injective and $k \geq 1$. If $\hat{\Lambda}_k$ is empty then so is $\hat{\Lambda}_{k+1}$ and if $\hat{\Lambda}_k$ has only cyclic components, then $\hat{\Lambda}_{k+1}$ is empty or has only cyclic components.*

Proof. There is an obvious "inclusion" of components, $\hat{\Lambda}_{k+1} \preceq \hat{\Lambda}_k$, induced by

$$\phi^{k+1}(F) \cap g\phi^{k+1}(F)g^{-1} \leq \phi^k(F) \cap g\phi^k(F)g^{-1}.$$

That $\hat{\Lambda}_k = \emptyset$ implies $\hat{\Lambda}_{k+1} = \emptyset$ is obvious. Suppose $\hat{\Lambda}_k$ has cyclic components, then $\hat{\Lambda}_{k+1}$ is empty or has cyclic components as the subgroups of a cyclic group are trivial or cyclic. $\square$

The **reduced rank** of a nontrivial finite rank free group H is $\mathrm{rr}(H) = \mathrm{rank}(H) - 1$ and the reduced rank of a pullback $H_1 \wedge H_2$, where H_1, H_2 are finitely generated free groups, is the sum of the reduced ranks of nontrivial intersections in $H_1 \wedge H_2$. The latter is denoted by $\mathrm{rr}(H_1 \wedge H_2)$. Since the restriction $\phi|_{\Lambda_k}$ gives natural isomorphisms of the components, the chain of injections produces a nondecreasing sequence of positive integers

$$\operatorname{rank}(F) - 1 = \operatorname{rr}(F) \leq \operatorname{rr}(\Lambda_1) \leq \operatorname{rr}(\Lambda_2) \leq \cdots .$$

Observe that $\operatorname{rr}(\Lambda_i) = \operatorname{rr}(\Lambda_{i+1})$ if and only if $\hat{\Lambda}_{i+1}$ is empty or has cyclic components only. By Lemma 8.2, the sequence becomes constant once two consecutive entries are equal. Walter Neumann used topological pullbacks to bound the reduced ranks of algebraic pullbacks thus improving Hanna Neumann's bound [10] on the rank of intersections.

Theorem 8.3 ([11, Proposition 2.1])**.** *If $H_1, H_2 \leq F$ are nontrivial finitely generated subgroups, then $\operatorname{rr}(H_1 \wedge H_2) \leq 2\operatorname{rr}(H_1)\operatorname{rr}(H_2)$.*

Remark. Although weaker than the *Strengthened Hanna Neumann Conjecture* (the Friedman-Mineyev Theorem [5, 7]), this bound is sufficient for our purposes.

Lemma 8.4. *If $\phi : F \to F$ is injective, then either $\hat{\Lambda}_k$ has cyclic components for all $k \geq 2\operatorname{rr}(F)^2$ or $\hat{\Lambda}_k = \emptyset$ for some k.*

Proof. Theorem 8.3 gives us a uniform bound on the reduced ranks of the pullbacks $\operatorname{rr}(\Lambda_k) \leq 2\operatorname{rr}(\phi^k(F))^2 = 2\operatorname{rr}(F)^2$ for all $k \geq 1$. By Lemma 8.2 and the uniform bound on the nondecreasing sequence of reduced ranks, $\operatorname{rr}(\Lambda_k)$ are all equal for $k \geq k_0 = 2\operatorname{rr}(F)^2$. Therefore, $\hat{\Lambda}_{k_0}$ is empty or has only cyclic components. The lemma follows by applying Lemma 8.2 again. □

We say $\phi : F \to F$ has an **invariant cyclic subgroup system with index $d \geq 1$** if there is an integer $k \geq 1$, element $x \in F$, and nontrivial cyclic subgroup $\langle c \rangle \leq F$ such that $\phi^k(\langle c \rangle) \leq x\langle c \rangle x^{-1}$ and has index d. We can now give the main result of this section:

Proposition 8.5. *If $\phi : F \to F$ is injective and $\hat{\Lambda}_k$ is nonempty for all $k \geq 1$, then ϕ has an invariant cyclic subgroup system with index $d \geq 2$.*

Proof. Let $k_0 = 2\,\mathrm{rr}(F)^2$ and $\phi : F \to F$ be an injective endomorphisma such that $\hat{\Lambda}_k$ is nonempty for all $k \geq 1$. By Lemma 8.4, $\hat{\Lambda}_k$ has cyclic components for $k \geq k_0$.

So far nothing in the section has used relative immersions but our main motivation for constructing them was this proposition. Note that $\hat{\Lambda}_1 \neq \emptyset$ automatically implies ϕ is not surjective. So ϕ is injective and nonsurjective. By Theorem 6.5, there is an expanding $\mathcal{A}^*$-relative immersion $f : T \to T$ for ϕ, where $\mathcal{A}^*$ is the ϕ-elliptic free factor system and T is an $(F, \mathcal{A}^*)$-tree. Elements of F will be called *elliptic* if they act with fixed point on T, i.e., their conjugacy classes are supported by $\mathcal{A}^*$. By construction of the ϕ-elliptic free factor system (Proposition 3.4), there is an $m \geq 1$ such that $[\phi^m(g)]$ is supported by the maximal ϕ-fixed free factor system $\mathcal{A}$ for all elliptic elements $g \in F$. By ϕ-invariance of $\mathcal{A}$, we get that $[\phi^k(g)]$ is supported by $\mathcal{A}$ for all $k \geq m$ and elliptic elements $g \in F$. As $\mathcal{A}$ is a ϕ-fixed free factor system, $\phi^k(\mathcal{A}) = \{\, \phi^k(A) \ : \ A \in \mathcal{A} \,\}$ are free factor systems of F equivalent to $\mathcal{A}$ in the poset of free factor systems for all k.

Suppose $[\phi^k(F) \cap g\phi^k(F)g^{-1}] \in \hat{\Lambda}_k$ for some $k \geq k_1 = k_0 + m$. As this component is cyclic, we may assume it has a representative $\phi^k(F) \cap g\phi^k(F)g^{-1}$ generated by a nontrivial element $x \in \phi^k(F)$. In particular, there exists $y \in \phi^k(F)$ and $g \in F \setminus \phi(F)$ such that $x = gyg^{-1}$. We first show that x is *loxodromic*, i.e., it is not elliptic. Suppose $x \in \phi^k(F)$ is elliptic. Then so is y and, furthermore, the two are in the free factor system $\phi^k(\mathcal{A})$ as $k \geq m$. Since $g \notin \phi(F)$, we have $g \notin \phi^k(A)$ for any $A \in \mathcal{A}$. But this contradicts the *malnormality* of free factor systems: elements in a free factor system, e.g. x, y in $\phi^k(\mathcal{A})$, cannot be conjugated by an element not in the free factor system, e.g. g not in $\phi^k(\mathcal{A})$. Therefore, x is loxodromic. The integer $k \geq k_1$ and component $\langle x \rangle \in \hat{\Lambda}_k$ were arbitrary, so all components of $\hat{\Lambda}_k$ are loxodromic for $k \geq k_1$.

Recall from the proof of Lemma 8.2 that each component of $\hat{\Lambda}_{k+1}$ is supported in a component of $\hat{\Lambda}_k$ for all $k \geq 1$. We denote this with an infinite descending chain:

$$\hat{\Lambda}_1 \succeq \hat{\Lambda}_2 \succeq \hat{\Lambda}_3 \succeq \cdots$$

Since $\hat{\Lambda}_{k_1}$ has finitely many components and the components are all cyclic, there is a cyclic component in $\hat{\Lambda}_{k_1}$ which supports some component of $\hat{\Lambda}_k$ for all $k \geq k_1$. Suppose this component has a representative generated by $c = \phi^{k_1}(x) \in \phi^{k_1}(F)$. Then for all $k \geq k_1$, there is a cyclic component of $\hat{\Lambda}_k$ with a representative generated by $\phi^k(x_k) \in \phi^k(F)$ such that $\langle c \rangle$ supports $\langle \phi^k(x_k) \rangle$.

If we let $\alpha \subset T$ be the axis for element c, then the previous sentence implies there are sequences of element $(x_k)_{k \geq k_1}$ and $(y_k)_{k \geq k_1}$ such that the (unoriented) axes of $\phi^k(x_k)$ are all translates $y_k \cdot \alpha$ of the (unoriented) axis α. For any $k \geq k_1$, replace x_k with its inverse if necessary so that the action of $\phi^k(x_k)$ on its axis is coherent (respects orientation) with the action of c on α. By passing to a strictly increasing subsequence $(k_i)_{i \geq 1}$, we may assume there is an edge e of T such that the axes α_{k_i} of $(x_{k_i})_{i \geq 1}$ all contain a translate of e. We now pass to the graph of groups $\Gamma = F \backslash T$ in order to avoid mentions of translates and orbits. The axis α will project to an immersed loop $\bar{\alpha}$ in Γ representing c and axes α_{k_i} project to immersed loops $\bar{\alpha}_{k_i}$ that represent x_{k_i} and whose $\bar{f}^{k_i}$-image is a power $\bar{\alpha}^{d_i}$ up to cyclic rotation/homotopy, where $d_i \geq 0$ and $\bar{f} : \Gamma \to \Gamma$ is the immersion induced by $f : T \to T$. The edge e projects to an edge $\bar{e}$ that is contained in all the loops $\bar{\alpha}_{k_i}$ for $i \geq 1$

The proof now mimics the proof of [8, Proposition 3.11]. Since f is an immersion, it maps axes in T onto axes and $\bar{f}$ maps immersed loops in Γ to immersed loops. So $\bar{f}^{k_i}(\bar{e})$ is a subpath of the immersed loop $\bar{f}^{k_i}(\bar{\alpha}_{k_i}) \simeq \bar{\alpha}^{d_i}$ for all i, and since f is expanding, $\bar{f}^{k_i}(\bar{e})$ contains arbitrarily long powers of $\bar{\alpha}$ as $i \to \infty$. Set N to be the number of subpaths of $\bar{\alpha}$ (up to rotation) that are also loops. Choose $n \gg 0$ such that $\bar{f}^{k_n}(\bar{e})$ contains the loop $\bar{\alpha}^{N+1}$ as a subpath. Then $\bar{f}^{k_{n+1}}(\bar{e})$ is a subpath of $\bar{\alpha}^{d_{n+1}}$ that contains the loop $\bar{f}^{k_{n+1}-k_n}(\bar{\alpha}^{N+1})$ as a subpath. In fact, for all positive integer $j \leq N+1$, the loop $\bar{f}^{k_{n+1}-k_n}(\bar{\alpha}^j)$ is a subpath of $\bar{\alpha}^{d_{n+1}}$. Thus, there is a sequence of loops $(\epsilon_j)_{j=1}^{N+1}$ that are subpaths of $\bar{\alpha}$ and strictly increasing positive integers $(s_j)_{j=1}^{N+1}$ such that $\bar{f}^{k_{n+1}-k_n}(\bar{\alpha}^j) \cdot \epsilon_j$ is $\bar{\alpha}^{s_j}$ up to rotation. By definition of N and pigeonhole principle, some $\epsilon_t = \epsilon_{t'} = \epsilon$ for some $t < t'$ and $\bar{f}^{k_{n+1}-k_n}(\bar{\alpha}^{t'-t})$ is $\bar{\alpha}^{s_{t'}-s_t}$ up to cyclic homotopy. Lifting back to the $(F, \mathcal{A})$-tree T, we find that $f^{k_{n+1}-k_n}$ maps the axis α to a translate of itself. So $\phi^{k_{n+1}-k_n}(c)$ is conjugate to a nontrivial power c^d and $d \geq 2$ since f is expanding. □

Remark. A careful examination of the proof reveals that it can be made more effective with the pigeonhole principle. So for any injective endomorphism ϕ, we can construct a specific number $k = k(\phi)$ for which $\hat{\Lambda}_k \neq \emptyset$ implies ϕ has an invariant cyclic subgroup system with index $d \geq 2$. Thus, one would not have to check infinitely many pullbacks to know that an invariant cyclic subgroup system with index $d \geq 2$ exists.

The **(algebraic) mapping torus** of an injective endomorphism ϕ is the ascending HNN extension of F with the presentation $F*_\phi = \langle F, t \mid t^{-1}xt = \phi(x), \forall x \in F \rangle$.

In the next section, we will actually be using the contrapositive of the proposition: if ϕ is injective and has no invariant cyclic subgroup system with index $d \geq 2$, then pullbacks stabilize. In this case, we get control on the types of annuli in the mapping torus $F*_\phi$, which allows us to prove the main theorem: $F*_\phi$ is word-hyperbolic if ϕ additionally has no fixed cyclic subgroup system.

Let us end the section with a third description of $\hat{\Lambda}_k$. The mapping torus $F*_\phi$ is the fundamental group of a circle of groups with one vertex group F and edge group F (See Appendix B). The edge monomorphisms for this circle of groups are the identity map $id_F : F \to F$ and endomorphism $\phi : F \to F$ and the corresponding Bass-Serre tree $\mathcal{T}$ has one orbit of edges and vertices. The tree also comes with a natural orientation where each vertex has exactly one outgoing oriented edge and the **stable letter** $t \in F*_\phi$ acts on its directed axis with positive translation. By construction, there is a unique vertex $\bullet$ of $\mathcal{T}$ whose stabilizer in $F*_\phi$ is exactly F. Let $\mathcal{T}^\bullet$ be the full subtree of $\mathcal{T}$ rooted at $\bullet$. The components of $\hat{\Lambda}_k$ are nontrivial stabilizers (up to conjugacy) of ($\phi^k(F)$-orbits of) geodesics in $\mathcal{T}^\bullet$ whose midpoint is $\bullet$. Colloquially, we can think of $\mathcal{T}^\bullet$ as a $\bullet$-descendants family-tree. In this view, the components of $\hat{\Lambda}_k$ are nontrivial conjugacy classes of simultaneous stabilizers for a pair of k^{th}-generation vertices in $\mathcal{T}^\bullet$ whose only common ancestor is $\bullet$.

9 Hyperbolic endomorphisms

We are finally ready to put all the pieces together. The first piece involves understanding the relationship between *annuli* in the mapping torus $F*_\phi$ and pullbacks of ϕ. The second piece involves building on Brinkmann's theorem to show that *atoroidal* injective endomorphisms are *hyperbolic*. In our previous work [8], we used these two pieces to give sufficient conditions for the mapping torus to be word-hyperbolic.

Theorem 9.1 ([8, Theorem 6.4])**.** *If $f : \Gamma \to \Gamma$ is a based-hyperbolic graph map and strictly bidirectional annuli in the mapping torus M_f are shorter than some positive even integer, then $\pi_1(M_f)$ is word-hyperbolic.*

We will define the new terms in the theorem as we go. Suppose $\phi : F \to F$ is an injective endomorphism and $f : \Gamma \to \Gamma$ is its topological representative. Recall that the (topological) mapping torus of f is the quotient space $M_f = (\Gamma \times [0,1]) / \sim_f$ with the identification $(x, 1) \sim_f (f(x), 0)$ for all $x \in \Gamma$ and the algebraic mapping torus $F*_\phi$ is isomorphic to the fundamental group $\pi_1(M_f)$. The **edge-space** of M_f will be the cross-section in M_f represented by $\Gamma \times \{\frac{1}{2}\}$.

Strictly bidirectional annuli of the mapping torus M_f with length $2L$ can be thought of as the pullbacks $\hat{\Lambda}_L$ of ϕ but it does take a bit of work to give the correspondence. Fix a basepoint $\star \in S^1$. For integers $L_1 < L_2$, an **(topological) annulus** in M_f of length $L = L_2 - L_1$ is a homotopy of loops $h : S^1 \times [L_1, L_2] \to M_f$ satisfying the following conditions:

1. it is transverse to the edge-space of M_f;
2. the h-preimage of the edge-space is $S^1 \times ([-L_1, L_2] \cap \mathbb{Z})$;
3. for integers $i \in [L_1, L_2]$, the **rings** of the annulus $h_i = h(\cdot, i) : S^1 \to M_f$ are locally injective every where except possibly at the basepoint $\star$;

4. and for the **trace** of the basepoint $h^\star = h(\star, \cdot) : [L_1, L_2] \to M_f$, no subpath between consecutive integer coordinates $[i, i+1]$ is homotopic rel. endpoints into the edge-space.

This is the definition used in [8]. In light of our last description of $\hat{\Lambda}_k$ in the previous section, we give alternative definition. Let $\mathcal{T}$ be the Bass-Serre tree for $F*_\phi$ and $\bullet$ be the point whose stabilizer is F. An **(algebraic) annulus** $([\alpha], p_\alpha)$ in $F*_\phi$ of length $L \geq 2$ is a choice a nontrivial conjugacy class $[\alpha]$ in $F*_\phi$ and an orbit of a geodesic path in $p_\alpha \subset \mathcal{T}$ of length L fixed by α. Since elements of $[\alpha]$ act on $\mathcal{T}$ with fixed points, we can always choose a representative $\alpha \in F$. Technically, we have defined a *conjugacy class* of algebraic annuli but the distinction will not be relevant for us.

Given a topological annulus h in M_f of length $L \geq 1$, then the generator of the image $\pi_1(h) : \mathbb{Z} \to \pi_1(M_f)$ determines a conjugacy class $[\alpha]$ in $\pi_1(M_f) \cong F*_\phi$. Condition (3) ensures $\pi_1(h)$ is injective and α is nontrivial. Let $\tilde{h} : \mathbb{R} \times [L_1, L_2] \to \tilde{M}_f$ be the lift of the annulus to the universal cover of M_f. Collapsing the $\tilde{\Gamma}$-direction of $\tilde{M}_f$ produces the Bass-Serre tree $\mathcal{T}$ and Condition (2) ensures the induced map $\bar{h} : \mathbb{R} \times ([L_1, L_2] \cap \mathbb{Z}) \to \mathcal{T}$ is constant on the first factor and its image is a collection of edge-midpoint; each ring h_i determines a conjugacy class in the stabilizer of the corresponding edge-midpoint. By Conditions (1) and (4), the midpoints extend to a geodesic edge-path p_α in $\mathcal{T}$ of length $L + 1$ fixed by α.

The other direction works in a similar fashion. For any conjugacy class $[\alpha]$ in $F*_\phi$ and two consecutive edge-midpoints in $\mathcal{T}^\bullet$ fixed by $\alpha \in F$, we can construct an annulus of length 1 as follows. Fix a basepoint in the edge-space and assume $\bullet$ is the vertex between the midpoints. If the midpoints are increasing/decreasing, then $\alpha = \phi(x) \in \phi(F)$ without loss of generality. Let σ, ρ be based loop in the edge-space representing $x, \phi(x) \in F$ and τ a based loop in M_f representing $t \in F*_\phi$. Then the based path $\sigma \cdot \tau \cdot \bar{\rho} \cdot \bar{\tau}$ is null-homotopic and can be extended to an annulus with ends σ, ρ and trace τ. If the midpoints are at the same height, then $\alpha = \phi(x) \in \phi(F)$ and $\alpha = g\phi(y)g^{-1} \in g\phi(F)g^{-1}$ for some $g \notin \phi(F)$. Let σ, ρ, γ be based loops in the edge-space representing $x, y, g \in F$ respectively and τ a based loop in M_f representing $t \in F*_\phi$. So the based path $\sigma \cdot \tau\gamma\bar{\tau} \cdot \bar{\rho} \cdot \tau\bar{\gamma}\bar{\tau}$ is null-homotopic and can be extended to an homotopy between σ, ρ satisfying Conditions (1)-(3) and having trace $\tau\gamma\bar{\tau}$. This trace satisfies Condition (4) because $tgt^{-1} \notin F$ and thus the homotopy is a topological annulus with ends σ, ρ. Given a geodesic path in $\mathcal{T}$ of length $L \geq 2$ fixed by α, we can replace the path with a translate in $\mathcal{T}^\bullet$ without affecting the class $[\alpha]$ and then construct a topological annulus in M_f of length $L - 1$ by concatenating the length 1 annuli from the preceding discussion. This concludes the correspondence between topological and algebraic annuli.

The natural orientation on $\mathcal{T}$ gives a dichotomy for algebraic annuli $([\alpha], p_\alpha)$ in $F*_\phi$:

1. all edges of p_α have the same orientation— we say α is **unidirectional**.

2. p_α switches from increasing to decreasing exactly once — we say α is **bidirectional**.

The reason each vertex of $\mathcal{T}$ has exactly one outgoing edge and hence the geodesic p_α cannot switch from decreasing to increasing follows from F having index 1 in F. For similar reasons, bidirectional annuli do not exist if and only if $\phi(F)$ is malnormal in F. The next proposition generalizes this equivalence of bidirectional annuli (or lack thereof) and malnormality.

An annulus $([\alpha], p_\alpha)$ in $F*_\phi$ is **strictly bidirectional** if the switch from increasing to decreasing occurs at the midpoint of p_α.

Lemma 9.2. *Let $\phi : F \to F$ be injective. For any integer $L \geq 1$, the mapping torus $F*_\phi$ has a strictly bidirectional annulus of length $2L$ if and only if $\hat{\Lambda}_L$ is nonempty.*

Proof. Algebraic annuli and the last description of $\hat{\Lambda}_L$ make this a trivial observation. If there is a strictly bidirection annulus $([\alpha], p_\alpha)$ in $F*_\phi$ of length $2L$, then we may assume the midpoint of p_α is $\bullet$ after replacing p_α with a translate. Then the stabilizer of p_α contains α and so it is nontrivial. Therefore, the conjugacy class of the stabilizer is a component of $\hat{\Lambda}_L$.

If $\hat{\Lambda}_L$ is not empty, then some path in $\mathcal{T}^\bullet$ of length $2L$ with midpoint at $\bullet$ has a nontrivial stabilizer. Choose a nontrivial element α in this stabilizer and $([\alpha], p_\alpha)$ is a strictly bidirectional annulus in $F*_\phi$ of length $2L$. □

Let $\phi : F \to F$ be an injective endomorphism with a topological representative $f : \Gamma \to \Gamma$. If ϕ has no invariant cyclic subgroup system with index $d \geq 2$, then there is an integer $L \geq 1$ for which $\hat{\Lambda}_L$ is empty (Proposition 8.5) and all strictly bidirectional annuli in M_f are shorter than $2L$ (Lemma 9.2). This sets up the second hypothesis of Theorem 9.1.

As for the first hypothesis, we begin by defining *based-hyperbolicity* and *hyperbolicity.* For a real number $\lambda > 1$ and integer $n \geq 1$, we say a graph map $f : \Gamma \to \Gamma$ is **(based-) (λ, n)-hyperbolic** if all (based) loops $\sigma : S^1 \to \Gamma$ (with the basepoint mapped to a vertex) satisfy the inequality

$$\lambda |f^n(\sigma)| \leq \max(\, |f^{2n}(\sigma)|, |\sigma|\,)$$

where $|\cdot|$ is the combinatorial length after tightening; whether tightening respects a basepoint (based homotopy) or not (free homotopy) will be apparent from the context.

When a graph map is (based-) (λ', n)-hyperbolic for some $\lambda' > 1, n \geq 1$, then it is (based-) (λ^k, nk)-hyperbolic for all $k \geq 1$ and $\lambda \in (1, \lambda']$. So the constants can be omitted and when we do need them, we can assume $\lambda > 1$ is any preferred integer. Hyperbolicity is a property of the homotopy class $[f]$ while based-hyperbolicity is a property of the map f. In the setting we will be interested in, the former will imply the latter.

A graph map $f : \Gamma \to \Gamma$ is **atoroidal** if there is no nontrivial loop σ in Γ and integer $k \geq 1$ such that $f^k(\sigma) \simeq \sigma$. This again is a property of the homotopy class $[f]$. Bestvina-Feighn-Handel showed, as a step in [2, Theorem 5.1], that hyperbolic atoroidal graph maps are based-hyperbolic. Their argument is reproduced here, modified for the possibility f is not a homotopy equivalence. This allows us to consider the growth rate of loops without basepoints for the rest of the section

Lemma 9.3. *If the graph map $f : \Gamma \to \Gamma$ is atoroidal and $(3, n)$-hyperbolic, then it is based-$(2, n')$-hyperbolic.*

Remark. To avoid context-ambiguity in the proof, we use $\|\cdot\|$ for lengths of free homotopy classes of loops and $|\cdot|$ for lengths of loops rel. basepoints. However, the distinction is not needed after the proof as all loops afterwards will be considered up to free homotopy.

Proof. Suppose $f : \Gamma \to \Gamma$ is atoroidal and $(3, n)$-hyperbolic for some integer $n \geq 1$. Set M to be the maximum length of $f^k(s)$ rel. basepoint over all embedded based loops s in Γ for $k \in \{0, n, 2n\}$.

Suppose $|f^n(\sigma)| \geq 4M$ for some immersed based loop σ and pick an embedded based loop s with same basepoint as σ so that the concatenation $s \cdot \sigma$ is an immersed loop, i.e., $\|s \cdot \sigma\| = |s| + |\sigma|$. As the graph map f is $(3, n)$-hyperbolic, we get

$$3\|f^n(s \cdot \sigma)\| \leq \max(\,\|f^{2n}(s \cdot \sigma)\|, \|s \cdot \sigma\|\,).$$

For a concatenation of based loops $\rho_1 \cdot \rho_2$, we get $|\,|\rho_1| - |\rho_2|\,| \leq \|\rho_1 \cdot \rho_2\| \leq |\rho_1| + |\rho_2|$.

Case 1. If $\|f^{2n}(s \cdot \sigma)\| \geq 3\|f^n(s \cdot \sigma)\|$, then

$$\begin{aligned} |f^{2n}(\sigma)| \geq \|f^{2n}(s \cdot \sigma)\| - |f^{2n}(s)| &\geq 3\|f^n(s \cdot \sigma)\| - M \\ &\geq 3|f^n(\sigma)| - 3|f^n(s)| - M \\ &\geq 3|f^n(\sigma)| - 4M \\ &\geq 2|f^n(\sigma)|. \end{aligned}$$

Case 2. If $\|s \cdot \sigma\| \geq 3\|f^n(s \cdot \sigma)\|$, then

$$|\sigma| = \|s \cdot \sigma\| - |s| \geq 3\|f^n(s \cdot \sigma)\| - M \geq 3|f^n(\sigma)| - 4M \geq 2|f^n(\sigma)|$$

Combining both cases: $2|f^n(\sigma)| \leq \max(\,|f^{2n}(\sigma)|, |\sigma|\,)$. If $|f^{nk}(\sigma)| \geq 4M$ for an immersed based loop σ in Γ and $k \geq 1$, then by induction

$$2^k \cdot |f^{nk}(\sigma)| \leq \max(\,|f^{2nk}(\sigma)|, |\sigma|\,).$$

For any bound B, there are only finitely many immersed based loops σ' in Γ with $|\sigma'| \leq B$. Since f is atoroidal, there is an integer $k \gg 0$ such that $|f^{nk}(\sigma')| \geq 8M$ for every based loop σ' with $|\sigma'| \leq 4M$ and we conclude that f is based-$(2, nk)$-hyperbolic . □

When p is a subpath of an immersed loop σ and $n \geq 1$, then $[f^n(p)]_\sigma$ is the subpath of $[f^n(p)]$ that survives in $[f^n(\sigma)]$ and $|f^n(p)|_\sigma$ is the length of $[f^n(p)]_\sigma$. Bounded cancellation implies $|f^n(p)| \leq |f^n(p)|_\sigma + 2C(f^n)$. The next lemma is based on Brinkmann's Lemma 4.2 in [4] with a few necessary changes made to account for ϕ possibly being nonsurjective.

Lemma 9.4. *Let $f : \Gamma \to \Gamma$ be a graph map and $R_* \subset \Gamma$ be an f-invariant union of roses such that the restriction $f|_{R_*} : R_* \to R_*$ is $(4, n)$-hyperbolic. For some constant L_c, if $p \subset R_*$ is a subpath (edge-path) of some immersed loop σ in Γ and $|f^n(p)|_\sigma \geq L_c$, then*

$$3|f^n(p)|_\sigma \leq \max\left(\,|f^{2n}(p)|_\sigma, |p|\,\right).$$

The number L_c is the **critical length** of the triple (f, Γ, R_*).

Proof. Let $f : \Gamma \to \Gamma$ be a graph map and $R_* \subset \Gamma$ be an f-invariant union of roses such that the restriction $f|_{R_*}$ is $(4, n)$-hyperbolic. Set M to be the maximum length of $f^k(s)$ rel. basepoint over all petals s in R_* and $k \in \{n, 2n\}$. Choose $L_c = 2C(f^{2n}) + 5M$ where $C(f^{2n})$ is the cancellation constant for f^{2n}. Recall the triangle inequality: $|\,|p_1| - |p_2|\,| \leq |p_1 \cdot p_2| \leq |p_1| + |p_2|$ for any path decomposition of a loop $p_1 \cdot p_2$. A remark on the context: paths $[p_i]$ are reduced rel. endpoints but the loop $[p_1 \cdot p_2]$ is freely reduced.

Given a subpath $p \subset R_*$ of some immersed loop σ in Γ, pick a petal s in R_* such that

$s \cdot p$ is an immersed loop in R_*. As the restriction to R_* is $(4, n)$-hyperbolic, we get

$$4|f^n(s \cdot p)| \leq \max(\,|f^{2n}(s \cdot p)|, |s \cdot p|\,).$$

If $4|f^n(s \cdot p)| \leq |f^{2n}(s \cdot p)|$, then

$$\begin{aligned} 4|f^n(p)|_\sigma \leq 4|f^n(p)| &\leq 4|f^n(s \cdot p)| + 4|f^n(s)| \\ &\leq |f^{2n}(s \cdot p)| + 4M \\ &\leq |f^{2n}(p)| + 5M \\ &\leq |f^{2n}(p)|_\sigma + 2C(f^{2n}) + 5M \quad \text{(bounded cancellation)} \end{aligned}$$

Similarly, if $4|f^n(s \cdot p)| \leq |s \cdot p|$, then

$$\begin{aligned} 4|f^n(p)|_\sigma &\leq |p| + 1 + 4M \\ &\leq |p| + 5M + 2C(f^{2n}) \quad \text{(since } M \geq 1\text{)} \end{aligned}$$

Since $L_c = 2C(f^{2n}) + 5M$, we have the desired implication:

$$|f^n(p)|_\sigma \geq L_c \implies 3|f^n(p)|_\sigma \leq \max\left(\,|f^{2n}(p)|_\sigma, |p|\,\right). \qquad \square$$

An injective endomorphism $\phi : F \to F$ is **atoroidal** if it has no fixed cyclic subgroup system, i.e., no invariant cyclic subgroup system with index $d = 1$. If $f : \Gamma \to \Gamma$ is a topological representative for ϕ, then f is atoroidal if and only if ϕ is atoroidal. The following proposition is an extension of Brinkmann's theorem [4, Proposition 7.1] and is the last technical result of the dissertation.

Proposition 9.5. *If $\phi : F \to F$ is injective and atoroidal, then it has a $(2, n)$-hyperbolic topological representative for some integer $n \geq 1$.*

Proof. Suppose $\phi : F \to F$ is an injective and atoroidal endomorphism. If ϕ is surjective, then the proposition is precisely Brinkmann's theorem. So we may assume ϕ is injective and nonsurjective. By Theorem 6.5, there is an expanding $\mathcal{A}^*$-relative immersion $g : T \to T$ for ϕ, where $\mathcal{A}^*$ is the ϕ-elliptic free factor system. Fix some $\mathcal{A}^*$-marked roses $(R_{\mathcal{A}^*}, \alpha_{\mathcal{A}^*})$ and set Γ to be the $(R_{\mathcal{A}^*}, \alpha_{\mathcal{A}^*})$-vertex blow-up of the graph of groups $F \backslash T$. The roses $R_{\mathcal{A}^*}$ form the *lower stratum* of Γ and the remaining edges the *top stratum*.

We outline the proof which follows closely the idea behind Brinkmann's proof of his theorem. Patch together a homotopy equivalence of the lower stratum with the expanding relative immersion to get some topological representative f of ϕ. By Brinkmann's theorem, the restriction of f to the lower stratum is hyperbolic. The expanding relative immersion on the top stratum means loops that are *mostly* top stratum will have uniform exponential growth under forward iteration. Lemma 9.4 implies loops that are *mostly* lower stratum will have uniform exponential growth under forward and/or backward iteration. The heart of the proof lies in specifying (quantifying) what being mostly top or lower stratum means and showing that all loops are one or the other. Of course, there are a few minor technicalities that need addressing; for instance, the restriction to the lower stratum is almost but not exactly a homotopy equivalence.

Recall that the maximal ϕ-fixed free factor system $\mathcal{A}$ is a subset of $\mathcal{A}^*$ and there is an integer $k_0 \geq 0$ such that $\phi^{k_0}(\mathcal{A}^*)$ is supported by $\mathcal{A}$ (Proposition 3.4 and Lemma 3.5). So we may find a topological representative $f_{\mathcal{A}^*} : R_{\mathcal{A}^*} \to R_{\mathcal{A}^*}$ for $\phi|_{\mathcal{A}^*}$ whose restriction to the periodic roses $R_{\mathcal{A}}$, denoted by $f_{\mathcal{A}}$, is a homotopy equivalence. As ϕ is atoroidal, the restriction $f_{\mathcal{A}}$ is $(4, n_0)$-hyperbolic for some $n_0 \geq 1$ (Brinkmann's theorem).

If σ is an immersed loop in $R_{\mathcal{A}^*}$, then $[f^k(\sigma)]$ is a loop in the periodic roses $R_{\mathcal{A}}$ for all $k \geq k_0$. Since the restriction $f_{\mathcal{A}}$ is $(4, n_0)$-hyperbolic and $f_{\mathcal{A}^*}^{n_0 k}(\sigma)$ is a loop in $R_{\mathcal{A}}$ for any loop σ in $R_{\mathcal{A}^*}$ and $k \geq k_0$, we get the inequality
$4 \cdot |f_{\mathcal{A}}^{n_0}(f_{\mathcal{A}^*}^{n_0 k}(\sigma))| \leq \max(\,|f_{\mathcal{A}}^{2n_0}(f_{\mathcal{A}^*}^{n_0 k}(\sigma))|, |f_{\mathcal{A}^*}^{n_0 k}(\sigma)|\,)$ for all loops σ in $R_{\mathcal{A}^*}$ and $k \geq k_0$.
Choose an integer $k_1 \geq 1$ so that $4^{k_1} \geq 4K^{n_0 k_0}$, where $K = K(f_{\mathcal{A}^*})$ is the Lipschitz constant for $f_{\mathcal{A}^*}$. Suppose σ is a loop in $R_{\mathcal{A}^*}$.

If $4 \cdot |f_{\mathcal{A}}^{n_0}(f_{\mathcal{A}^*}^{n_0(k_0+k_1-1)}(\sigma))| \leq |f_{\mathcal{A}}^{2n_0}(f_{\mathcal{A}^*}^{n_0(k_0+k_1-1)}(\sigma))|$, then by induction $4^{k_0+k_1} \cdot |f_{\mathcal{A}^*}^{n_0(k_0+k_1)}(\sigma)| \leq |f_{\mathcal{A}^*}^{2n_0(k_0+k_1)}(\sigma)|$.

If $4 \cdot |f_{\mathcal{A}}^{n_0}(f_{\mathcal{A}^*}^{n_0(k_0+k_1-1)}(\sigma))| \leq |f_{\mathcal{A}^*}^{n_0(k_0+k_1-1)}(\sigma)|$, then by induction and Lipschitz property $4^{k_1} \cdot |f_{\mathcal{A}^*}^{n_0(k_0+k_1)}(\sigma)| \leq |f_{\mathcal{A}^*}^{n_0 k_0}(\sigma)| \leq K^{n_0 k_0} \cdot |\sigma|$ and $4 \cdot |f_{\mathcal{A}^*}^{n_0(k_0+k_1)}(\sigma)| \leq |\sigma|$ by choice of k_1.

Therefore, the lower stratum map $f_{\mathcal{A}^*}$ is $(4, n_1)$-hyperbolic with $n_1 = n_0(k_0 + k_1)$.

Let $f : \Gamma \to \Gamma$ be a topological representative for ϕ that extends $f_{\mathcal{A}^*}$ to the top stratum and induces the expanding $\mathcal{A}^*$-relative immersion $g : T \to T$ upon collapsing the lower stratum in the universal cover $\tilde{\Gamma}$. For an arbitrary immersed loop σ in Γ, define σ_{top} (σ_{low} resp.) to be the collection of maximal subpaths of σ in the top (lower resp.) stratum. For all $n \geq 1$, define $[f^n(\sigma_{top})]_\sigma$ ($[f^n(\sigma_{low})]_\sigma$ resp.) to be the collection of paths $[f^n(p)]_\sigma$ where p is some path in σ_{top} (σ_{low} resp.). That f induces an immersion g upon collapsing the lower stratum implies that the top stratum is *persistent*: if σ is an immersed loop in Γ, then $f(\sigma)_{top}$ survives in $[f(\sigma)]$.

As the relative immersion $g : T \to T$ is expanding, there is an integer $k_2 \geq 1$, such that $g^{k_2}(e)$ has length ≥ 2 for all edges e in T; and as f induces g, for any immersed loop σ in Γ and path p in σ_{top}, we get $2|p|_\sigma \leq |f^{k_2}(p)|_\sigma$. We may replace n_1 and k_2 with a common multiple and assume $n_1 = k_2$. A similar inequality holds in the lower stratum. By the $(4, n_1)$-hyperbolicity of $f|_{R_{\mathcal{A}^*}}$ and Lemma 9.4, there is a critical length $L_c = L_c(f, \Gamma, R_{\mathcal{A}^*})$ such that for any immersed loop σ in Γ and path p in σ_{low},

$$|f^{n_1}(p)|_\sigma \geq L_c \implies 3|f^{n_1}(p)|_\sigma \leq \max(\,|f^{2n_1}(p)|_\sigma, |p|\,).$$

Set M to be the maximal length amongst all paths in $f^{n_1}(e)_{low}$ for all top stratum edges e of Γ. For any integer $k \geq 1$, we distinguish two cases:

Case 1. If $|f^{n_1 k}(\sigma)_{low}| < (L_c + 6M)\,|f^{n_1 k}(\sigma)_{top}|$, then

$|f^{n_1k}(\sigma)| = |f^{n_1k}(\sigma)_{low}| + |f^{n_1k}(\sigma)_{top}| < (L_c + 6M + 1)\,|f^{n_1k}(\sigma)_{top}|$ and

$$2^k\,|f^{n_1k}(\sigma)_{top}| \leq |f^{n_1k}(f^{n_1k}(\sigma)_{top})|_\sigma \leq |f^{2n_1k}(\sigma)|.$$

Additionally, if $2^k \geq 2(L_c + 6M + 1)$, then $2\,|f^{n_1k}(\sigma)| < |f^{2n_1k}(\sigma)|$.

Case 2. Suppose $|f^{n_1k}(\sigma)_{low}| \geq (L_c + 6M)\,|f^{n_1k}(\sigma)_{top}|$. Set m to be the number of paths in $f^{n_1k}(\sigma)_{low}$. Then m is also the number of paths in $f^{n_1k}(\sigma)_{top}$ and $m \leq |f^{n_1k}(\sigma)_{top}|$. By the pigeonhole principle, some path ρ in $f^{n_1k}(\sigma)_{low}$ satisfies $|\rho| \geq L_c + 6M$. As $|\rho| \geq 6M$, we have $3(|\rho| - 2M) \geq 2|\rho|$. Set $\sigma' = f^{n_1(k-1)}(\sigma)$. By definition of M and persistence of $f^{n_1}(\sigma')_{top}$, there must be a path p' in σ'_{low} such that $[f^{n_1}(p')]_{\sigma'}$ is a subpath of ρ, $|f^{n_1}(p')|_{\sigma'} \geq |\rho| - 2M \geq L_c$, and $3\,|f^{n_1}(p')|_{\sigma'} \leq \max(\,|f^{2n_1}(p')|_{\sigma'}, |p'|\,)$.

If $|f^{2n_1}(p')|_{\sigma'} \geq 3|f^{n_1}(p')|_{\sigma'}$, then $|f^{2n_1}(p')|_{\sigma'} \geq 2|\rho|$ and $|f^{n_1(k+1)}(p')|_{\sigma'} \geq 3^{k-1} \cdot 2|\rho|$.

If $|p'| \geq 3|f^{n_1}(p')|_{\sigma'}$, then $|p'| \geq 2|\rho|$. By inducting on the same argument used at the start of the case, there must be a path p in σ_{low} such that $|f^{n_1k}(p)|_\sigma$ is a subpath of ρ and $|p| \geq 2^k\,|\rho|$. In either case, we get $2^k|\rho| \leq \max(\,|f^{n_1(k+1)}(p')|_{\sigma'}, |p|\,)$. Define $f^{n_1k}(\sigma)_{crit}$ to be the set of paths ρ in $f^{n_1k}(\sigma)_{low}$ with $|\rho| \geq L_c + 6M$. Altogether, we have shown:

$$2^k|f^{n_1k}(\sigma)_{crit}| \leq \max(\,|f^{2n_1k}(\sigma)|, |\sigma|\,).$$

The following computation is lifted from Brinkmann [4, Proof of Proposition 7.1]. Set $A = |f^{n_1k}(\sigma)_{crit}|$ to be the total length of paths in $f^{n_1k}(\sigma)_{crit}$, $B = |f^{n_1k}(\sigma)_{low}| - A$ to be the total length of the remaining paths in $f^{n_1k}(\sigma)_{low}$, and $C = |f^{n_1k}(\sigma)_{top}|$. We now find a positive lower bound of $\frac{A}{A+B+C}$ that is independent of σ and k. We assumed $A + B \geq (L_c + 6M)\,C$ and so $\frac{A}{A+B+C} \geq \frac{A(L_c+6M)}{(A+B)(L_c+6M+1)}$ and we can focus on the factor $\frac{A}{A+B} = 1 - \frac{B}{A+B}$. Recall $m \leq C$, so $A + B \geq (L_c + 6M)\,m$. Since each path p in $f^{n_1k}(\sigma)_{low}$ but not in $f^{n_1k}(\sigma)_{crit}$ satisfies $|p| < L_c + 6M$ and there are at most m of them, $B \leq m\,(L_c + 6M - 1)$. Combining the last two inequalities gives the bound

$1-\frac{B}{A+B} \geq 1-\frac{m\,(L_c+6M-1)}{m\,(L_c+6M)} \geq \frac{1}{L_c+6M}$. Altogether, $\frac{A}{A+B+C} \geq \frac{1}{L_c+6M+1}$.

Additionally, if $2^k \geq 2(L_c+6M+1)$, then $2|f^{n_1 k}(\sigma)| \leq \max(\,|f^{2n_1 k}(\sigma)|, |\sigma|\,)$.

Choose $k \geq 1$ so that $2^k \geq 2(L_c+6M+1)$; the two exhaustive cases above imply f is $(2, n_1 k)$-hyperbolic. □

All the heavy lifting is done and we have proved our main theorem

Theorem 9.6. *Let $\phi : F \to F$ be an injective endomorphism. Then the following statements are equivalent:*

1. *$F{*_\phi}$ is word-hyperbolic;*
2. *$F{*_\phi}$ contains no $BS(1,d)$ subgroups with $d \geq 1$;*
3. *ϕ has no invariant cyclic subgroup system with index $d \geq 1$;*
4. *ϕ has a based-hyperbolic topological representative and all strictly bidirectional annuli in its mapping torus are shorter than some positive even integer.*

Proof. This proof is just a matter of bookkeeping.

(1) $\Longrightarrow$ (2): $BS(1,d)$ subgroups are well-known obstructions to word-hyperbolicity.

(2) $\Longrightarrow$ (3): if ϕ has an invariant cyclic subgroup system with index d, then there is a subgroup of $F{*_\phi}$ isomorphic to a quotient of $BS(1,d)$; use normal forms to show the subgroup is in fact isomorphic to $BS(1,d)$ (See [6, Lemma 2.3] for details).

(3) $\Longrightarrow$ (4): Proposition 8.5 with Lemma 9.2 give the second part of (4) and Proposition 9.5 with Lemma 9.3 give the first part.

(4) $\Longrightarrow$ (1) — Theorem 9.1: the proof is a calculation showing M_f satisfies the *annuli flaring condition* and then invoking Bestvina-Feighn's combination theorem [1]. □

A Marshall Hall's theorem applied

Stallings' paper [14] is the standard reference for the material in this appendix.

A group $\mathcal{G}$ is **subgroup separable/LERF** if, for every finitely generated subgroup $H \leq \mathcal{G}$ and element $g \in \mathcal{G} \setminus H$, there is a finite index subgroup $\hat{H} \leq G$ that contains H but not g. By passing to the kernel of the action of $\mathcal{G}$ on cosets of $\hat{H}$, it follows that this definition is equivalent to the definition given in the proof of Lemma 2.4.

Marshall Hall's theorem states that free groups are subgroup separable. Since every finitely generated subgroup of a free group F is contained in a finite rank free factor of F, it is enough to prove the theorem for finite rank free groups. The proof we give here is due to Stallings.

Theorem A.1 ([14, Theorem 6.1])**.** *Finite rank free groups are subgroup separable.*

Sketch proof. Let F be a finite rank free group, (R, α) be an F-marked rose, $H \leq F$ be a finitely generated free group, and $g \in F \setminus H$. The Stallings subgroup graph for $S(H)$ with respect to (R, α) is a core graph uniquely defined for the conjugacy class $[H]$. We can extend it to a possibly non-core finite graph $S(H, \star)$ that is uniquely defined for the subgroup H. This *pointed* Stallings graph comes with an immersion $\iota : S(H, \star) \to R$ that maps the marked point $\star$ to the vertex of R such that the image of $\pi_1(\iota)$ equals $\alpha(H)$ in $\pi_1(R, \iota(\star))$. The graph $S(H, \star)$ and immersion ι can be extended to a graph $S(H, g)$ and immersion ι_g such that the equality between the image of $\pi_1(\iota_g)$ and $\alpha(H)$ still holds and there is an edge-path p_g in $S(H, g)$ with distinct endpoints and starts at $\star$ for which $\iota_g(p_g)$ is a loop in (R, α) representing $\alpha(g)$.

As H is finitely generated, $S(H, g)$ is a finite graph and the immersion ι_g *can be extended to a finite cover* $\hat{\imath} : \hat{R} \to R$. Let $\hat{H}$ be the corresponding finite index subgroup. As $S(H, g) \subset \hat{R}$, we have $H \leq \hat{H}$ and, since the path p_g has distinct endpoints in $\hat{R}$, it follows that $g \notin \hat{H}$. □

To turn this sketch proof into a complete proof, the portion in the proof that is italicized ("ι_g can be extended to a finite cover") needs to be proven. The following proposition is a stronger version of Lemma 2.4.

Proposition A.2. *Let $\mathcal{G}$ be a finitely generated subgroup separable group and $\psi : \mathcal{G} \to \mathcal{G}$ be an automorphism. If $H \leq \mathcal{G}$ is finitely generated and $\psi(H) \leq H$, then $\psi(H) = H$.*

Proof. Let $\mathcal{G}$ be a finitely generated supgroup separable group, $\psi : F \to F$ an automorphism, and $H \leq \mathcal{G}$ a finitely generated subgroup such that $\psi(H) \leq H$. Suppose, for a contradiction, $g \in H \setminus \psi(H)$ and invoke subgroup separability for the pair $(\psi(H), g)$. There exists a finite index subgroup $\hat{H}_0$ that contains $\psi(H)$ but not g. Let $\hat{H}_1$ be the intersection of all finite index subgroups of $\mathcal{G}$ with the same index as $\hat{H}_0$. As $\mathcal{G}$ is finitely generated, there are finitely many such finite index subgroups and they are permuted by any automorphism of $\mathcal{G}$. Therefore, $\hat{H}_1$ is a finite index *characteristic* subgroup and $\psi(\hat{H}_1) = \hat{H}_1$. By construction, $g \notin \hat{H}_1$ and no nontrivial element $h \in \psi(H) \leq \hat{H}_0$ is in the $\hat{H}_1$-coset $g\hat{H}_1$ — if $g\hat{h} = h \in \psi(H)$ for some $\hat{h} \in \hat{H}_1 \leq \hat{H}_0$, then $g = h\hat{h}^{-1} \in \hat{H}_0$, which is a contradiction.

Let $\pi : \mathcal{G} \to G$ be the quotient map with kernel $\hat{H}_1$, then ψ induces an automorphism $\bar{\psi} : G \to G$ as $\psi(\hat{H}_1) = \hat{H}_1$. The conditions on g and nontrivial elements of $\psi(H)$ imply $\pi(g) \in \pi(H)$ but $\pi(g) \notin \bar{\psi}(\pi(H)) = \pi(\psi(H))$. But this is a contradiction since $\bar{\psi}$ is an automorphism of a finite group G, so $\psi(H) \leq H$ implies $\bar{\psi}(\pi(H)) = \pi(H)$. $\square$

B Bass-Serre theory reviewed

Serre's book [13] is the standard reference for the material in this appendix.

Let G be a group and T be a simplicial tree. Simplicial/isometric G-actions on T will be assumed to be minimal, i.e., the action doesn't induce an action on any proper subtree, and *orientation preserving*, i.e., G acts without edge-inversion. These are not restrictive assumptions since any nonminimal G-action with edge-inversion can be subdivided and restricted to get a minimal and orientation preserving G-action on a another tree. As this is enough for the results in the first part of this dissertation, let us assume no nontrivial element of G fixes an edge of any G-tree T.

Let T be a G-tree with a finite/compact quotient graph $\mathcal{Q} = G\backslash T$. Choose a maximal tree $\mathcal{T}$ in $\mathcal{Q}$ and lift it to a subtree in $\tilde{\mathcal{T}}$. We are effectively choosing a representative for each G-orbit of vertices in T. Let $V = V(\tilde{\mathcal{T}})$ be the vertices of $\tilde{\mathcal{T}}$. For each $v \in V$, let $G_v \leq G$ be the subgroup of elements that fix v. Lastly, let $E = \mathcal{T}^c$ be the set of edges of $\mathcal{Q}$ not in the maximal tree $\mathcal{T}$. The *fundamental theorem of Bass-Serre theory* uses this information to construct a presentation: $G \cong \langle G_v, e \mid v \in V, e \in E \rangle$.

This presentation is known as a *graph of groups decomposition* of G. More carefully, a **graph of groups** $\mathcal{G}$ is a graph whose vertices are labelled by groups. The vertex-labeling groups are aptly called **vertex groups** and elements of G conjugate into vertex groups are known as **elliptic elements**. Conversely, elements of G that are not conjugate into the vertex groups, i.e., the elements that act freely, are the **loxodromic elements**. The quotient graph $\mathcal{Q}$ considered with vertex groups $\{G_v \mid v \in V\}$ is an example of a graph of groups. Part of the fundamental theorem of Bass-Serre theory states that all graph of groups can be constructed this way and their related trees are called **Bass-Serre trees**. Graph of groups have a related notion of fundamental groups and, in this language, the previous isomorphism is a canonical identification between the "deck group" for the Bass-Serre $T \to \mathcal{Q}$ and the "fundamental group" $\pi_1(\mathcal{Q})$.

When $G = F$ is free, the vertex groups form a free factor system of F and the graph of groups decomposition given by the finite quotient of the F-tree is also known as a **free splitting** of F.

When the action of G on T does have a nontrivial element that fixes an edge, then the presentation for a fundamental group of a graph of groups is a bit more involved and better understood inductively. Assume the quotient graph $\mathcal{Q}$ has a single edge t let $G_t \geq G$ be the subgroup of elements that fix (stabilizer for) some orbit representative $\tilde{t}$ of the lift of t to T. There are two cases:

Case 1. $\mathcal{Q}$ has two vertices x, y. Then let $G_x, G_y \leq G$ be the stabilizers for the corresponding endpoints of $\tilde{t}$. Note that G_t naturally include into G_x, G_y and we will denote the inclusions by i_x, i_y. By the fundamental theorem of Bass-Serre theory, G is an amalgamated free product of G_x and G_y over the subgroup G_t:

$$G \cong \pi_1(\mathcal{Q}) = \langle G_x, G_y \mid i_x(z) = i_y(z), \forall z \in G_t \rangle$$

Case 2. $\mathcal{Q}$ has one vertex x. Then let $G_x, G_y \leq G$ be the stabilizer for the origin and terminal endpoint of $\tilde{t}$ respectively and denote as before the inclusions of G_t to G_x, G_y by i_x, i_y. Since there only one orbit of vertices, G_x and G_y are conjugate and the conjugation acts on i_y to give a map $i'_y : G_t \to G_x$. By the fundamental theorem, G is the HNN extension of G_x over isomorphic subgroups $i_x(G_t)$ and $i'_y(G_t)$:

$$G \cong \pi_1(\mathcal{Q}) = \langle G_x, t \mid i_x(z) = t i'_y(z) t^{-1}, \forall z \in G_t \rangle$$

In the latter part of this dissertation, we start studying the ascending HNN extension $G = F*_\phi$, which can be thought of as an example of this second case: set $G_x = G_t = F$, $i_x = id_F : F \to F$, and $i'_y = \phi : F \to F$. So the group $F*_\phi$ naturally acts on the Bass-Serre tree T that has one orbit of edges and vertices.

C Bestvina-Handel's algorithm relativized

Bestvina-Handel's paper [3] is the reference for the material in this appendix.

The objective in this appendix is to sketch the proof of the fact that irreducible relative representatives with minimal stretch factor are train tracks. Bestvina-Handel's construction of train tracks for irreducible automorphisms [3] translates verbatim to the non-free forest setting.

Let $\phi : F \to F$ be an injective endomorphism, $\mathcal{A} \prec \mathcal{B}$ be ϕ-invariant free factor systems, and T_* be a $(\mathcal{B}, \mathcal{A})$-forest. We allow forests to have bivalent vertices. An **$\mathcal{A}$-relative (topological) representative** for the restriction $\phi|_{\mathcal{B}}$ is a $\phi|_{\mathcal{B}}$-equivariant graph map $f_* : T_* \to T_*$ with no pretrivial edges. Additionally, we say the relative representative is **minimal** if it has no orbit-closed invariant subforests with bounded components.

For any $\mathcal{A}$-relative representative f_*, we get the **transition matrix** $A(f_*)$. An $\mathcal{A}$-relative representative f_* is **irreducible** if the matrix $A(f_*)$ is irreducible and, in this case, the **stretch factor** of f_*, denoted by $\lambda(f_*)$, is the Perron-Frobenius eigenvalue of $A(f_*)$. An **$\mathcal{A}$-relative train track** for $\phi|_{\mathcal{B}}$ is an $\mathcal{A}$-relative representative f_* for $\phi|_{\mathcal{B}}$ that additionally satisfies the property: the edge-paths $f_*^n(e)$ are immersed for all edges e in T_* and integers $n \geq 1$. We have set the stage for the theorem.

Theorem C.1 ([3, Theorem 1.7])**.** *If $\phi|_{\mathcal{B}}$ is irreducible relative to $\mathcal{A}$, T_* is a $(\mathcal{B}, \mathcal{A})$-forest with no bivalent vertices, and $f_* : T_* \to T_*$ is an irreducible $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$ with minimal stretch factor, then f_* is an irreducible $\mathcal{A}$-relative train track.*

We now state the lemmas that will be fundamental steps in the proof. The argument relies on understanding how various moves on an irreducible $\mathcal{A}$-relative representative f_* affect $\lambda(f_*)$ and invoking minimality of $\lambda(f_*)$ to conclude that no such moves are possible. Note that although the moves are described locally, they must be performed equivariantly if we want the resultant forests to be $(\mathcal{B}, \mathcal{A})$-forests. The proofs of these moves/lemmas are omitted since they are the same as the proofs in [3].

The first move is *subdivision*, which occurs at an interior point of an edge that is in the preimage of vertices under the representative.

Lemma C.2 ([3, Lemma 1.10]). *If $f_* : T_* \to T_*$ is an irreducible $\mathcal{A}$-relative representative for $\phi|_\mathcal{B}$ and $f'_* : T'_* \to T'_*$ is induced by a subdivision, then f'_* is an irreducible $\mathcal{A}$-relative representative and $\lambda(f'_*) = \lambda(f_*)$.*

The next move is *bivalent homotopy*, which occurs at a bivalent vertex and decreases the number of edges.

Lemma C.3 ([3, Lemma 1.13]). *If $f_* : T_* \to T_*$ is an irreducible $\mathcal{A}$-relative representative for $\phi|_\mathcal{B}$ and $f''_* : T''_* \to T''_*$ is an irreducible $\mathcal{A}$-relative representative induced by a bivalent homotopy followed by collapse of a maximal invariant subforest with bounded components, then $\lambda(f''_*) \leq \lambda(f_*)$.*

The last move we need is *folding*, which occurs between a pair of oriented edges originating from the same vertex that have the same image under the representative.

Lemma C.4 ([3, Lemma 1.15]). *Suppose $f_* : T_* \to T_*$ is an irreducible $\mathcal{A}$-relative representative for $\phi|_\mathcal{B}$ and $f'_* : T'_* \to T'_*$ is induced by a fold. If f'_* is an $\mathcal{A}$-relative representative, then it is irreducible and $\lambda(f'_*) = \lambda(f_*)$. Otherwise, if $f''_* : T''_* \to T''_*$ is an irreducible $\mathcal{A}$-relative representative induced by a homotopy of f'_* that makes the final map locally injective on the interior of edges, followed by collapse of a maximal invariant subforest with bounded components, then $\lambda(f''_*) < \lambda(f_*)$.*

Sketch proof of Theorem C.1. Let $\phi|_\mathcal{B}$ be irreducible relative to $\mathcal{A}$, T_* be a $(\mathcal{B}, \mathcal{A})$-forest with no bivalent vertices, and $f_* : T_* \to T_*$ be an irreducible $\mathcal{A}$-relative representative with minimal stretch factor. If $\lambda(f_*) = 1$, then f_* is a simplicial embedding and we are done. So we may assume $\lambda(f_*) > 1$.

Suppose, for a contradiction, that f_* is not an $\mathcal{A}$-relative train track, then the edge-path $f_*^n(e)$ is not immersed for some edge e in T_* and integer $n \geq 1$. Let n be the smallest such integer and assume $\star$ is an interior point of an edge e at which f_*^n fails to be locally injective. We appropriately subdivide T_* so that a neighborhood U of $\star$ and its iterates $f_*^k(U)$ $(1 \leq k \leq n)$ satisfy nice properties: 1) U is an interval whose boundary consists of distinct vertices; 2) f_*^k is locally injective on U for $1 \leq k < n$; 3) f_*^n folds U at $\star$ to an edge; and 4) $\star \notin f_*^k(U)$ for $1 \leq k \leq n$. We can then iteratively fold $f_*^{n-1}(U)$, ..., $f_*^2(U)$, and $f_*(U)$. By minimality of n, all the folds except the last one induce an irreducible $\mathcal{A}$-relative representative. By the first case of Lemma C.4, this irreducible $\mathcal{A}$-relative representative has the same stretch factor as f_*. By construction, the last fold induces a map f_*' that fails to be an $\mathcal{A}$-relative representative as it fails to be locally injective at $\star$. We can apply a tightening homotopy on f_*' to make it locally injective at $\star$, then collapse a maximal invariant subforest with bounded components to get $f_*'' : T_*'' \to T_*''$, a minimal $\mathcal{A}$-relative representative for $\phi|_{\mathcal{B}}$. By Lemma 5.4, the map f_*'' is irreducible. By the second case of Lemma C.4, the stretch factor is strictly smaller: $\lambda(f_*'') < \lambda(f_*)$. We then sequentially apply bivalent homotopies and collapse maximal invariant subforests with bounded components until we get an irreducible $\mathcal{A}$-relative representative f_*''' such that T_*''' has no bivalent vertices. The stretch factor satisfies $\lambda(f_*''') \leq \lambda(f_*'') < \lambda(f_*)$ by Lemma C.3. However, this contradicts the minimality of $\lambda(f_*)$. So f_* must have been an expanding irreducible $\mathcal{A}$-relative train track. □

References

[1] Mladen Bestvina and Mark Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.

[2] Mladen Bestvina, Mark Feighn, and Michael Handel. Laminations, trees, and irreducible automorphisms of free groups. *Geom. Funct. Anal.*, 7(2):215–244, 1997.

[3] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math. (2)*, 135(1):1–51, 1992.

[4] Peter Brinkmann. Hyperbolic automorphisms of free groups. *Geom. Funct. Anal.*, 10(5):1071–1089, 2000.

[5] Joel Friedman. Sheaves on graphs, their homological invariants, and a proof of the Hanna Neumann conjecture: with an appendix by Warren Dicks. *Mem. Amer. Math. Soc.*, 233(1100):xii+106, 2015.

[6] Ilya Kapovich. Mapping tori of endomorphisms of free groups. *Comm. Algebra*, 28(6):2895–2917, 2000.

[7] Igor Mineyev. Submultiplicativity and the Hanna Neumann conjecture. *Ann. of Math. (2)*, 175(1):393–414, 2012.

[8] Jean Pierre Mutanguha. Hyperbolic immersions of free groups. *Groups Geom. Dyn.*, to appear.

[9] Jean Pierre Mutanguha. Irreducible nonsurjective endomorphisms of F_n are hyperbolic. *Bull. Lond. Math. Soc.*, to appear.

[10] Hanna Neumann. On the intersection of finitely generated free groups. Addendum. *Publ. Math. Debrecen*, 5:128, 1957.

[11] Walter D. Neumann. On intersections of finitely generated subgroups of free groups. In *Groups—Canberra 1989*, volume 1456 of *Lecture Notes in Math.*, pages 161–170. Springer, Berlin, 1990.

[12] Patrick R. Reynolds. *Dynamics of irreducible endomorphisms of F_N*. PhD thesis, University of Illinois at Urbana-Champaign, 2011.

[13] Jean-Pierre Serre. *Arbres, amalgames,* SL_2. Société Mathématique de France, Paris, 1977. Avec un sommaire anglais, Rédigé avec la collaboration de Hyman Bass, Astérisque, No. 46.

[14] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[15] William P. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982.