



RESEARCH ARTICLE

Limit pretrees for free group automorphisms: existence

Jean Pierre Mutanguha 

Max Planck Institute for Mathematics, Vivatsgasse 7, Bonn 53111, Germany; E-mail: mutanguha@mpim-bonn.mpg.de.

Received: 25 July 2023; Revised: 3 March 2024; Accepted: 28 March 2024

2020 Mathematics Subject Classification: Primary – 20F65; Secondary – 20E05, 20E36

Abstract

To any free group automorphism, we associate a real pretree with several nice properties. First, it has a rigid/non-nesting action of the free group with trivial arc stabilizers. Secondly, there is an expanding pretree-automorphism of the real pretree that represents the free group automorphism. Finally and crucially, the loxodromic elements are exactly those whose (conjugacy class) length grows exponentially under iteration of the automorphism; thus, the action on the real pretree is able to detect the growth type of an element.

This construction extends the theory of metric trees that has been used to study free group automorphisms. The new idea is that one can equivariantly blow up an isometric action on a real tree with respect to other real trees and get a rigid action on a treelike structure known as a real pretree. Topology plays no role in this construction as all the work is done in the language of pretrees (intervals).

Contents

Introduction	1
1 Preliminaries	3
1.1 Free splittings and topological representatives	3
1.2 Free group systems and automorphisms	4
1.3 Train track theory	5
2 Trees and real pretrees	6
2.1 Trees and index theory	6
2.2 Pretrees and rigid actions	7
Interlude – same trees, different views	8
3 Limit pretrees: exponentially growing automorphisms	10
4 Blow-ups: simple, naïve, and ideal	14
4.1 Simple patchwork	14
4.2 Naïve stitching	16
4.3 Ideal stitching	19
Epilogue – are all limits the same?	21
A Index theory for small rigid actions	22
B Solutions to some exercises	27

Introduction

The study of free group outer automorphisms shares a lot with the theory of mapping class groups of surfaces. The relevant dictionary replaces compact surfaces with finite graphs (i.e., finite 1-dimensional CW-complexes) and surface homeomorphisms with graph homotopy equivalences. Although many

© The Author(s), 2024. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

results in surface theory have analogues in the free group setting, the situation for mapping classes tends to be comparatively well-behaved as one is working with surface homeomorphisms. On the other hand, an infinite order outer automorphism of a free group can never be represented by a homeomorphism of a graph. As a consequence, we have fewer tools at our disposal to study free group automorphisms. One such important missing tool is a complete analogue of Nielsen–Thurston theory.

Nielsen–Thurston theory

The main goal of the current project is defining an analogue of singular measured foliations in the free group setting. Let us start with a summary of the Nielsen–Thurston theory for surface homeomorphisms.

Suppose S is a compact surface with negative Euler characteristic, and let $f: S \rightarrow S$ be a homeomorphism. William Thurston proved [21, Theorem 4] there is

1. a *canonical* (potentially empty) union γ of essential simple closed curves;
2. a homeomorphism $g: S \rightarrow S$ isotopic to f that leaves a regular neighbourhood $N(\gamma)$ of the multicurve γ invariant; and
3. the restriction of g to permuted components of the complement $S \setminus N(\gamma)$ is either
 - a *pseudo-Anosov* homeomorphism (i.e., has a *canonical* invariant measured singular foliation whose transverse measure is scaled by a stretch factor $\lambda > 1$); or
 - a *linear* homeomorphism (i.e., reducible (possibly non-canonically) to a finite order homeomorphism) – Birman–Lubotzky–McCarthy later proved the latter reduction can be made canonical [4, Theorem C], but that is not in the standard expositions of Thurston’s result.

Thus, up to isotopy, we may assume a given surface homeomorphism $S \rightarrow S$ has a canonical decomposition into pseudo-Anosov components and their complements, and each orbit S' of pseudo-Anosov components has an invariant measured singular foliation whose transverse measure is scaled by some stretch factor $\lambda_{S'} > 1$. Lifting this decomposition to the universal cover \tilde{S} gives a canonical partition of the cover into lifts of the foliation’s leaves and the unfoliated complementary regions.

One may now endow this universal cover with a pseudo-metric using the lift of the transverse measure of the foliation. Morgan–Shalen proved [18, Section 1] that the corresponding metric space is an \mathbb{R} -tree, or simply *tree* as we shall call it in this paper, and the action of $\pi_1(S)$ on the universal cover by deck transformations induces an action on this tree by isometries, or simply *isometric* action. Finally, any lift of the surface homeomorphism induces an expanding ‘dilation’ of the tree: the tree decomposes into finitely many orbits of subtrees, and the restriction of the induced map to each subtree is an expanding homothety – the expansion factor may vary with the subtree’s orbit.

The minimal isometric $\pi_1(S)$ -action on the tree has two interesting properties:

- arc (pointwise) stabilizers are trivial; and
- an element of $\pi_1(S)$ is *elliptic* if and only if its conjugacy class is represented by a closed curve in the complement of the pseudo-Anosov components.

Moreover, the tree’s equivariant dilation class is canonical by construction.

Our motivation and main theorem

When we started this project, we wanted to prove a direct free group analogue of this statement about the canonical limit tree – the dilation requirement on the expanding homeomorphism would have been relaxed and complements of pseudo-Anosov components replaced by *polynomially growing* subgroups. This proved to be extremely elusive, and we now doubt such a metric analogue always exists!

Fortunately, we did manage to prove a variation to this analogy. Let us return to the limit \mathbb{R} -tree T given by the canonical decomposition of the universal cover \tilde{S} and the transverse measure. Closed arcs in T determine a *real pretree* structure; a real pretree is essentially what you get when you try to define an \mathbb{R} -tree without a metric or topology (see Section 2.2). The isometric $\pi_1(S)$ -action on the \mathbb{R} -tree

determines a *rigid* $\pi_1(S)$ -action on its real pretree. Rigid actions are commonly known as *non-nesting* actions in the literature [12], but we find the name ‘rigid’ more evocative.

There is no simple way to capture the ‘expanding’ nature of the homotheties without a metric on the real pretree. We introduce the term *(free group)-expanding* to describe the induced pretree-automorphisms (see Section 2.2). With this new terminology, the free group analogue for the real pretree statement becomes the following:

Main Theorem (Theorem 3.3). *If $\phi: F \rightarrow F$ is an automorphism of a finitely generated free group, then there is:*

1. *a minimal rigid F -action on a real pretree T with trivial arc stabilizers;*
2. *a ϕ -equivariant F -expanding pretree-automorphism $f: T \rightarrow T$; and*
3. *an element in F is T -elliptic if and only if it grows polynomially under $[\phi]$ -iteration.*

The nontrivial point stabilizers are finitely generated. Moreover, these subgroups are proper and have rank strictly less than that of F if and only if $[\phi]$ is exponentially growing.

Remark. The third condition can be restated in terms of loxodromic elements. The actual theorem will be stated and proven in a ‘relative setting’. The growth types of elements and automorphisms are defined in Chapter 3.

We suspect that the action’s rigidity can be strengthened: perhaps, the induced action on the *pretree completion* is a *convergence action*. To complete the analogy with the surface setting, we discuss in the epilogue the extent to which a limit pretree produced by this theorem is canonical – this is proven in the sequel [19]. This is of much interest as it is a canonical representation of a free group automorphism. For instance, it allows us to classify automorphisms in terms of the limit pretree’s *index* (see Appendix A).

In the long run, we would also like to find interesting canonical representations of polynomially growing automorphisms – any pretree of Theorem 3.3 is a point for these automorphisms.

Proof outline for the main theorem. Using irreducible train tracks, construct an isometric action on a limit tree with trivial arc stabilizers. This action admits an expanding homothety that represents the given free group automorphism. In particular, polynomially growing elements are elliptic; however, it is possible for some exponentially growing elements to be elliptic as well. By considering restrictions of the automorphism to point stabilizers, we get a hierarchy of isometric actions on trees. Most importantly, elements grow polynomially if and only if they are elliptic in each tree in the hierarchy.

The key step describes how to combine this hierarchy of trees, through blow-ups, into one *treelike* structure – a real pretree. While intuitive, the details of this construction get a bit technical. If done appropriately, this new structure will admit an F -expanding pretree-automorphism. Since we have a hierarchy of expanding homotheties, the contraction mapping theorem implies the blow-up construction can indeed be done appropriately!

1. Preliminaries

In this paper, F will denote a nontrivial free group of finite rank. Note that subscripts will never indicate the rank but will instead be mostly used to index a collection of free groups.

1.1. Free splittings and topological representatives

A *free splitting* of F is a simplicial tree T (i.e., 1-dimensional contractible CW-complex) and a *minimal* (left) F -action by simplicial automorphisms with trivial edge stabilizers. We will also assume the simplicial tree T has no bivalent vertices. Choose a maximal set of orbit representatives \mathcal{V} for the set of vertices of T whose *half-edge neighbourhood* is connected. If the simplicial tree T is not a singleton, then the collection of nontrivial stabilizers \mathcal{G} of vertices in \mathcal{V} is a *proper free factor system* of F by Bass–Serre theory [20]. After labelling the vertices of the finite quotient graph $\Gamma = F \setminus T$ with \mathcal{G} , we

get a graph of groups decomposition (Γ, \mathcal{G}) of F with trivial edge groups. In fact, by the fundamental theorem of Bass–Serre theory, there is a one-to-one correspondence between free splittings of F and graph of groups decompositions of F with trivial edge groups. For free splittings, we may use (Γ, \mathcal{G}) as a synonym for T .

A (relative) topological representative of an automorphism $\phi: F \rightarrow F$ on a free splitting T of F is a map $f: T \rightarrow T$ satisfying the following conditions:

- *cellular*: f maps vertices to vertices and is injective on edges; and
- *ϕ -equivariant*: $f(x \cdot p) = \phi(x) \cdot f(p)$ for all $x \in F$ and $p \in T$.

By definition, a topological representative on T induces a cellular map $[f]: \Gamma \rightarrow \Gamma$ on the quotient graph Γ .

Suppose $f: T \rightarrow T$ is a topological representative of an automorphism $\phi: F \rightarrow F$, $v \in \mathcal{V}$ a vertex orbit representative, and $G_v \in \mathcal{G}$ the stabilizer of v . Then $f(v) = x_v \cdot w$ for some element $x_v \in F$ and vertex orbit representative $w \in \mathcal{V}$. The stabilizer of $f(v)$ is $x_v G_w x_v^{-1}$ and, if G_v is not trivial, then the restriction of ϕ to G_v is an isomorphism $G_v \rightarrow x_v G_w x_v^{-1}$. By post-composing with the inner automorphism $\text{inn}(x_v^{-1}): F \rightarrow F$ that maps $y \mapsto x_v^{-1} y x_v$, we get a homomorphism $\phi_v: G_v \rightarrow G_w$. The *outer class* $[\phi_v]$ is independent of the chosen element $x_v \in F$ with $f(v) = x_v \cdot w$ (i.e., the homomorphism ϕ_v is unique up to post-composition with an inner automorphism of G_w). The collection $\{\phi_v : v \in \mathcal{V}\}$ is denoted by $\phi|_{\mathcal{G}}$ and called a restriction of ϕ to \mathcal{G} .

For the rest of the paper, we shall restrict \mathcal{V} to the subset consisting of vertices with nontrivial stabilizers. This is mostly a stylistic choice made to simplify the exposition. For instance, the restriction $\phi|_{\mathcal{G}}$ will permute the nontrivial stabilizers in \mathcal{G} and, under this assumption, can be considered an automorphism of \mathcal{G} .

1.2. Free group systems and automorphisms

To formalize the last statement, we define a (countable) group system \mathcal{H} to be a disjoint union $\bigsqcup_{i \in \mathcal{I}} H_i$ of countably many nontrivial countable groups; the latter are the *components* of \mathcal{H} . The empty system is the group system with empty index set \mathcal{I} . If all component groups H_i have [property-?], then we shall call \mathcal{H} a ‘[property-?] group system’. For instance, we will mainly work with *subgroup* systems and *free group* systems. In some ambient group, a subgroup system \mathcal{G} carries another subgroup system \mathcal{G}' if each \mathcal{G}' -component is contained in a conjugate of some \mathcal{G} -component.

The complexity of a (possibly trivial) free group H is $c(H) = 2 \cdot \text{rank}(H) - 1$, which takes values in $\mathbb{Z}_{\geq -1} \cup \{\infty\}$. For a (possibly empty) free group system \mathcal{H} ,

$$c(\mathcal{H}) := \sum_{i \in \mathcal{I}} c(H_i) \quad \in \mathbb{Z}_{\geq 0} \cup \{\infty\}.$$

A group system has finite type if the index set is finite and components are finitely generated. In particular, a free group system has finite type exactly when its complexity is finite. In this paper, \mathcal{F} will denote a nonempty free group system of finite type.

The collection of nontrivial vertex stabilizers \mathcal{G} for a free splitting of F can and will be viewed as a subgroup (or rather, free factor) system of finite type (even if empty). Similarly, we define a free splitting T of F to be a ‘disjoint union’ of free splittings of the components of F . A free splitting T is degenerate if all component simplicial trees T_i are singletons. By passing to the (representatives of) nontrivial vertex stabilizers in a nondegenerate free splitting, we can inductively form a descending chain of free factor systems with strictly decreasing complexity. Starting with F , the length of such a chain is at most $c(F) + 1$. We will mostly state and prove the results in terms of free group systems to facilitate induction on complexity.

An automorphism $\psi: \mathcal{H} \rightarrow \mathcal{H}$ of a group system \mathcal{H} is a disjoint union of isomorphisms $\psi_i: H_i \rightarrow H_{\sigma \cdot i}$ ($i \in \mathcal{I}$) where $\sigma \in \text{Sym}(\mathcal{I})$. A subgroup system $\mathcal{G} = \bigsqcup_{j \in \mathcal{J}} G_j$ of \mathcal{H} is $[\psi]$ -invariant if $\psi(G_j) =$

$h_j G_{\alpha \cdot j} h_j^{-1}$ for all $j \in \mathcal{J}$, for some choice $(h_j \in \mathcal{H} : j \in \mathcal{J})$ and $\alpha \in \text{Sym}(\mathcal{J})$. The collection of isomorphisms $\psi|_{\mathcal{G}} = \{\text{inn}(h_j^{-1}) \circ \psi|_{G_j} : G_j \rightarrow G_{\alpha \cdot j}\}$ will be called a restriction of ψ to \mathcal{G} ; note that $\psi|_{\mathcal{G}}$ involves an implicit choice $(h_j : j \in \mathcal{J})$.

Let $\psi : \mathcal{F} \rightarrow \mathcal{F}$ be an automorphism and \mathcal{T} a free splitting of \mathcal{F} . A topological representative of ψ on \mathcal{T} is a disjoint union of ψ_i -equivariant cellular maps $f_i : T_i \rightarrow T_{\sigma \cdot i}$.

1.3. Train track theory

We conclude the chapter with some preliminary results from Bestvina–Handel’s theory of train tracks [3, Section 1]. A relative train track for an automorphism $\psi : \mathcal{F} \rightarrow \mathcal{F}$ is a topological representative of ψ whose iterates are all topological representatives as well; equivalently, a train track is a topological representative whose restrictions of iterates to any edge are all injective. A topological representative $f : \mathcal{T} \rightarrow \mathcal{T}$ is irreducible if for any pair of edges $[e], [e']$ in the quotient graph $\Gamma_* = \mathcal{F} \setminus \mathcal{T}$, there is an integer $n \geq 1$ for which $[f^n(e')]$ contains $[e]$. To any topological representative f , we can associate a nonnegative integer square matrix $A(f)$: rows and columns are indexed by the (unoriented) edges $[e], [e']$ resp. and matrix entry is the number of times $[e]$ appears in the path $[f(e')]$; then f is irreducible if and only if $A(f)$ is irreducible.

Let $f : \mathcal{T} \rightarrow \mathcal{T}$ be an irreducible topological representative. By Perron–Frobenius theory, the matrix $A(f)$ has a unique eigenvalue $\lambda(f) \geq 1$ with a positive eigenvector $v(f)$; it follows from the theory that f is a simplicial automorphism if $\lambda(f) = 1$. Using $v(f)$, we can equip \mathcal{T} with an invariant metric (i.e., \mathcal{F} acts by isometries with respect to this metric); furthermore, after applying an equivariant isotopy, the restriction of f to any edge will be a $\lambda(f)$ -homothety. The metric on \mathcal{T} will be referred to as the eigenmetric. We have set the stage for the foundational theorem due to Bestvina–Handel:

Theorem 1.1 (cf. [3, Theorem 1.7]). *If $\psi : \mathcal{F} \rightarrow \mathcal{F}$ is an automorphism of a free group system \mathcal{F} and \mathcal{G} a $[\psi]$ -invariant proper free factor system of \mathcal{F} , then there is an irreducible train track $\tau : \mathcal{T} \rightarrow \mathcal{T}$ for ψ defined on some nondegenerate free splitting \mathcal{T} of \mathcal{F} whose vertex stabilizers carry \mathcal{G} .*

See also Rylee Lyman’s [16, Theorem A] for a very general version of this theorem.

Proof outline. Section 1 of [3] describes an algorithm that takes a topological representative of $\phi : F \rightarrow F$ on a free splitting of F with trivial vertex stabilizers as input and finds either an irreducible train track $\tau : T \rightarrow T$ on a free splitting with trivial vertex stabilizers or a topological representative on a nondegenerate free splitting with nontrivial vertex stabilizers. However, trivial vertex stabilizers are not crucial to the algorithm, and it can be adapted into one that takes a topological representative on a nondegenerate free splitting with vertex stabilizers \mathcal{G} as input and finds either an irreducible train track on a free splitting with the same stabilizers or a topological representative on a nondegenerate free splitting whose vertex stabilizers *properly carry* \mathcal{G} .

Starting with an initial input of a topological representative on a nondegenerate free splitting with vertex stabilizers \mathcal{G} , the modified algorithm can then be repeatedly applied to its own outputs until it finds an irreducible train track on some nondegenerate free splitting whose vertex stabilizers carry \mathcal{G} . This process will stop after at most $(c(F) - c(\mathcal{G}))$ repetitions. Finally, we note that F being ‘connected’ was not crucial to the algorithm and F can be safely replaced with a system \mathcal{F} . \square

Let $f : \mathcal{T} \rightarrow \mathcal{T}$ be a topological representative. We say an immersed path σ in \mathcal{T} is f -legal if restrictions of f -iterates to σ are all injective. Thus, a train track is a topological representative τ whose edges are τ -legal. Besides knowing edges (and their forward iterates) are legal, it is useful to have legal elements (i.e., *loxodromic elements with legal axes*).

Proposition 1.2. *If $\psi : \mathcal{F} \rightarrow \mathcal{F}$ is an automorphism and $\tau : \mathcal{T} \rightarrow \mathcal{T}$ an irreducible train track for ψ , then each nondegenerate component of \mathcal{T} contains a τ -legal axis of some loxodromic element of \mathcal{F} .*

Proof. If $\lambda(\tau) = 1$, then τ is a simplicial automorphism and all immersed paths in \mathcal{T} are τ -legal. So we may assume $\lambda(\tau) > 1$. Choose an oriented edge e_i in a component $T_i \subset \mathcal{T}$. Irreducibility of train track

τ and $\lambda(\tau) > 1$ imply we can find distinct translates $x_1 \cdot e_i, x_2 \cdot e_i$ in the oriented edge-path $\tau^n(e_i)$ for some large n . The \mathcal{T} -loxodromic element $x_2 x_1^{-1}$ has a τ -legal axis in \mathcal{T} : a fundamental domain for the axis is the τ -legal subpath of $\tau^n(e_i)$ joining the midpoints of the chosen two translates of e_i . \square

2. Trees and real pretrees

Expanding irreducible train tracks are very useful for understanding the dynamics of an automorphism since the iterates of edges all expand by the same factor under the eigenmetric. Unfortunately, since the train track necessarily fails to be injective near some vertices, not all paths will similarly expand: for instance, there are so-called ‘periodic Nielsen paths’ whose lengths (after reduction) remain uniformly bounded under iteration.

One way to get around this is to promote the expanding irreducible train track to an expanding homothety; however, this promotion requires leaving the category of simplicial trees and working in the metric category instead (Proposition 3.2).

2.1. Trees and index theory

A (metric) tree is a 0-hyperbolic geodesic metric (nonempty) space – 0-hyperbolic means the union of any two sides of a geodesic triangle contains the third side of the triangle; this is also known as an \mathbb{R} -*tree*. More generally, a forest \mathcal{T} is a disjoint union $\bigsqcup_{i \in \mathcal{I}} T_i$ of trees. Let T be a tree; a direction at a point $p \in T$ is a component of the complement $T \setminus \{p\}$. A branch point of T is a point with at least three directions.

A λ -homothety $f: T \rightarrow T$ is expanding if $\lambda > 1$. Being expanding is invariant under iteration and composition with an isometry. It follows from the contraction mapping theorem that an expanding λ -homothety of a *complete* tree has a unique fixed point and it is a *repellor*. Generally, a λ -homothety $f: \mathcal{T} \rightarrow \mathcal{T}$ of a forest $\mathcal{T} = \bigsqcup_{i \in \mathcal{I}} T_i$ is a disjoint union of λ -homotheties $f_i: T_i \rightarrow T_{\sigma \cdot i}$ ($i \in \mathcal{I}$), where $\sigma \in \text{Sym}(\mathcal{I})$.

An isometry of a tree is either elliptic if it fixes a point or loxodromic if it has a unique minimal invariant subtree isometric to \mathbb{R} , known as the *axis*, on which it acts by a nontrivial translation. We will mostly only care about *isometric actions* on trees with *finite arc stabilizers* (i.e., actions by isometries where the stabilizers of nondegenerate intervals are finite). When the acting group is torsion-free, we can equivalently say isometric actions with *trivial arc stabilizers*.

Suppose d_T is the metric on T ; an isometry $\iota: T \rightarrow T$ has translation distance given by

$$\|\iota\| = \inf_{p \in T} d_T(p, \iota(p)).$$

Any isometric H -action on a tree has a nonnegative function $\|\cdot\|: H \rightarrow \mathbb{R}_{\geq 0}$ given by the translation distances. For a group system $\mathcal{H} = \bigsqcup_{i \in \mathcal{I}} H_i$ and a forest $\mathcal{T} = \bigsqcup_{i \in \mathcal{I}} T_i$, an isometric \mathcal{H} -action on \mathcal{T} is a disjoint union of isometric H_i -actions on T_i .

A characteristic subtree for an isometric action on a tree T is a minimal invariant subtree in T . An isometric action on a tree T is minimal if T is the characteristic subtree.

Suppose F acts minimally on a nondegenerate tree T by isometries with trivial arc stabilizers. We now introduce the *index theory* for such actions. For each F -orbit of points $[p] \in F \setminus T$, let $G_p \leq F$ denote the stabilizer of $p \in T$ and $\#\text{dir}[p]$ denote the number of G_p -orbits of directions at p . The (local) *index* at $[p]$ is

$$i[p] := c(G_p) - 1 + \#\text{dir}[p] \in \mathbb{Z}_{\geq 0} \cup \{\infty\}.$$

The (global) index of $F \setminus T$ is

$$i(F \setminus T) := \sum_{[p] \in F \setminus T} i[p].$$

One of our main tools will be Gaboriau–Levitt’s index inequality [8, Theorem III.2]:

$$i(F \setminus T) < c(F).$$

For example, the index of any free splitting of F is $c(F) - 1$. In Appendix A, we will sketch Gaboriau–Levitt’s proof in a metric-agnostic setting. Corollaries of the inequality: there are finitely many F -orbits of directions at branch points and the point stabilizers subgroup system is represented by a system that has strictly lower complexity than F . For a free group system of finite type, the index of its isometric action on a forest is the sum of indices for each component, and the index inequality still holds.

2.2. Pretrees and rigid actions

Brian Bowditch’s paper [5] is a good survey about pretrees and other ‘treelike’ structures. A pretree is a (nonempty) set T and a function $[\cdot, \cdot]: T \times T \rightarrow \mathcal{P}(T)$ (i.e., an association of a subset $[p, q] \subset T$ to each pair of points $p, q \in T$) that satisfies the *pretree axioms*: for all $p, q, r \in T$,

1. (symmetric) $[p, q] = [q, p]$ contains $\{p, q\}$;
2. (thin) $[p, r] \subset [p, q] \cup [q, r]$; and
3. (linear) if $r \in [p, q]$ and $q \in [p, r]$, then $q = r$.

The subsets $[p, q]$ from the definition will be referred to as *closed intervals* and should be thought of as encoding a ‘betweenness’ relation on T . Define the *open interval* (p, q) to be the subset $[p, q] \setminus \{p, q\}$. Similarly, define *half-open intervals* $[p, q) = (q, p) = [p, q] \setminus \{q\}$. Naturally, the real line \mathbb{R} is a pretree; a tree has a (canonical) pretree structure where closed intervals are the closed geodesic segments; also, any subset $S \subset T$ of a pretree inherits a pretree structure given by $[\cdot, \cdot]_S := [\cdot, \cdot] \cap S$.

A direction at $p \in T$ is a maximal subset $D_p \subset T$ for which $p \notin [q, r]$ for all $q, r \in D_p$. As we did for trees, a branch point for a pretree is a point with at least three directions. The *observers’ topology* on a pretree is the canonical topology generated by the subbasis of directions; Bowditch calls it the *order topology* [5, Section 7]. Generally, a direction at a nonempty subset $S \subset T$ is a maximal $D_S \subset T$ for which $[q, r] \cap S = \emptyset$ for all $q, r \in D_S$.

Suppose T and T' have pretree structures $[\cdot, \cdot]$ and $[\cdot, \cdot]'$, respectively. A set-bijection $f: T \rightarrow T'$ is a pretree-isomorphism if $[f(p), f(q)]' = f([p, q])$ for all $p, q \in T$. An injection $f: T \rightarrow T'$ is a pretree-embedding if it is a pretree-isomorphism onto the image $f(T)$ with the inherited pretree structure; we will only need pretree-embeddings for Appendix A. Note that pretree-isomorphisms induce homeomorphisms of the observers’ topologies. There is another canonical topology finer than the observers’ topology (see the interlude chapter); we will not need either canonical topologies for the results of the paper.

For a pretree-automorphism $f: T \rightarrow T$, the fixed-point set $\text{Fix}_T(f)$ is the subset of points in T fixed by f . A pretree-automorphism f is rigid if either $\text{Fix}_T(f)$ is empty or f fixes no direction at $\text{Fix}_T(f)$. An isometry of a (subset of a) tree is a rigid pretree-automorphism of its (inherited) pretree structure. A subset $C \subset T$ is convex if $[p, q] \subset C$ for all $p, q \in C$. The fixed-point set $\text{Fix}_T(f)$ of a rigid pretree-automorphism is convex.

A pretree is real if every closed interval $[p, q]$ is pretree-isomorphic to a closed interval of \mathbb{R} . The pretree structure of a tree is real. A nondegenerate convex subset $A \subset T$ is an arc if any $x, y, z \in A$ simultaneously lie in some closed interval. A pretree is complete if every arc is an interval. Notably, any real pretree T embeds in a canonical complete pretree \widehat{T} known as the pretree completion [5, Lemma 7.14]. A real pretree is short if every arc is pretree-isomorphic to an arc in \mathbb{R} . The real pretree structure of a tree is short. For any short real pretree T , the pretree completion \widehat{T} is real [5, Lemma 7.15]. For example, the pretree \mathbb{R} is not complete (it is an arc but not an interval), but it is short by definition; the pretree completion of \mathbb{R} is the extended real line, which has two additional points $\{\pm\infty\}$, may be denoted $[-\infty, \infty]$, and is pretree-isomorphic to a closed interval of \mathbb{R} . The long line is the prototype of

a real pretree whose pretree completion, the *extended long line*, is not real: it is a closed interval not pretree-isomorphic to a closed interval of \mathbb{R} .

A pretree-automorphism $f: T \rightarrow T$ of a real pretree is elliptic if it has a fixed point and loxodromic otherwise; in the latter case, there is a maximal arc $A \subset T$, known as the *axis*, preserved by f . For an elliptic pretree-automorphism f of a real pretree T , the complement $T \setminus \text{Fix}_T(f)$ is ‘open’ in the following sense: the complement $(p, q) \setminus \text{Fix}_T(f)$ is a union of open intervals for any $p, q \in T$; in particular, a direction d at $\text{Fix}_T(f)$ has an *attaching point* $p_d \in \text{Fix}_T(f)$. If f is also rigid, then the direction d has a unique attaching point due to convexity. In the literature, rigid pretree-automorphisms of real pretrees have been studied as *non-nesting homeomorphisms* [12] – Levitt’s results are stated with \mathbb{R} -trees, but the metrics are never used and so the results apply to pretrees as well. Again, we mostly only care about *rigid actions* on real pretrees with *finite arc stabilizers* (i.e., actions by rigid pretree-automorphisms where the (pointwise) stabilizers of arcs are finite).

A characteristic convex subset for a rigid action on a real pretree T is a minimal invariant nonempty convex subset of T . A rigid action on a real pretree T is minimal if T is the characteristic convex subset. Note that a real pretree T that admits a minimal rigid action by a countable group must be short: T will be a countable union of closed intervals; so its arcs are countable ascending unions of closed intervals and hence pretree-isomorphic to arcs in \mathbb{R} .

Suppose we have a minimal rigid F -action on a nondegenerate real pretree T with trivial arc stabilizers. Then we can define the index of $F \setminus T$ exactly as we did for minimal isometric F -actions with trivial arc stabilizers. Gaboriau–Levitt’s index inequality still holds since their proof extends almost verbatim to this setting of minimal rigid actions with trivial arc stabilizers (see Appendix A for the sketch).

We now introduce a new term that will be our replacement for expanding homotheties in the real pretree setting. Let $\phi: F \rightarrow F$ be an automorphism and T a real pretree with a chosen minimal rigid F -action whose arc stabilizers are trivial. Recall that $f: T \rightarrow T$ is ϕ -equivariant if $f(x \cdot p) = \phi(x) \cdot f(p)$ for all $x \in F$ and $p \in T$. A ϕ -equivariant pretree-automorphism $f: T \rightarrow T$ expands at $p \in \text{Fix}_T(f)$ if each orbit $G_p \cdot d$ of directions at p contains some half-open interval $(p, q_d]$ and each $(p, q_d]$ properly embeds in a G_p -translate of $f((p, q_{d'})]$ for some orbit $G_p \cdot d'$ of directions at p . By the index inequality, there are finitely many G_p -orbits of directions at p ; so for some $n \geq 1$, each $(p, q_d]$ properly embeds in a G_p -translate of $f^n((p, q_d])$ – this justifies the term ‘expanding’.

A ϕ -equivariant pretree-automorphism $f: T \rightarrow T$ is F -expanding if for any element $x \in F$ and integer $n \geq 1$, the composition $x \circ f^n: T \rightarrow T$ is either 1) elliptic with a unique fixed point and it expands at this fixed point or 2) loxodromic with an axis that is not shared with any T -loxodromic element in F . For the motivation behind this definition, note that ϕ -equivariant expanding homotheties are (by the contraction mapping theorem) F -expanding pretree-automorphism of the canonical pretree structure.

Interlude – same trees, different views

This interlude is meant to describe the different perspectives on ‘simplicial trees’. As it was borne out of my own failure to appreciate the differences before starting this project, I am writing the interlude from a personal point of view.

What is a *simplicial tree*? Figure 1 is visual representation of a simplicial tree that will be the running example for the interlude.

The most elementary definition describes it as a combinatorial object. A *simple graph* is a pair (V, E) , where V is the vertex set and the edge set E is a collection of size 2 subsets of V . I could use this to define *reduced paths* and *cycles*; then a simple tree would be a path-connected cycle-free simple graph. For example, this was the language used by Serre in [20], and it is the quickest way to define *free splittings*. As a simple tree, the running example can be defined as: $V = \mathbb{Z}_{\geq 0}$ and $E = \{ \{0, n\} : n \geq 1 \}$. This definition is unique up to a simplicial automorphism: a set-bijection of vertex sets that induces a set-bijection of the edge sets.

This perspective is especially useful if you are interested in algorithmic questions. The downside is that it can get quite cumbersome to describe maps between simple trees that send edges to paths.

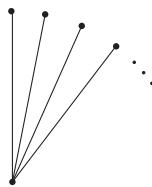


Figure 1. A tree with one ‘branch point’ and (countably) infinitely many edges attached.

Instead, it helps to work with topological spaces where the language of maps and their deformations is already well-established.

The first topological definition: a cellular tree is a 1-dimensional contractible CW-complex. This was the parenthetical definition I gave at the start of Chapter 1. This definition assumes you already know what a CW-complex is. For our purposes, a 1-dimension CW-complex is a quotient of the disjoint union of a discrete space (vertex set) and copies of closed intervals $[0, 1]$ with an equivalence relation identifying each endpoint of the closed intervals with some vertex. The complex is endowed with the quotient topology, and I will skip defining contractibility. If you were a pedantic reader, then you may have noticed that I abused terminology while defining free splittings: I define a simplicial tree as a ‘cellular tree’, yet I require that the free group act by ‘simplicial automorphism’; I never exactly explained what a simplicial automorphism of a CW-complex is! The moral of the story is that I have prioritized brevity. Anyway, as a cellular tree, the running example is defined as the quotient of $\{v\} \sqcup (\mathbb{Z}_{\geq 1} \times [0, 1])$ with the equivalence relation generated by $v \sim (n, 0)$ for all $n \geq 1$. The set of equivalence classes (without the topology) will be denoted T , and let π be the quotient function onto T .

Now that I am dealing with topological spaces, I can discuss concepts like convergence or compactness. Let $x_n = (1, \frac{1}{n})$, $y_n = (n, \frac{1}{n})$, and $z_n = (n, 1)$. Then the sequence $(x_n)_{n \geq 1}$ converges to $\pi(v)$ in the CW/quotient topology, while the sets $\{y_n\}_{n \geq 1}$ and $\{z_n\}_{n \geq 1}$ have no limit points; moreover, this topology is not metrizable.

The metric definition: a discrete (metric) tree is a (metric) tree that is the convex hull of its singular points (i.e., number of directions at the point is not 2), and the subspace of these points is discrete. Metric spaces have a topology generated by the basis of open balls. In practice, this topology is secondary, and it is usually more convenient to work with the metric directly. To view T as a discrete tree, equip it with the combinatorial *convex* metric:

$$d_T(\pi(n, a), \pi(m, b)) = \begin{cases} |a - b| & \text{if } m = n \\ a + b & \text{otherwise.} \end{cases}$$

The metric topology on T is *strictly coarser* than the CW-topology as it contains fewer open sets. This time, the sequences $(x_n)_{n \geq 1}$ and $(y_n)_{n \geq 1}$ both converge to $\pi(v)$ in the metric topology. However, the set $\{z_n\}_{n \geq 1}$ still has no limit points. So the metric topology is not compact.

As you shall see in the next chapter, the metric setting helps us understand what an automorphism does under iteration, especially at the ‘limit’. This starts with equipping free splittings (simple or cellular trees) with the eigenmetric of an irreducible train track and taking the *projective-limit* of iterating the train track to get a useful and probably nondiscrete tree.

The final topological definition: a (separable) ‘visual’ tree is a connected subspace of a dendrite (i.e., compact separable metric tree) that is the convex hull of its own singular points and whose branch points cannot accumulate to a branch point along one direction. I could remove the separability condition by replacing ‘dendrite’ with ‘dendron’, but that would necessitate a definition for dendrons. There is no metric characterization for dendrons as there is for dendrites, and, unfortunately, their purely topological definition is beyond the scope of this paper; see [5] for details. A pretree characterization for dendrons can be the observers’ topology on a complete real pretree. As a visual tree, the running example is the observers’ topology on the cellular or metric tree T . The observers’ topology is even coarser than the

metric topology, and it makes the sequences $(x_n)_{n \geq 1}$, $(y_n)_{n \geq 1}$ and $(z_n)_{n \geq 1}$ all converge to $\pi(v)$. In fact, T with the observers' topology is compact – it is a dendrite! As a visual tree, the topology is metrizable, and one compatible convex metric is given by

$$d'_T(\pi(m, a), \pi(n, b)) = \begin{cases} \frac{1}{m}|a - b| & \text{if } m = n \\ \frac{a}{m} + \frac{b}{n} & \text{otherwise.} \end{cases}$$

The underlying thesis of this project is that the ‘metric category’ might not be right for defining *blow-ups* of nondiscrete trees. It is not clear to me that it is even possible in general! However, defining blow-ups in the ‘visual category’ is rather natural. In fact, we avoid topology altogether and carry out the construction in the ‘pretree category’!

Finally, a topological-combinatorial hybrid definition: a *simple pretree* is a pretree whose closed intervals are finite sets; and a simple real pretree is real pretree that is the convex hull of its singular points, and these points form a simple pretree. The simple pretree definition is simply paraphrasing that of a simple tree, while the simple real pretree definition simultaneously captures the combinatorial nature of trees and allows topological approaches. Simple real pretrees have two canonical topologies: the coarser one is the observers’ topology, and the finer one is more-or-less the CW-topology.

So what is a simplicial tree? Well, it could be a simple tree, cellular tree, discrete tree, visual tree or simple real pretree; the answer depends on what I need it for. For instance, if I want to discuss ‘simplicial actions’ on \mathbb{R} , then only the first two definitions are applicable – \mathbb{R} has no singular points!

3. Limit pretrees: exponentially growing automorphisms

Returning to free group automorphisms, we can now use translation distances in free splittings (with invariant combinatorial metrics) to classify free group automorphisms.

Fix an automorphism $\psi: \mathcal{F} \rightarrow \mathcal{F}$ and a $[\psi]$ -invariant proper free factor system \mathcal{G} of \mathcal{F} . By Theorem 1.1, there is an irreducible train track $\tau^{(1)}$ for ψ on a free splitting $(\Gamma_*^{(1)}, \mathcal{F}_2)$ of $\mathcal{F}_1 = \mathcal{F}$, where \mathcal{F}_2 carries \mathcal{G} . The train track $\tau^{(1)}$ determines a restriction $\psi|_{\mathcal{F}_2}$; applying the theorem repeatedly, we get a finite *descending sequence* of irreducible train tracks $\tau^{(i)}$ for restrictions $\psi|_{\mathcal{F}_i}$ on free splittings $(\Gamma_*^{(i)}, \mathcal{F}_{i+1})$ of \mathcal{F}_i , and the final train track $\tau^{(k)}$ in the sequence is defined rel. \mathcal{G} (i.e., $\mathcal{F}_{k+1} = \mathcal{G}$). The next proposition collects standard facts that immediately follow from this train track theory; we sketch the proof mainly to highlight the study of automorphisms through the *blow-up of a free splitting rel. another free splitting*.

Proposition 3.1. *Let $\psi: \mathcal{F} \rightarrow \mathcal{F}$ be an automorphism, \mathcal{G} a $[\psi]$ -invariant proper free factor system of \mathcal{F} , and (Λ_*, \mathcal{G}) a free splitting of \mathcal{F} . For any element x in \mathcal{F} , the limit inferior of the sequence $\left(\sqrt[n]{\|\psi^n(x)\|_{\Lambda_*}}\right)_{n \geq 1}$ is finite and independent of (Λ_*, \mathcal{G}) .*

Let $(\tau^{(i)})_{i=1}^k$ be a descending sequence of irreducible train tracks with $\tau^{(k)}$ defined rel. \mathcal{G} . Then the following are equivalent:

1. $\lambda(\tau^{(i)}) = 1$ for all $i = 1, \dots, k$;
2. for any $x \in \mathcal{F}$, the sequence $(\|\psi^n(x)\|_{\Lambda_*})_{n \geq 1}$ is bounded by a polynomial in n ; and
3. $\liminf_{n \rightarrow \infty} \sqrt[n]{\|\psi^n(x)\|_{\Lambda_*}} \leq 1$ for all $x \in \mathcal{F}$.

An automorphism is polynomially growing rel. \mathcal{G} if these conditions hold; otherwise, it is exponentially growing rel. \mathcal{G} .

Sketch of proof. The automorphism ψ has a Lipschitz topological representative on the free splitting (Λ_*, \mathcal{G}) of \mathcal{F} . So for any element x in \mathcal{F} , we can set $\underline{\lambda}_x := \liminf_{n \rightarrow \infty} \sqrt[n]{\|\psi^n(x)\|_{\Lambda_*}}$ in $\mathbb{R}_{\geq 0}$. By $[\psi]$ -invariance of \mathcal{G} , $\underline{\lambda}_x = 0$ if $\|x\|_{\Lambda_*} = 0$. Since 0 is isolated in the image of $\|\cdot\|_{\Lambda_*}$, we get $\underline{\lambda}_x \geq 1$ if $\|x\|_{\Lambda_*} > 0$. Any pair

of free splittings (Λ_*, \mathcal{G}) and (Ω_*, \mathcal{G}) of \mathcal{F} have equivariant Lipschitz maps between them, and so their translation distance functions are comparable; that is,

$$\frac{1}{K} \|\cdot\|_{\Lambda_*} \leq \|\cdot\|_{\Omega_*} \leq K \|\cdot\|_{\Lambda_*} \text{ (pointwise) for some constant } K \geq 1.$$

This implies $\underline{\lambda}_x$ is independent of the free splitting (Λ_*, \mathcal{G}) and concludes the first part.

(1 \implies 2): Suppose $\lambda(\tau^{(i)}) = 1$ for all $i = 1, \dots, k$. So the train tracks $\tau^{(i)}$ are simplicial automorphisms; in particular, all edge-paths in $(\Omega_*^{(k)}, \mathcal{G}) = (\Gamma_*^{(k)}, \mathcal{G})$ have constant growth. For induction, assume all edge-paths in $(\Omega_*^{(i)}, \mathcal{G})$ have at most degree $(k-i)$ polynomial growth. Let $(\Omega_*^{(i-1)}, \mathcal{G})$ be a blow-up of $(\Gamma_*^{(i-1)}, \mathcal{F}_i)$ with respect to $(\Omega_*^{(i)}, \mathcal{G})$. As $\tau^{(i-1)}$ is a simplicial automorphism, each ‘top-stratum’ edge gains a predetermined prefix and suffix in $(\Omega_*^{(i)}, \mathcal{G})$ under iteration; these prefixes and suffixes have at most degree $(k-i)$ polynomial growth. So edge-paths in $(\Omega_*^{(i-1)}, \mathcal{G})$ have at most degree $(k-i+1)$ polynomial growth. By induction and comparability again, edge-paths in (Λ_*, \mathcal{G}) have at most degree $(k-1)$ polynomial growth.

(2 \implies 3): Polynomials are sub-exponential.

(3 \implies 1): Suppose $\lambda = \lambda(\tau^{(i)}) > 1$ for some i . Equip $(\Gamma_*^{(i)}, \mathcal{F}_{i+1})$ with the eigenmetric d_i , and let $x \in \mathcal{F}_i$ be a $\tau^{(i)}$ -legal element (Proposition 1.2). The choice of metric and legality of x imply $\|\psi^n(x)\|_{d_i} = \lambda^n \cdot \|x\|_{d_i} > 0$. Let $(\Omega_*^{(1)}, \mathcal{G})$ be the free splitting constructed above. Then $\|\cdot\|_{d_i}$ and $\|\cdot\|_{\Gamma_*^{(i)}}$ are comparable, $\|\cdot\|_{\Gamma_*^{(i)}} \leq \|\cdot\|_{\Omega_*^{(i)}}$ pointwise, and $\|\cdot\|_{\Omega_*^{(i)}}$ is the restriction of $\|\cdot\|_{\Omega_*^{(1)}}$. So $\underline{\lambda}_x \geq \lambda > 1$. \square

For any element x in \mathcal{F} , define $\underline{\lambda}(x; \psi, \mathcal{G}) := \liminf_{n \rightarrow \infty} \sqrt[n]{\|\psi^n(x)\|_{\mathcal{G}}}$. We say $x \in \mathcal{F}$ grows exponentially rel. \mathcal{G} if $\underline{\lambda}(x; \psi, \mathcal{G}) > 1$; otherwise, it grows polynomially rel. \mathcal{G} . The ‘rel. \mathcal{G} ’ will be omitted when \mathcal{G} is empty. Suppose \mathcal{H} is a $[\psi]$ -invariant subgroup system of finite type that carries \mathcal{G} . By passing to the characteristic subforest for \mathcal{H} in a free splitting (Γ_*, \mathcal{G}) of \mathcal{F} , we see that $\underline{\lambda}(x; \phi|_{\mathcal{H}}, \mathcal{G}) = \underline{\lambda}(x; \psi, \mathcal{G})$ for elements x in \mathcal{H} . We use this observation whenever we pass to invariant subgroup systems of finite type.

Remark. It was known since the introduction of train tracks that the first sequence in the statement of Proposition 3.1 converges (see [3, Remark 1.8]). Gilbert Levitt gave a finer classification for an element’s growth rate [13, Theorem 6.2].

The next proposition appears in Levitt–Lustig’s paper [14, Proposition 3.2]. Our proof takes a slightly different approach using the *blow-up of a free splitting rel. expanding forest*. This is done to highlight ideas crucial to the proof of the main theorem.

Proposition 3.2. *Let $\psi: \mathcal{F} \rightarrow \mathcal{F}$ be an automorphism and \mathcal{G} a $[\psi]$ -invariant proper free factor system of \mathcal{F} . If $\psi: \mathcal{F} \rightarrow \mathcal{F}$ is exponentially growing rel. \mathcal{G} , then there is a minimal isometric \mathcal{F} -action on a nondegenerate forest \mathcal{Y} with trivial arc stabilizers for which \mathcal{G} is elliptic, and a ψ -equivariant expanding homothety $h: \mathcal{Y} \rightarrow \mathcal{Y}$.*

In particular, \mathcal{Y} -loxodromic elements in \mathcal{F} grow exponentially rel. \mathcal{G} .

Proof (modulo a black box). Suppose $\psi: \mathcal{F} \rightarrow \mathcal{F}$ is exponentially growing rel. \mathcal{G} , and set $\mathcal{F}_1 = \mathcal{F}$. Then there is a finite sequence of irreducible train tracks $\tau^{(i)}$ for restrictions $\psi|_{\mathcal{F}_i}$ on nondegenerate free splittings $(\Gamma_*^{(i)}, \mathcal{F}_{i+1})$ of \mathcal{F}_i with $\lambda(\tau^{(i)}) > 1$ for some $i \geq 1$; we equip the free splittings with eigenmetrics d_i . After truncating the sequence if necessary, assume $\lambda(\tau^{(i)}) = 1$ for $i = 1, \dots, k-1$ and $\lambda = \lambda(\tau^{(k)}) > 1$. By Proposition 1.2, there is a $\tau^{(k)}$ -legal element x_0 in \mathcal{F}_k .

As $\tau^{(k)}$ is λ -Lipschitz, we can define a limit function

$$\|\cdot\|^{(k)}: \mathcal{F}_k \rightarrow \mathbb{R}_{\geq 0} \quad \text{by} \quad x \mapsto \lim_{n \rightarrow \infty} \lambda^{-n} \|\psi^n(x)\|_{\Gamma_*^{(k)}}.$$

Culler–Morgan proved $\|\cdot\|^{(k)}$ is a translation distance function for a minimal isometric \mathcal{F}_k -action on a forest $\mathcal{T}^{(k)}$ with *cyclic arc stabilizers* [6, Theorem 5.3]; the forest is not degenerate as

$\|x_0\|^{(k)} = \|x_0\|_{\Gamma_v^{(k)}} > 0$. Lustig proved the \mathcal{F}_k -action on $\mathcal{T}^{(k)}$ has trivial arc stabilizers [9, Appendix]. There is a $\psi|_{\mathcal{F}_k}$ -equivariant λ -homothety $h^{(k)}: \mathcal{T}^{(k)} \rightarrow \mathcal{T}^{(k)}$ since $\|\psi(x)\|^{(k)} = \lambda\|x\|^{(k)}$ for all elements x in \mathcal{F}_k [6, Theorem 3.7]. By construction, the $[\psi]$ -invariant proper free factor system \mathcal{G} is $\mathcal{T}^{(k)}$ -elliptic.

The main step of the proof is a construction that allows us to ‘merge’ the higher *polynomial strata* and the *exponential stratum* at k . If $k = 1$, then set $\mathcal{Y} = \mathcal{T}^{(k)}$, $h = h^{(k)}$, and we are done. Otherwise, $k \geq 2$ and, for some $i \leq k$, there is a minimal isometric \mathcal{F}_i -action on a nondegenerate forest $\mathcal{T}^{(i)}$ with trivial arc stabilizers and an equivariant copy of $\mathcal{T}^{(k)}$, and a $\psi|_{\mathcal{F}_i}$ -equivariant λ -homothety $h^{(i)}: \mathcal{T}^{(i)} \rightarrow \mathcal{T}^{(i)}$.

In the next chapter (simple patchwork 4.1, black box), we construct a unique forest $\mathcal{T}^{(i-1)}$ that is ‘an equivariant blow-up and simplicial collapse’ of $(\Gamma_*^{(i-1)}, \mathcal{F}_i)$ with respect to $\tau^{(i-1)}$ and $h^{(i)}$. In particular, there is a minimal isometric \mathcal{F}_{i-1} -action on $\mathcal{T}^{(i-1)}$ with trivial arc stabilizers and an equivariant copy of $\mathcal{T}^{(i)}$, and a $\psi|_{\mathcal{F}_{i-1}}$ -equivariant λ -homothety $h^{(i-1)}$ on $\mathcal{T}^{(i-1)}$ induced by $\tau^{(i-1)}$ and $h^{(i)}$. By induction, we have a minimal isometric \mathcal{F} -action on a forest $\mathcal{Y} = \mathcal{T}^{(1)}$ with trivial arc stabilizers, an equivariant copy of $\mathcal{T}^{(k)}$ and a ψ -equivariant λ -homothety $h = h^{(1)}$ on \mathcal{Y} .

For the last part of the proposition, choose a free splitting (Λ_*, \mathcal{G}) for \mathcal{F} . Let $\|\cdot\|_{\Lambda_*}$ and $\|\cdot\|_{\mathcal{Y}}$ be the translation distance functions for (Λ_*, \mathcal{G}) and \mathcal{Y} , respectively. Since \mathcal{G} is \mathcal{Y} -elliptic, there is an equivariant Lipschitz map $(\Lambda_*, \mathcal{G}) \rightarrow \mathcal{Y}$ and

$$\|\cdot\|_{\mathcal{Y}} \leq K \|\cdot\|_{\Lambda_*} \text{ pointwise for some constant } K \geq 1.$$

Therefore, $\lambda^n \|\cdot\|_{\mathcal{Y}} = \|\psi^n(\cdot)\|_{\mathcal{Y}} \leq K \|\psi^n(\cdot)\|_{\Lambda_*}$. So \mathcal{Y} -loxodromic elements in \mathcal{F} grow exponentially rel. \mathcal{G} . \square

The main theorem of this paper associates to a free group automorphism a real pretree with a minimal rigid F -action whose loxodromic elements are precisely the elements that grow exponentially. This time, we need a more delicate idea than the ones highlighted in the proofs of Propositions 3.1 and 3.2: the *blow-up of an expanding real pretree rel. an expanding forest*.

Theorem 3.3. *If $\phi: F \rightarrow F$ is an automorphism and \mathcal{G} a $[\phi]$ -invariant proper free factor system of F , then there is*

1. *a minimal rigid F -action on a real pretree T with trivial arc stabilizers;*
2. *a ϕ -equivariant F -expanding pretree-automorphism $f: T \rightarrow T$; and*
3. *an element in F is T -loxodromic if and only if it grows exponentially rel. \mathcal{G} .*

The T -point stabilizers give a canonical $[\phi]$ -invariant subgroup system \mathcal{H} of finite type, and any restriction $\phi|_{\mathcal{H}}$ is polynomially growing rel. \mathcal{G} (if \mathcal{H} is not empty). This system has strictly lower complexity than F if and only if ϕ is exponentially growing rel. \mathcal{G} .

A (forward) limit pretree of a free group automorphism is a real pretree satisfying the theorem’s conclusion.

Proof (modulo a black box). If ϕ is polynomially growing rel. \mathcal{G} , then let T be the degenerate real pretree (i.e., a singleton) with a trivial F -action; Conditions 1–3 hold automatically. For the rest of the proof, we assume that ϕ is exponentially growing rel. \mathcal{G} .

Let $Y^{(1)}$ be a nondegenerate tree for ϕ given by Proposition 3.2. By Gaboriau–Levitt’s index inequality, $i(F \setminus Y^{(1)}) < c(F)$, we can define \mathcal{V} to be a finite set of representatives for the F -orbits of branch points v with nontrivial stabilizers G_v . The index inequality implies the subgroup system $\mathcal{G}^{(2)} = \bigsqcup_{v \in \mathcal{V}} G_v$ has complexity $c(\mathcal{G}^{(2)}) < c(F)$ as the tree $Y^{(1)}$ is not degenerate and the F -action is minimal. The system is $[\phi]$ -invariant and has a restriction automorphism $\phi|_{\mathcal{G}^{(2)}}: \mathcal{G}^{(2)} \rightarrow \mathcal{G}^{(2)}$ that is unique up to post-composition with inner automorphisms of $\mathcal{G}^{(2)}$ -components – just as we showed for free splittings in Chapter 1. If $\mathcal{G}^{(2)}$ is not empty and the restriction $\phi|_{\mathcal{G}^{(2)}}$ is exponentially growing rel. \mathcal{G} , then we can repeatedly apply Proposition 3.2 to the restrictions.

The complexities of the point stabilizer systems are strictly descending, and in the end, we get a finite sequence of nondegenerate forests $\mathcal{Y}^{(i)}$, minimal isometric $\mathcal{G}^{(i)}$ -actions with trivial arc stabilizers and $\phi|_{\mathcal{G}^{(i)}}$ -equivariant expanding homotheties $h^{(i)}: \mathcal{Y}^{(i)} \rightarrow \mathcal{Y}^{(i)}$. An element in F grows exponentially rel. \mathcal{G} if and only if it is conjugate to a $\mathcal{Y}^{(i)}$ -loxodromic element in some system $\mathcal{G}^{(i)}$. An almost identical argument is in the preliminaries of [13].

Set $T^{(1)} := Y^{(1)}$ with its canonical pretree structure and $f^{(1)} := h^{(1)}$. So there is a minimal rigid F -action on a real pretree $T^{(i-1)}$ with trivial arc stabilizers (for some $i \geq 2$), a ϕ -equivariant F -expanding pretree-automorphism $f^{(i-1)}$ and the $T^{(i-1)}$ -point stabilizers are represented by $\mathcal{G}^{(i)}$.

The novel construction in the paper (ideal stitching 4.4, black box) defines a unique real pretree $T^{(i)}$ that is an equivariant blow-up of $T^{(i-1)}$ with respect to $f^{(i-1)}$ and $h^{(i)}$. In particular, there is

1. a minimal rigid F -action on $T^{(i)}$ with trivial arc stabilizers;
2. a ϕ -equivariant F -expanding pretree-automorphism $f^{(i)}: T^{(i)} \rightarrow T^{(i)}$; and
3. $x \in F$ is $T^{(i)}$ -loxodromic if and only if it is $\mathcal{Y}^{(j)}$ -loxodromic for some $j \leq i$.

By induction, we may assume we have a minimal rigid F -action on a real pretree T with trivial arc stabilizers; a ϕ -equivariant F -expanding pretree-automorphism $f: T \rightarrow T$; and an element of F is T -loxodromic if and only if it is loxodromic in some forest $\mathcal{Y}^{(i)}$. By construction of the sequence of forests, the last condition translates to the following: an element in F is T -loxodromic if and only if it grows exponentially rel. \mathcal{G} , as required.

Thus, the T -point stabilizers are represented by a canonical $[\phi]$ -invariant subgroup system \mathcal{H} . Since ϕ is exponentially growing rel. \mathcal{G} , T -loxodromic elements exist and T is not degenerate. By the index inequality, we get $c(\mathcal{H}) < c(\mathcal{F})$. \square

We can now give the more natural characterization of polynomially growing elements:

Corollary 3.4. *Let $\psi: \mathcal{F} \rightarrow \mathcal{F}$ be an automorphism, \mathcal{G} a $[\psi]$ -invariant proper free factor system of \mathcal{F} , and (Λ_*, \mathcal{G}) a free splitting of \mathcal{F} . An element x in \mathcal{F} grows polynomially rel. \mathcal{G} if and only if the sequence $(\|\psi^n(x)\|_{\Lambda_*})_{n \geq 1}$ is bounded by a polynomial in n .*

Proof. The reverse direction is immediate: polynomials are subexponential. Suppose a nontrivial element $x \in \mathcal{F}$ grows polynomially rel. \mathcal{G} . Then it has a conjugate in \mathcal{H} , the canonical $[\phi]$ -invariant subgroup system of finite type given by Theorem 3.3(3). As any restriction $\phi|_{\mathcal{H}}$ is polynomially growing rel. \mathcal{G} , we are done by Proposition 3.1(3 \Rightarrow 2). \square

We also give a dendrological characterization of *purely exponentially growing* automorphisms. An automorphism $\phi: F \rightarrow F$ is atoroidal if there are no $[\phi]$ -periodic conjugacy classes of nontrivial elements in F .

Corollary 3.5. *Suppose $\phi: F \rightarrow F$ is a free group automorphism with a limit pretree T . Then the following are equivalent:*

1. ϕ is atoroidal;
2. all nontrivial elements in F grow exponentially; and
3. the F -action on T is free.

Proof. (1 \Rightarrow 2): Set $\mathcal{G} := \emptyset$. If F has a nontrivial element that grows polynomially, then the canonical $[\phi]$ -invariant subgroup system \mathcal{H} of finite type given by Theorem 3.3 is not empty and any restriction $\phi|_{\mathcal{H}}$ is polynomially growing. By Proposition 3.1(3 \Rightarrow 1), conjugacy classes of nontrivial elements in the lowest stratum of $\phi|_{\mathcal{H}}$ are $[\phi]$ -periodic.

(2 \Rightarrow 1): Any nontrivial element in a $[\phi]$ -periodic conjugacy class has polynomial (in fact, ‘constant’) growth by definition.

(2 \Longleftrightarrow 3): This is Theorem 3.3(3). \square

4. Blow-ups: simple, naïve, and ideal

We now define the equivariant blow-ups that were the key steps in the proofs of Proposition 3.2 and Theorem 3.3. Although we present the constructions as two separate ideas (simplicial and non-simplicial), the underlying principle is the same:

1. start with an F -action on a tree (or real pretree) T and a \mathcal{G} -action on a forest \mathcal{T}_V where \mathcal{G} are conjugacy representatives of the nontrivial point stabilizers of T ;
2. then define a naïve equivariant blow-up of T with respect to \mathcal{T}_V ;
3. there is a lot of freedom in the naïve construction, but most choices will not be useful;
4. let $f: T \rightarrow T$ and $f_V: \mathcal{T}_V \rightarrow \mathcal{T}_V$ be ϕ - and $\phi|_{\mathcal{G}}$ -equivariant homeomorphisms, respectively;
5. consider the equivariant copies of \mathcal{T}_V in some blow-up tree (or real pretree) T^* ; then f induces a ϕ -equivariant function $f^*: T^* \rightarrow T^*$ whose restriction to the copies is f_V – in a way, f^* is formed by *stitching* f and f_V together;
6. for most blow-ups, f^* will not be a homeomorphism; but if f_V is *expanding*, then a unique fixed point theorem produces the ‘ideal’ blow-up T^* whose corresponding map f^* is a homeomorphism.

4.1. Simple patchwork

In this case, a blow-up is easy but tedious to define; ensuring the induced map is a homothety is the tricky bit.

Theorem 4.1. *Let $\psi: \mathcal{F} \rightarrow \mathcal{F}$ be an automorphism of a free group system \mathcal{F} and (Γ_*, \mathcal{G}) a free splitting of \mathcal{F} . Assume there is*

1. a ψ -equivariant simplicial automorphism $\tau: (\Gamma_*, \mathcal{G}) \rightarrow (\Gamma_*, \mathcal{G})$,
2. a minimal isometric \mathcal{G} -action on a forest \mathcal{T}_V with trivial arc stabilizers; and
3. a $\psi|_{\mathcal{G}}$ -equivariant expanding λ -homothety $h_V: \mathcal{T}_V \rightarrow \mathcal{T}_V$.

Then there is

1. a minimal isometric \mathcal{F} -action on a forest \mathcal{T} with trivial arc stabilizers, an equivariant copy of \mathcal{T}_V , and
2. a ψ -equivariant expanding λ -homothety $h: \mathcal{T} \rightarrow \mathcal{T}$ induced by τ and h_V .

In fact, the \mathcal{F} -action on \mathcal{T} decomposes as a graph of actions whose underlying simplicial action is the free splitting (Γ_*, \mathcal{G}) , and the vertex actions are the given \mathcal{G} -action on \mathcal{T}_V . The homothety h acts by τ on the underlying simplicial action and h_V on the vertex actions. Any pair (\mathcal{T}', h') satisfying this conclusion admits an equivariant isometry $\mathcal{T}' \rightarrow \mathcal{T}$ that conjugates h' to h .

We state and prove the theorem in terms of homotheties due to the specific needs in Proposition 3.2, but the argument actually holds if ‘expanding homothety’ is replaced with ‘expansions (or contractions)’. We only need a hypothesis that lets us use the contraction mapping theorem.

Proof. Suppose (Γ_*, \mathcal{G}) is a free splitting of \mathcal{F} . Let \mathcal{V} be a set of orbit representatives for vertices in the free splitting with nontrivial stabilizers. For each $v \in \mathcal{V}$, let \mathcal{D}_v be a set of orbit representatives for half-edges originating from v and \bar{T}_v , the metric completion of the \mathcal{T}_V -component corresponding to G_v .

Suppose $(p_d \in \bar{T}_v : v \in \mathcal{V}, d \in \mathcal{D}_v)$ is a choice of *attaching points* in \bar{T}_V .

(graph of actions) To simplify the discussion, we will pretend $\mathcal{F} = F$ is connected for the moment. Let E be the complement of the F -orbit of \mathcal{V} in (Γ, \mathcal{G}) ; essentially, we just want the half-edges in \mathcal{D}_v to now have distinct origins. E inherits a free F -action from the free splitting. Set $\mathcal{V}^* := F \times \bar{T}_V$; F acts on \mathcal{V}^* by left-multiplication on the first factor. The equivariant blow-up $T^* = T^*(p_d : v \in \mathcal{V}, d \in \mathcal{D}_v)$ of (Γ, \mathcal{G}) with respect to the forest \mathcal{T}_V is defined through the quotient $\mathcal{V}^* \sqcup E \xrightarrow{\iota} T^*$ given by the following identifications:

1. for each $v \in \mathcal{V}$, $d \in \mathcal{D}_v$ and $x \in F$, the origin of the half-edge $x \cdot d$ in E is identified with (x, p_d) in \mathcal{V}^* (i.e., the half-edge d is *attached* to p_d equivariantly); and
2. for each $v \in \mathcal{V}$, $p \in \bar{T}_v$, $s \in G_v$ and $x \in F$, the point (xs, p) in \mathcal{V}^* is identified with $(x, s \cdot p)$ in \mathcal{V}^* .

As the F -actions on \mathcal{V}^* and E respect the identifications, the blow-up T^* inherits an F -action by homeomorphisms of the quotient topology. For each $s \in G_v$, the image $\iota(s, \bar{T}_v)$ is a G_v -equivariant copy of \bar{T}_v since (Γ, \mathcal{G}) has trivial edge stabilizers; in general, we would need each attaching point $p_d \in \bar{T}_v$ to be fixed by the stabilizer of d in G_v .

The blow-up T^* with the quotient topology has a natural projection $\pi: T^* \rightarrow (\Gamma, \mathcal{G})$ whose point-preimages are connected: single point or a copy of \bar{T}_v . This implies T^* is *uniquely arcwise connected*. Since the vertex set of (Γ, \mathcal{G}) is a discrete subspace with finitely many F -orbits, any closed arc in T^* decomposes into finitely many subarcs that are in $\iota(E)$ or $\iota(\mathcal{V}^*)$. Thus, closed arcs in T^* inherit lengths from (Γ_*, \mathcal{G}) and $\bar{\mathcal{T}}_V$; this path metric makes T^* a (metric) tree. By the equivariant construction, the tree T^* inherits an isometric F -action with trivial arc stabilizers and an equivariant copy of \mathcal{T}_V .

The F -action on T^* is what Levitt calls a *graph of actions* [11, p. 32] – the underlying simplicial action is the free splitting (Γ, \mathcal{G}) . Different choices $(p_d : v \in \mathcal{V}, d \in \mathcal{D}_v)$ of attaching points may produce drastically different graphs of actions.

Returning to the general case where \mathcal{F} is possibly disconnected, we can apply the blow-up construction componentwise to get a forest \mathcal{T}^* with an isometric \mathcal{F} -action.

We have yet to use the ψ -equivariant simplicial automorphism $\tau: (\Gamma_*, \mathcal{G}) \rightarrow (\Gamma_*, \mathcal{G})$ or the $\psi|_{\mathcal{G}}$ -equivariant expanding λ -homothety $h_V: \mathcal{T}_V \rightarrow \mathcal{T}_V$.

Remark. We give two related proofs of the conclusion to the theorem. The first proof is short, but it does not generalize to non-metric settings – we only sketch it. The second proof is thorough as it contains the ideas needed later for the main construction.

(*projective limit*) Suppose we have a naïve blow-up \mathcal{T}^* (made from an arbitrary choice of attaching points), and let $\|\cdot\|^*$ be its translation distance function. Since τ was a simplicial automorphism and h_V is a $\psi|_{\mathcal{G}}$ -equivariant λ -homothety, we get $\lambda^{-n} \|\psi(\cdot)\|^* \rightarrow \|\cdot\|$, where the limit function $\|\cdot\|$ is the translation distance function for the required \mathcal{F} -action.

For the second proof, we get the ideal attaching points without taking projective limits.

(*ideal stitching*) As the simplicial automorphism τ permutes the (finitely many) orbits of vertices and half-edges, it induces permutations $\beta \in \text{Sym}(\mathcal{V})$ and $\partial \in \text{Sym}(\bigcup_{v \in \mathcal{V}} \mathcal{D}_v)$. The maps τ and h_V induce a ψ -equivariant *PL-map* $f^*: \mathcal{T}^* \rightarrow \mathcal{T}^*$:

o let $v: \mathcal{V}^* \rightarrow \iota(\mathcal{V}^*)$ be given by $v(x, p) := \iota(\psi(x)x_v, \bar{h}_v(p))$, where $(x_v \in \mathcal{F}: v \in \mathcal{V})$ was the implicit choice used to define the restriction $\psi|_{\mathcal{G}}$, and hence h_V and its extension $\bar{h}_V: \bar{\mathcal{T}}_V \rightarrow \bar{\mathcal{T}}_V$. Observe that $v(xs, p) = v(x, s \cdot p)$ when $s \in G_v$ (exercise).

Let $\epsilon: E \rightarrow \mathcal{T}^*$ be given by the following:

- o for each half-edge e in E not in the \mathcal{F} -orbit of $\bigcup_{v \in \mathcal{V}} \mathcal{D}_v$, ϵ maps e onto $\iota(\tau(e)) \subset E$ along an orientation-preserving isometry.
- o for each $v \in \mathcal{V}$, $d \in \mathcal{D}_v$, and element x in the \mathcal{F} -component that contains G_v , ϵ maps $x \cdot d$ onto the concatenation of $\iota(\psi(x)x_v, [\bar{h}_v(p_d), s_{\partial \cdot d} \cdot p_{\partial \cdot d}]_{\partial \cdot v})$ and $\iota(\psi(x) \cdot \tau(d))$ along an orientation-preserving linear map, where $\tau(d) = x_v s_{\partial \cdot d} \cdot \partial(d) - s_{\partial \cdot d} \in G_{\beta \cdot v}$ is unique as edge stabilizers are trivial. Notice that the origin of $\epsilon(x \cdot d)$ is $v(x, p_d)$.

By definition, $v \sqcup \epsilon: \mathcal{V}^* \sqcup E \rightarrow \mathcal{T}^*$ factors through ι to induce a PL-map f^* .

The PL-map f^* is injective if and only if $\bar{h}_v(p_d) = s_{\partial \cdot d} \cdot p_{\partial \cdot d}$ for all $v \in \mathcal{V}$, $d \in \mathcal{D}_v$. This finite system of equations with unknowns $(p_d \in \bar{T}_v : v \in \mathcal{V}, d \in \mathcal{D}_v)$ is restated as

$$p_d = h_{v,d}(p_{\partial \cdot d}) \text{ for all } v \in \mathcal{V}, d \in \mathcal{D}_v, \quad \text{where } h_{v,d} := \bar{h}_v^{-1} \circ s_{\partial \cdot d}: T_{\beta \cdot v} \rightarrow T_v.$$

Note that minimality of \mathcal{T}_V implies h_V is surjective and hence invertible. So $h_{v,d}: T_{\beta \cdot v} \rightarrow T_v$, the composition of an isometry with a *contracting* homothety, is a contracting homothety. As $\bigcup_{v \in \mathcal{V}} \mathcal{D}_v$ is finite and $\bar{\mathcal{T}}_V$ is complete, our system of equations has a unique solution by the contraction mapping theorem (i.e., there is a unique tuple $\vec{p} = (p_d \in \bar{T}_v : v \in \mathcal{V}, d \in \mathcal{D}_v)$ with $p_d = h_{v,d}(p_{\partial \cdot d})$ for all $v \in \mathcal{V}, d \in \mathcal{D}_v$).

Let $\mathcal{T}^* = \mathcal{T}^*(\vec{p})$ be the blow-up given by this unique solution. Then the homeomorphism f^* is isometric on permuted components of the simplicial part $\iota(E)$ and an expanding λ -homothety on components of the non-simplicial part $\iota(\mathcal{V}^*)$.

(simplicial collapse) To finish the proof, we collapse the simplicial edges to get an isometric \mathcal{F} -action on a forest with trivial arc stabilizers. Let \mathcal{T} be the characteristic subforest for \mathcal{F} (i.e., the \mathcal{F} -action on \mathcal{T} is minimal). Since we only collapsed the simplicial part, the PL-map f^* induces a ψ -equivariant expanding λ -homothety $h: \mathcal{T} \rightarrow \mathcal{T}$.

(uniqueness) Suppose some \mathcal{F} -forest \mathcal{T}' decomposed as a graph of actions with underlying simplicial action (Γ_*, \mathcal{G}) and vertex actions \mathcal{T}_V and admitted a homothety h' that acts by τ on (Γ_*, \mathcal{G}) and h_V on \mathcal{T}_V . This forest must arise from the above construction and the equivariant isometric identification with \mathcal{T} follows from uniqueness of the solution \vec{p} . \square

4.2. Naïve stitching

Due to the involved nature of the proofs when dealing with non-simplicial trees, the blow-up and stitching constructions are split into two sections; let us start with blow-ups.

Suppose a countable group G acts rigidly on a real pretree T . Recall that if the action is minimal, then T must be short and the pretree completion \widehat{T} is itself real. Given any subset $P \subset \widehat{T}$, we define $T[P]$ to be the union in \widehat{T} of T and the G -orbit of P . Generally, assume a countable group system $\mathcal{G} = \bigsqcup_{i \in \mathcal{I}} G_i$ acts minimally on real pretrees $\mathcal{T}_{\mathcal{I}}$. Then for any subset $P \subset \widehat{T}$, we can similarly define $\mathcal{T}[P]$ and let $T_i[P]$ be the component of $\mathcal{T}[P]$ indexed by $i \in \mathcal{I}$. We now make a simple but crucial observation:

Lemma 4.2. *Let \mathcal{G} be a countable group system that admits a minimal rigid action on real pretrees \mathcal{T} . If a subset $P \subset \widehat{T}$ contains no points fixed by a \mathcal{T} -loxodromic element of \mathcal{G} , then the induced \mathcal{G} -action on the real pretrees $\mathcal{T}[P]$ is rigid.*

Additionally, if the rigid action on \mathcal{T} has finite arc stabilizers, then so does the rigid action on $\mathcal{T}[P]$.

Sketch of proof. $\mathcal{T}[P]$ -loxodromics are precisely \mathcal{T} -loxodromics as no point in P is fixed by a \mathcal{T} -loxodromic element of \mathcal{G} . Also, observe that for all elements g in \mathcal{G} , the $\mathcal{T}[P]$ -directions at $\text{Fix}_{\mathcal{T}[P]}(g)$ are in one-to-one equivariant bijection with the \mathcal{T} -directions at $\text{Fix}_{\mathcal{T}}(g)$. So the induced action on $\mathcal{T}[P]$ is rigid. Finally, any arc of $\mathcal{T}[P]$ contains an arc of \mathcal{T} ; if the latter have finite stabilizers, then the rigid action on $\mathcal{T}[P]$ has finite arc stabilizers. \square

For a given countable group H and rigid H -action on a real pretree T with finite arc stabilizers, fix a set of orbit representatives $\mathcal{V} \subset T$ for branch points with infinite stabilizers \mathcal{G} . By countability of H and finiteness of arc stabilizers, the set \mathcal{V} is countable. For each $v \in \mathcal{V}$, consider the G_v -action on the directions at v and fix a set of orbit representatives \mathcal{D}_v for these directions. For a direction d at v , we let $\text{Stab}_{G_v}(d)$ denote the (setwise) stabilizer of d in G_v .

Proposition 4.3. *Suppose a real pretree T has a rigid H -action with finite arc stabilizers, \mathcal{G} represents the infinite T -point stabilizers, and there is a minimal rigid \mathcal{G} -action on real pretrees \mathcal{T}_V .*

If $\vec{p} = (p_d \in \widehat{T}_v : v \in \mathcal{V}, d \in \mathcal{D}_v)$ is a choice of points and each point p_d is fixed by the direction stabilizer $\text{Stab}_{G_v}(d)$, then there is

1. *an H -action on a real pretree $T^* = T^*(\vec{p})$ by pretree-automorphisms;*
2. *equivariantly collapsing the copies of $\mathcal{T}_V[\vec{p}]$ in T^* recovers the real pretree T ;*

If no attaching point p_d in \vec{p} is fixed by a \mathcal{T}_V -loxodromic element in \mathcal{G} , then

3. *the H -action on T^* is rigid;*
4. *the characteristic convex subset of T^* for \mathcal{G} is (an equivariant copy of) \mathcal{T}_V ; and*
5. *an element in H is T^* -loxodromic if and only if it is T - or \mathcal{T}_V -loxodromic;*

moreover, if the \mathcal{G} -action on \mathcal{T}_V has finite arc stabilizers, then so does the H -action on T^ . If the rigid H -action on T is minimal, then so is the rigid H -action on T^* .*

For the sake of compartmentalizing the proposition's long proof, we break it into two parts.

Proof of real pretree's existence. Following the discussion preceding the theorem's statement, let T_v ($v \in \mathcal{V}$) be the component of \mathcal{T}_V corresponding to G_v .

Assume $\vec{p} = (p_d \in \widehat{T}_v : v \in \mathcal{V}, d \in \mathcal{D}_v)$ is a choice of points where each point p_d is fixed by the direction stabilizer $\text{Stab}_{G_v}(d)$; these points will be referred to as the *attaching points*. The real pretrees T and $T_v[\vec{p}]$ (for $v \in \mathcal{V}$) have pretree structures denoted $[\cdot, \cdot]$ and $[\cdot, \cdot]_v$, respectively. The equivariant blow-up $T^* = T^*(\vec{p})$ of T with respect to $\mathcal{T}_V[\vec{p}]$ is defined through the quotient

$$\iota: T \setminus (H \cdot \mathcal{V}) \sqcup H \times \mathcal{T}_V[\vec{p}] \longrightarrow T^*,$$

given by these identifications:

- for each $v \in \mathcal{V}$, $p \in T_v[\vec{p}]$, $s \in G_v$, and $x \in H$, identify (xs, p) with $(x, s \cdot p)$.

Just as in the simplicial setting, T^* inherits an H -action – at least by set-bijections. By construction, $T_v[\vec{p}]$ is equivariantly identified with $\iota(G_v \times T_v[\vec{p}])$ for $v \in \mathcal{V}$, and equivariantly collapsing the translates of $\mathcal{T}_V[\vec{p}]$ produces an equivariant surjection $\pi: T^* \rightarrow T$. This almost establishes Condition 2: we need to define a pretree structure that is H -invariant and projects to the pretree structure of T under π .

For a pretree structure, we need a function $[\cdot, \cdot]^*: T^* \times T^* \rightarrow \mathcal{P}(T^*)$. Let $p^*, q^* \in T^*$.

Case 1: $\pi(p^*) = \pi(q^*) \in H \cdot V$. Choose $y \in H$ so that $\pi(p^*) = y \cdot v$ with $v \in \mathcal{V}$. Then $p^* = \iota(y, p)$ and $q^* = \iota(y, q)$ for some $p, q \in T_v[\vec{p}]$. Define

$$[p^*, q^*]^* := \iota(y, [p, q]_v).$$

Let us quickly check that $[p^*, q^*]^*$ is independent of our choice of $y \in H$. If $s \in G_v$, then $\pi(p^*) = ys \cdot v$, $p^* = \iota(ys, p')$, and $q^* = \iota(ys, q')$ where $s \cdot p' = p$ and $s \cdot q' = q$ by ι 's definition. So $\iota(ys, [p', q']_v) = \iota(y, s \cdot [p', q']_v) = \iota(y, [p, q]_v)$.

In this case, we are equivariantly extending the pretree-structure of $\mathcal{T}_V[\vec{p}]$ to its orbit in T^* (i.e., the image $\iota(H \times \mathcal{T}_V[\vec{p}])$).

Case 2: both $\pi(p^*) = p$, $\pi(q^*) = q$ are not in $H \cdot V$. Consider the closed interval $[p, q] \subset T$. Then

$$[p, q] = [p, q] \setminus (H \cdot \mathcal{V}) \cup \bigcup_{v \in \mathcal{V}, k \geq 1} \{y_{v,k} \cdot v\},$$

for some (possibly finite or empty) subset $\{y_{v,k}\}_{v \in \mathcal{V}, k \geq 1}$ of H . For each $v \in \mathcal{V}$ and $k \geq 1$, the interval $[p, y_{v,k} \cdot v]$ is contained in the direction $y_{v,k} s_{v,k_-} \cdot d(v, k_-)$ at $y_{v,k} \cdot v$ for some $s_{v,k_-} \in G_v$ and $d(v, k_-) \in \mathcal{D}_v$. The element s_{v,k_-} is unique up to right-multiplication by $\text{Stab}_{G_v}(d(v, k_-))$. So the point $s_{v,k_-} \cdot p_{d(v, k_-)}$ is independent of the choice of s_{v,k_-} since the attaching point $p_{d(v, k_-)}$ is fixed by $\text{Stab}_{G_v}(d(v, k_-))$. Similarly, $(y_{v,k} \cdot v, q]$ is in $y_{v,k} s_{v,k_+} \cdot d(v, k_+)$ at $y_{v,k} \cdot v$.

Define

$$[p^*, q^*]^* := \iota \left([p, q] \setminus (H \cdot \mathcal{V}) \cup \bigcup_{v \in \mathcal{V}, k \geq 1} \{y_{v,k}\} \times [s_{v,k_-} \cdot p_{d(v, k_-)}, s_{v,k_+} \cdot p_{d(v, k_+)}]_v \right).$$

Again, $[p^*, q^*]^*$ is independent of the choice of $y_{v,k}$. In this case (and the rest), we are equivariantly ‘attaching’ directions d to points p_d to get the pretree $T^*(\vec{p})$.

Case 3: $\pi(p^*) = p \notin H \cdot V, \pi(q^*) \in H \cdot V$ (or vice-versa). Choose an element $y_\omega \in H$ so that $\pi(q^*) = y_\omega \cdot w$ with $w \in \mathcal{V}$. Then $q^* = \iota(y_\omega, q)$ for some $q \in T_w$. Consider the closed interval $[p, y_\omega \cdot w] \subset T$. Then

$$[p, y_\omega \cdot w] = [p, y_\omega \cdot w] \setminus (H \cdot \mathcal{V}) \cup \{y_\omega \cdot w\} \cup \bigcup_{\substack{v \in \mathcal{V}, k \geq 1 \\ y_{v,k} \cdot v \neq y_\omega \cdot w}} \{y_{v,k} \cdot v\}.$$

As in the previous case, for each $v \in \mathcal{V}$ and $k \geq 1$, the interval $[p, y_\omega \cdot w]$ determines directions $y_{v,k} s_{v,k_\pm} \cdot d(v, k_\pm)$ at the interior point $y_{v,k} \cdot v$. Additionally, the interval determines some direction $y_\omega s_\omega \cdot d(\omega)$ at $y_\omega \cdot w$.

Define $[p^*, q^*]^*$ to be the ι -image of

$$[p, y_\omega \cdot w] \setminus (H \cdot \mathcal{V}) \cup (\{y_\omega\} \times [s_\omega \cdot p_{d(\omega)}, q]_w) \cup \bigcup_{\substack{v \in \mathcal{V}, k \geq 1 \\ y_{v,k} \cdot v \neq y_\omega \cdot w}} \{y_{v,k}\} \times [s_{v,k_-} \cdot p_{d(v,k_-)}, s_{v,k_+} \cdot p_{d(v,k_+)}]_v.$$

Case 4: $\pi(p^*), \pi(q^*) \in H \cdot V$ with $\pi(p^*) \neq \pi(q^*)$. Choose $y_\alpha, y_\omega \in H$ so that $\pi(p^*) = y_\alpha \cdot a$ and $\pi(q^*) = y_\omega \cdot w$ with $a, w \in \mathcal{V}$. The rest is just like Case 3 except that this time, $[y_\alpha \cdot a, y_\omega \cdot w]$ determines directions $y_\alpha s_\alpha \cdot d(\alpha)$ and $y_\omega s_\omega \cdot d(\omega)$ at $y_\alpha \cdot a$ and $y_\omega \cdot w$, respectively. We omit the redundant details.

By construction, $\pi([p^*, q^*]^*) = [\pi(p^*), \pi(q^*)]$. From this, we can directly verify that

1. $[\cdot, \cdot]^*$ satisfies the pretree axioms (see Appendix B for details).
2. H acts on T^* by pretree-automorphisms – that is, for all $x \in H$ and $p^*, q^* \in T^*$, we have $[x \cdot p^*, x \cdot q^*]^* = x \cdot [p^*, q^*]^*$ (exercise).

Remember that $(T, [\cdot, \cdot])$ and $(T_v[\vec{p}], [\cdot, \cdot]_v)$ ($v \in \mathcal{V}$) were real pretrees. Let us now check that the pretree $(T^*, [\cdot, \cdot]^*)$ is real (i.e., any closed interval $I^* = [p^*, q^*]^*$ is pretree-isomorphic to a closed interval of \mathbb{R}). By the definition of $[\cdot, \cdot]^*$, I^* is formed by removing countably many points from a closed interval $I \subset T$ and replacing them with closed intervals pretree-isomorphic to closed intervals of $T_v[\vec{p}]$ -components. Note that I is pretree-isomorphic to a closed interval of \mathbb{R} since T is a real pretree. For the same reason, closed intervals of $T_v[\vec{p}]$ -components are pretree-isomorphic to closed intervals of \mathbb{R} . Similar to the Cantor function, we can use the given pretree-isomorphisms (and the axiom of choice) to construct an isomorphism between I^* and a closed interval of \mathbb{R} . We leave the details as an exercise. So T^* is a real pretree. \square

Remark. Although we invoked the axiom of (countable) choice to conclude that the blow-up T^* is real, this can be avoided in the setting we are most interested in: $H = F$ and the action is minimal. In this case, Gaboriau–Lustig’s index theory implies there are only finitely many orbits of branch points and hence only finitely many choices to be made.

Proof that conditions are satisfied.... Now assume none of the attaching points \vec{p} are fixed by a T_v -loxodromic element in \mathcal{G} .

- Condition 3, rigidity of H -action on T^* :

Suppose $x \in H$ has a nonempty fixed-point set $\text{Fix}_{T^*}(x) \subset T^*$. We need to show that x fixes no T^* -direction at $\text{Fix}_{T^*}(x)$.

First, we show that the fixed-point set $\text{Fix}_{T^*}(x)$ is convex. By the equivariance of π and rigidity of the H -action on T , $\text{Fix}_T(x)$ is nonempty and convex. As x acts as a pretree-automorphism, it must also fix the limit points of $\iota(\text{Fix}_T(x) \setminus (H \cdot \mathcal{V}))$ in $\iota(H \times T_v[\vec{p}])$. By Lemma 4.2, the \mathcal{G} -action on $T_v[\vec{p}]$ is rigid and so $\text{Fix}_{T^*}(x)$ is convex: this used the assumption \vec{p} was a choice of points not fixed by T_v -loxodromic elements in \mathcal{G} . Thus, each T^* -direction d at $\text{Fix}_{T^*}(x)$ has a unique attaching point p^* in $\text{Fix}_{T^*}(x)$.

Finally, pick a T^* -direction d at $\text{Fix}_{T^*}(x)$, and let $p^* \in \text{Fix}_{T^*}(x)$ be its attaching point. If $\pi(p^*) \notin H \cdot \mathcal{V}$, then d corresponds to a T -direction at $\text{Fix}_T(x)$ and the latter is not fixed by x as the H -action on T is rigid. If $\pi(p^*) = y \cdot v$ for some $y \in H$ and $v \in \mathcal{V}$, then d corresponds to either a $\iota(y, T_v)$ -direction at $\iota(y, T_v) \cap \text{Fix}_{T^*}(x)$ or a T -direction at $\text{Fix}_T(x)$ (possible when $p^* \in \iota(y, \vec{p})$). Again, in either case, the direction is not fixed due to rigidity of actions. As d was arbitrary, we are done: x fixes no T^* -direction at $\text{Fix}_{T^*}(x)$.

- Condition 4, characteristic convex subset of T^* for \mathcal{G} :

Fix a point $v \in \mathcal{V}$. We need to show that the characteristic convex subset of T^* for G_v is $\iota(G_v \times T_v)$. Since the rigid H -action on T has finite arc stabilizers and the T -elliptic subgroup G_v is infinite,

the characteristic convex subset of T for G_v is the singleton $\{v\}$. Thus, by the equivariance of π , a characteristic convex subset of T^* for G_v is contained in $\pi^{-1}(v) = \iota(G_v \times T_v[\vec{p}])$. But $\pi^{-1}(v)$ and $T_v[\vec{p}]$ are equivariantly pretree-isomorphic (Condition 2). The characteristic convex subset of $T_v[\vec{p}]$ for G_v is T_v since we assumed the G_v -action on T_v is minimal; therefore, the characteristic convex subset of T^* for G_v is $\iota(G_v \times T_v)$, as needed.

- Condition 5, characterizing T^* -loxodromics – or equivalently, T^* -elliptics:

Suppose x is T - and \mathcal{T}_V -elliptic (i.e., after replacing it with a conjugate if necessary, $x \in G_v$ for some $v \in \mathcal{V}$ and is also T_v -elliptic). So $\iota(x, \text{Fix}_{T_v}(x)) \subset \text{Fix}_{T^*}(x)$ and x is T^* -elliptic.

Conversely, suppose x is T - or \mathcal{T}_V -loxodromic. By the equivariance of π , $\pi(\text{Fix}_{T^*}(x))$ is contained in $\text{Fix}_T(x)$. If x is T -loxodromic (i.e., $\text{Fix}_T(x) = \emptyset$), then it is also T^* -loxodromic. Suppose x is T -elliptic but \mathcal{T}_V -loxodromic. Then $\text{Fix}_T(x)$ is a singleton since the rigid H -action on T has finite arc stabilizers. After replacing it with a conjugate if necessary, assume $x \in G_v$ and is T_v -loxodromic for some $v \in \mathcal{V}$. As $\text{Fix}_{T^*}(x)$ is contained in $\pi^{-1}(v)$ and the pretrees $\pi^{-1}(v)$ and $T_v[\vec{p}]$ are equivariantly identified, we get $\text{Fix}_{T^*}(x) = \iota(x, \text{Fix}_{T_v[\vec{p}]}(x))$. But x is $T_v[\vec{p}]$ -loxodromic since it is T_v -loxodromic, and it fixes none of the attaching points \vec{p} by assumption; therefore, x is T^* -loxodromic.

- Finiteness of arc stabilizers:

Assume the rigid \mathcal{G} -action on \mathcal{T}_V has finite arc stabilizers. Let $A^* \subset T^*$ be an arc; then $A := \pi(A^*) \subset T$ is either a singleton or an arc. If A is an arc, then, by the equivariance of π and rigidity of actions (Condition 3), the stabilizer of A^* is the stabilizer of A ; so it is finite. If $A = \{p\}$ is a singleton, then $p = y \cdot v$ for some $y \in H$ and $v \in \mathcal{V}$. Thus, $A^* = \iota(y, A_v)$ for some arc $A_v \subset T_v[\vec{p}]$ and its stabilizer in H is conjugate to the stabilizer of A_v in G_v ; therefore, it is finite by Lemma 4.2.

- Minimality:

Assume the rigid H -action on T is minimal. By the equivariance of π , any characteristic convex subset of T^* for H must map onto T under the projection π . In particular, any characteristic convex subset will contain $\iota(T \setminus (H \cdot \mathcal{V}))$ and the attaching points $\iota(H \times \vec{p})$. By Condition 4, any characteristic convex subset will contain $\iota(H \times \mathcal{T}_V)$; therefore, T^* is the characteristic convex subset for H (i.e., the H -action on T^* is minimal). \square

4.3. Ideal stitching

The crucial step in proving Theorem 4.1 was using the contraction mapping theorem to find the right blow-up that induced a homothety. We use the same idea for the next theorem, which is the heart of the paper.

For a minimal rigid F -action on a real pretree with trivial arc stabilizers, let \mathcal{V} be a set of orbit representatives of branch points with nontrivial stabilizers G_v . The point stabilizer system $\mathcal{G} = \bigsqcup_{v \in \mathcal{V}} G_v$ has finite type by Gaboriau–Levitt’s index inequality.

Theorem 4.4. *Suppose a real pretree T has a minimal rigid F -action with trivial arc stabilizers, and let $\phi: F \rightarrow F$ be an automorphism.*

Assume $f: T \rightarrow T$ is a ϕ -equivariant pretree-automorphism and \mathcal{G} represents the nontrivial T -point stabilizers with

- a minimal isometric \mathcal{G} -action on a forest \mathcal{Y}_V with trivial arc stabilizers; and
- a $\phi|_{\mathcal{G}}$ -equivariant expanding homothety $h_V: \mathcal{Y}_V \rightarrow \mathcal{Y}_V$.

Then there is:

1. a minimal rigid F -action on a real pretree T^* with trivial arc stabilizers, where
2. the characteristic convex subset of T^* for \mathcal{G} is (an equivariant copy of) \mathcal{Y}_V ;
3. an element in F is T^* -loxodromic if and only if it is T - or \mathcal{Y}_V -loxodromic; and
4. $f^*: T^* \rightarrow T^*$ is the ϕ -equivariant pretree-automorphism induced by f and h_V – the restriction of f^* to \mathcal{Y}_V is h_V and equivariantly collapsing \mathcal{Y}_V recovers f .

Any pair (T', f') satisfying the same conclusion admits an equivariant pretree-isomorphism $T' \rightarrow T^*$ that conjugates f' to f^* ; moreover, if f is F -expanding, then so is f^* .

Proof. For each $v \in \mathcal{V}$, fix a set of F -orbit representatives \mathcal{D}_v for T -directions at v . The set \mathcal{D}_v is finite as well, again by the index inequality. Since the ϕ -equivariant pretree-automorphism f permutes the F -orbits of T -branch points and T -directions, it induces permutations $\beta \in \text{Sym}(\mathcal{V})$ and $\partial \in \text{Sym}(\bigcup_{v \in \mathcal{V}} \mathcal{D}_v)$ with $\partial(d) \in \mathcal{D}_{\beta \cdot v}$ if $d \in \mathcal{D}_v$.

By viewing \mathcal{Y}_V as real pretrees, the minimal isometric \mathcal{G} -action is also a minimal rigid \mathcal{G} -action. For any choice $\vec{p} = (p_d \in \widehat{Y}_v : v \in \mathcal{V}, d \in \mathcal{D}_v)$ of points not fixed by a \mathcal{Y}_V -loxodromic element in \mathcal{G} , we can apply Proposition 4.3 to construct a real pretree $T^* = T^*(\vec{p})$ that satisfies Conditions 1–3. To find the choice \vec{p} that implies the remaining Condition 4, we must use the pretree-automorphism $f: T \rightarrow T$ and expanding homothety $h_V: \mathcal{Y}_V \rightarrow \mathcal{Y}_V$.

The maps f and h_V induce a ϕ -equivariance set-bijection $f^*: T^* \rightarrow T^*$:

- o for $p \in T \setminus (F \cdot \mathcal{V})$, define $f^*(\iota(p)) := \iota(f(p))$; and
- o for $(x, p) \in F \times \mathcal{Y}_V[\vec{p}]$, define $f^*(\iota(x, p)) := \iota(\psi(x)x_v, \hat{h}_v(p))$, where the implicit choice $(x_v \in F : v \in \mathcal{V})$ was used to define the restriction $\phi|_{\mathcal{G}}$, and hence h_V and its extension $\hat{h}_V: \mathcal{Y}_V[\vec{p}] \rightarrow \mathcal{Y}_V[\vec{p}]$.

For each representative $v \in \mathcal{V}$ and direction $d \in \mathcal{D}_v$, let $s_{\partial \cdot d} \in G_{\beta \cdot v}$ be an element such that $f(d) = x_v s_{\partial \cdot d} \cdot \partial(d)$; the element $s_{\partial \cdot d}$ is unique (trivial arc stabilizers). The set-bijection f^* is a pretree-automorphism if and only if $\hat{h}_v(p_d) = s_{\partial \cdot d} \cdot p_{\partial \cdot d}$ for all $v \in \mathcal{V}$ and $d \in \mathcal{D}_v$ (see Appendix B for details).

As $\bigcup_{v \in \mathcal{V}} \mathcal{D}_v$ is finite and h_V is an expanding homothety, the system of equations has a unique solution $\vec{p} = (p_d \in \widehat{Y}_v : v \in \mathcal{V}, d \in \mathcal{D}_v)$ in the metric completion of \mathcal{Y}_V ; these points are not fixed by any \mathcal{Y}_V -loxodromic element in \mathcal{G} . Let $T^* = T^*(\vec{p})$ be the real pretree given by this solution. Then the induced function f^* is a ψ -equivariant pretree-automorphism, as required for Condition 4. Any pair (T', f') satisfying the same conclusion arises from the same construction, and uniqueness follows from uniqueness of \vec{p} .

Finally, assume f is F -expanding. For all $x \in F$ and $n \geq 1$, we need to show the composition $\gamma^* := x \circ (f^*)^n$ is either 1) elliptic with a unique fixed point and it expands at this fixed point or 2) loxodromic with an axis that is not shared with any T^* -loxodromic element in F .

Let $\gamma := x \circ f^n: T \rightarrow T$; then we can use the F -expanding assumption on f .

For the first case, assume γ is loxodromic with an axis that is not shared with any T -loxodromic element in F . Then by construction of T^* and Condition 3, γ^* is loxodromic with an axis that is not shared with any T^* -loxodromic element in F . For the rest of the proof, we may assume γ is elliptic with a unique fixed point $p \in T$ and it expands at p .

For the second case, suppose p has a trivial stabilizer. Then γ^* is elliptic with a unique fixed point $\iota(p) \in T^*$, and it expands at $\iota(p)$.

For the final case, suppose p has a nontrivial stabilizer. Then $p = y \cdot v$ for some $y \in F$ and $v \in \mathcal{V}$. In particular, v is fixed by $\beta^n \in \text{Sym}(\mathcal{V})$ since $p = y \cdot v$ is fixed by $\gamma = x \circ f^n$. In fact, $y \cdot v = x \circ f^n(y \cdot v) = x\psi^n(y)\psi^{n-1}(x_v) \cdots x_{\beta^{n-1} \cdot v} \cdot v$ implies

$$s_v := y^{-1}x\psi^n(y)\psi^{n-1}(x_v) \cdots x_{\beta^{n-1} \cdot v} \text{ is in } G_v,$$

and $\gamma^*(\iota(y, q)) = (y, s_v \cdot \hat{h}_{\beta^{n-1} \cdot v} \circ \cdots \circ \hat{h}_v(q))$ for all $q \in Y_v[\vec{p}]$.

Let $\gamma_V := s_v \circ h_{\beta^{n-1} \cdot v} \circ \cdots \circ h_v$. As h_V is an expanding homothety, the extension of the composition γ_V to the metric completion has a unique fixed point $q_v \in \widehat{Y}_v$ that is repelling. If $q_v \in Y_v[\vec{p}]$, then γ^* is elliptic with a unique fixed point $\iota(y, q_v)$ and it expands at $\iota(y, q_v)$. Otherwise, $q_v \notin Y_v[\vec{p}]$ and γ^* is loxodromic. As γ is elliptic, γ^* does not share its axis (in T^*) with any T -loxodromic element in F . As $\iota(y, q_v)$ is an end of the axis for γ^* and q_v is in the metric completion \widehat{Y}_v , γ^* cannot share its axis with a conjugate (in F) of a Y_v -loxodromic element of G_v . By Condition 3, γ^* does not share an axis with any T^* -loxodromic element in F , and we are done. \square

Epilogue – are all limits the same?

As a preview for the sequel [19], we end the paper discussing whether the real pretree of Theorem 3.3 is *canonical*. There are two ways ‘canonical’ can be interpreted:

1. the real pretree produced by the proof of the theorem does not depend on any choices made in the proof – in a sense, the real pretree was *predetermined*;
2. limit pretrees (i.e., pretrees satisfying the conclusion of the theorem) are equivariantly pretree-isomorphic.

A priori, the second requirement seems stronger as, presumably, there might be limit pretrees that do not arise from the blow-up construction of the theorem. While the jury is still out on the second interpretation, we prove the first one to be true in the sequel!

A careful rereading of Theorem 3.3’s proof reveals that Theorem 1.1 is the only source of indeterminacy: both blow-up constructions in Theorems 4.1 and 4.4 use uniquely determined attaching points! In other words, once the hierarchy of expanding forests (produced by Proposition 3.2) is fixed, the real pretree given at the end of the proof is already determined. Unfortunately, the irreducible train tracks (produced by Theorem 1.1) used in Proposition 3.2 do depend on choices; in general, there are no ‘canonical’ train tracks and our main motivation was to address this issue.

To this end, we also suspect that Theorem 3.3 can be strengthened by requiring the real pretree’s completion to admit an equivariant surjection from the Cantor set boundary of the free group ∂F that is continuous with respect to the observers’ topology. First and foremost, this would immediately allow us to upgrade the rigid action on the real pretree to a convergence action on the completion. Secondly, this would imply the completion of the limit pretree is dual to certain (relative) laminations in the free group (rel. the maximal elliptic subgroup system). Using the dynamics of the automorphism acting on ∂F , we might be able to show that these laminations, and hence their dual pretrees, are unique. Alternatively, Bestvina–Feighn–Handel defined *topmost attracting laminations* [2, Section 6] that can be used to prove a canonical version of Proposition 3.2, thereby making the limit pretrees of Theorem 3.3 canonical – this is the approach taken in the sequel.

Fortunately, there is already precedent for these approaches. When the free group outer automorphism is *irreducible* with infinite order, Bestvina–Feighn–Handel proved that a limit tree produced by the proof of Proposition 3.2 is unique up to equivariant homothety [1, Lemma 3.4]. In particular, the limit tree is independent of the chosen train track used to define it. In fact, the limit trees are canonical in a stronger sense: all trees satisfying the conclusion of Proposition 3.2 are equivariantly homothetic. Importantly, they do not assume beforehand that these trees are projective limits of iterating a topological representative. Later, Levitt–Lustig proved the corresponding observers’ compactifications of these trees admit canonical equivariant quotient maps from the boundary ∂F [15].

For an irreducible atoroidal automorphism, Kapovich–Lustig [10, Proposition 4.5] recently showed that the *dynamical* lamination dual to the limit tree is exactly the *geometric* lamination defined by Mahan Mj (formerly Mitra) [17, p. 388] using the hyperbolic geometry of the corresponding mapping torus $F \rtimes \mathbb{Z}$.

When we drop irreducibility but keep the atoroidal assumption, there still are promising results. Following works of Mj and Bowditch, Elizabeth Field has recently used the hyperbolicity of the mapping torus to construct a canonical dendrite that is dual to Mj’s geometric lamination [7]. We note that the index of the dendrite’s interior is an invariant of the atoroidal outer automorphism. It might follow from the construction that Field’s dendrite is a completion of a limit pretree. On the other hand, Uyanik has proven that atoroidal automorphisms act with generalized north-south dynamics on the projective space of currents [22, Theorem 1.4]. As an extension of Kapovich–Lustig’s result, the geometric lamination should be the support of the repelling simplex of projective currents!

Finally, when we drop the atoroidal assumption, it seems reasonable to conjecture that there are generalizations of Field’s and Uyanik’s results that use relative hyperbolicity and relative currents, respectively, and these can be unified.

A. Index theory for small rigid actions

In this appendix, we sketch Gaboriau–Levitt’s proof of the index inequality [8]. Although the inequality was originally stated for isometric actions, their proof only used the rigidity of actions – the metric plays no other important role. Thus, we sketch this general version of the index theory.

A rigid F -action on a real pretree T is *small* if arc stabilizers are cyclic. Fix a minimal small rigid F -action on a real pretree T . For each orbit of points $[p] \in F \setminus T$, let $G_p \leq F$ be the stabilizer for $p \in T$ and $\#\text{dir}_1[p]$ denote the number of G_p -orbits of directions at p with trivial stabilizers. Recall that the complexity of G_p is $c(G_p) = 2 \cdot \text{rank}(G_p) - 1$. The index at $[p]$ is

$$i[p] := c(G_p) - 1 + \#\text{dir}_1[p] \in \mathbb{Z}_{\geq 0} \cup \{\infty\}.$$

Nonnegativity follows from minimality of the action (F is not trivial). The index is

$$i(F \setminus T) := \sum_{[p] \in F \setminus T} i[p] \in \mathbb{Z}_{\geq 0} \cup \{\infty\}.$$

Gaboriau–Levitt proved the following index inequality.

Theorem A.7 (cf. [8, Theorem III.2]) $i(F \setminus T) < c(F)$.

• The real pretree associated to a rigid system

A real pretree is finite if it is the *convex hull* of a finite subset. A function $j: T \rightarrow T'$ of real pretrees is a morphism if every closed T -interval is a finite union of closed subintervals that are pretree-embedded into T' by j .

Let K be a finite real pretree and $x_i: A_i \rightarrow B_i$ ($1 \leq i \leq n$) be pretree-isomorphisms of finite real pretrees in K . Alternatively, $x_i: K \rightarrow K$ are *partial* pretree-automorphisms whose domains are convex and complete. Denote this system by \mathcal{K} . A system $\mathcal{K} = (K, \{x_i\})$ is nontrivial if no domain $A_i := \text{dom}(x_i)$ is empty.

Let $F_{\mathcal{K}}$ be free group generated by the set $\{x_i\}_{i=1}^n$. Note that any element $x \in F_{\mathcal{K}}$ (i.e., reduced word in $x_i^{\pm 1}$) determines a partial pretree-automorphism $x: K \rightarrow K$ with a possibly empty domain. The length of $x \in F_{\mathcal{K}}$ as a reduced word in $x_i^{\pm 1}$ is denoted $|x|_{\mathcal{K}}$.

Theorem A.1 (cf. [8, Theorem I.1]). *For a nontrivial system $\mathcal{K} = (K, \{x_i\})$, there is*

1. *an $F_{\mathcal{K}}$ -action on a real pretree $T_{\mathcal{K}}$ by pretree-automorphisms;*
2. *a pretree-embedding $\iota: K \rightarrow T_{\mathcal{K}}$ with $x_i \cdot \iota(a) = \iota(x_i(a))$ for $a \in A_i$ ($1 \leq i \leq n$); and*
3. *the $F_{\mathcal{K}}$ -orbit of $\iota(K)$ covers $T_{\mathcal{K}}$, i.e. $T_{\mathcal{K}} = F_{\mathcal{K}} \cdot \iota(K)$.*

If another pretree $F_{\mathcal{K}}$ -action on a real pretree T' satisfies Condition 2 with $\iota': K \rightarrow T'$, then there is a unique equivariant morphism $j: T_{\mathcal{K}} \rightarrow T'$ with $j(\iota(a)) = \iota'(a)$ for $a \in K$.

The last part states that the pair of the pretree $T_{\mathcal{K}}$ and $F_{\mathcal{K}}$ -action is universal with respect to Condition 2. The real pretree $T_{\mathcal{K}}$ (along with its $F_{\mathcal{K}}$ -action and pretree-embedding $\iota: K \rightarrow T_{\mathcal{K}}$) will be referred to as the system’s associated real pretree.

Sketch of proof. The set $T_{\mathcal{K}}$ is defined through the quotient

$$\pi: F_{\mathcal{K}} \times K \rightarrow T_{\mathcal{K}},$$

given by the following identifications:

- for each $a \in A_i$ and $x \in F_{\mathcal{K}}$, (xx_i, a) is identified with $(x, x_i(a))$.

$T_{\mathcal{K}}$ inherits a set $F_{\mathcal{K}}$ -action from the left-multiplication on the first factor of $F_{\mathcal{K}} \times K$. We need to define a pretree structure on $T_{\mathcal{K}}$.

Let the length of $(x, a) \in F_{\mathcal{K}} \times K$ be $|x|_{\mathcal{K}}$. Every point $p \in T_{\mathcal{K}}$ has a unique shortest element $\delta(p) := (x, a) \in \pi^{-1}(p)$. Suppose $p, q \in T_{\mathcal{K}}$ have representatives $\delta(p) = (x, a)$ and $\delta(q) = (y, b)$.

Case 1: $x = y$. Define $[p, q] := \pi(x, [a, b]_K)$.

Case 2: $|y^{-1}x|_{\mathcal{K}} = m \geq 1$. Let $y^{-1}x \in F_{\mathcal{K}}$ be the reduced word $y_m \cdots y_1$ in $x_i^{\pm 1}$. Set c_1 to be the *projection* of a to the domain $\text{dom}(y_1) \neq \emptyset$. Assume $c_j \in \text{dom}(y_j)$ for $1 \leq j < m$, and set c_{j+1} to be the projection of $y_j(c_j)$ to $\text{dom}(y_{j+1}) \neq \emptyset$. Define

$$[p, q] := \pi(x, [a, c_1]_K) \cup \left(\bigcup_{j=1}^{m-1} \pi(xy_1^{-1} \cdots y_j^{-1}, [y_j(c_j), c_{j+1}]_K) \right) \cup \pi(y, [y_m(c_m), b]_K).$$

We leave it as an exercise to check that $(T_{\mathcal{K}}, [\cdot, \cdot])$ is a real pretree and $F_{\mathcal{K}}$ acts on $T_{\mathcal{K}}$ by pretree-automorphisms. This concludes Theorem A.1(1).

Let $1 \in F_{\mathcal{K}}$ be the identity element. The map $\iota: K \rightarrow T_{\mathcal{K}}$ given by $a \mapsto \pi(1, a)$ is a pretree-embedding (see Lemma A.2 below). If $a \in A_i$, then by construction,

$$x_i \cdot \iota(a) = \pi(x_i, a) = \pi(1, x_i(a)) = \iota(x_i(a)).$$

Also by construction, $F_{\mathcal{K}} \cdot \iota(K) = \pi(F_{\mathcal{K}} \times K) = T_{\mathcal{K}}$. This gives us Theorem A.1(2–3).

Finally, we also leave the proof of the universal property as an exercise: use the fact that any closed interval of $T_{\mathcal{K}}$ decomposes (by construction) into finitely many closed subintervals, each contained in a translate of $\iota(K)$. \square

Lemma A.2 (cf. [8, Proposition I.4]). *Let \mathcal{K} be a nontrivial system and $T_{\mathcal{K}}$ be the associated real pretree. For $y \in F_{\mathcal{K}}$ and $a, b \in K$,*

$$y \cdot \iota(a) = \iota(b) \text{ if and only if } y(a) = b.$$

Sketch of proof. The identifications on $F_{\mathcal{K}} \times K$ used to define $T_{\mathcal{K}}$ are symmetric but not transitive. Since the reduced word in $x_i^{\pm 1}$ representing an element of $F_{\mathcal{K}}$ is unique, the identifications generate the following equivalence relation: for $a, b \in K$ and $x, y \in F_{\mathcal{K}}$, $(xy, a) \sim (x, b)$ if and only if $y(a) = b$. \square

Corollary A.3 (cf. [8, Proposition I.5]). *Let \mathcal{K} be a nontrivial system and $T_{\mathcal{K}}$ be the associated real pretree. For any nontrivial $y \in F_{\mathcal{K}}$ and any morphism $j: T_{\mathcal{K}} \rightarrow T'$ as in the last part of Theorem A.1, the restriction of j to the fixed-point set $\text{Fix}_{T_{\mathcal{K}}}(y)$ is injective.*

Sketch of proof. As j is equivariant, ι' is injective, and $j(\iota(a)) = \iota'(a)$ for $a \in K$, it is enough to show that $\text{Fix}_{T_{\mathcal{K}}}(y)$ is contained in some translate of $\iota(K)$. Without loss of generality, we assume $y \in F_{\mathcal{K}}$ is a cyclically reduced nonempty word in $x_i^{\pm 1}$ and show $\text{Fix}_{T_{\mathcal{K}}}(y)$ is contained in $\iota(K)$ – or rather, we prove the contrapositive.

Let $p \in T_{\mathcal{K}}$ be fixed by a nontrivial $y \in F_{\mathcal{K}}$, and assume $p \notin \iota(K)$ (i.e., $\delta(p) = (x, a)$ has positive length). Since (yx, a) and (x, a) represent p , we get $x^{-1}yx(a) = a$ by Lemma A.2. Let $z := x^{-1}yx \in F_{\mathcal{K}}$ be the reduced nonempty word $z_l \cdots z_1$ in $x_i^{\pm 1}$. Then $(xz_l^{-1}, z_1(a))$ and $(xz_l, z_l^{-1}(a))$ also represent p . By minimality of $\delta(p) = (x, a)$, the element $x \in F_{\mathcal{K}}$ as a reduced word in $x_i^{\pm 1}$ cannot end with z_1 nor z_l^{-1} . So $y = xzx^{-1}$ as a reduced word in $x_i^{\pm 1}$ is not cyclically reduced. \square

A partial pretree-automorphism is rigid if either its fixed-point set is empty or it fixes no direction at its fixed-point set. A nontrivial system \mathcal{K} is rigid if every element of $F_{\mathcal{K}}$ determines a rigid partial pretree-automorphism of K .

Lemma A.4. *If \mathcal{K} is a rigid system, then the pretree $F_{\mathcal{K}}$ -action on the associated real pretree $T_{\mathcal{K}}$ is rigid.*

Proof. Suppose $x \in F_{\mathcal{K}}$ has fixed points in $T_{\mathcal{K}}$, and let d be an arbitrary $T_{\mathcal{K}}$ -direction at $\text{Fix}_{T_{\mathcal{K}}}(x)$. Let $p_d \in \text{Fix}_{T_{\mathcal{K}}}(x)$ be an attaching point for d . By construction, some nondegenerate interval $[p_d, q] \subset$

$\{p_d\} \cup d$ is contained in a translate $y \cdot \iota(K)$ for some $y \in F_K$. Let $p_d = y \cdot \iota(a_d)$ and $q = y \cdot \iota(b)$ for some $a_d, b \in K$.

For any $p' = y \cdot \iota(a') \in [p_d, q]$, Lemma A.2 implies $x \cdot p' = p'$ if and only if $y^{-1}xy(a') = a'$. In particular, $a_d \in \text{Fix}_K(y^{-1}xy)$. Since no point in d , and hence $(p_d, q] = y \cdot \iota((a_d, b]_K)$, is fixed by x , no point in $(a_d, b]_K$ is fixed by the partial pretree-automorphism $y^{-1}xy$. So $(a_d, b]_K$ determines a K -direction d_K at $\text{Fix}_K(y^{-1}xy)$.

By rigidity of the system, $y^{-1}xy$ (partially applied to K) does not fix d_K ; therefore, the element $x \in F_K$ (acting on T_K) does not fix d (Lemma A.2), and we are done. \square

Assume the system \mathcal{K} is rigid and T_K is the associated real pretree. Let $S \subset K$ be the finite subset of vertices (i.e., branch points or *endpoints*) in K , A_i and B_i ($1 \leq i \leq n$).

Lemma A.5 (cf. [8, Proposition I.8]). *If $p \in T_K$ is a branch point, then the F_K -orbit of p contains a point of $\iota(S)$, and there are finitely many G_p -orbits of directions at p .*

Sketch of proof. Suppose $[p, q] \cap \iota(K) = \{p\}$, and let $p = \iota(a)$ with $a \in K$. Let $y \in F_K$ be the shortest reduced word $y_m \cdots y_1$ in $x_i^{\pm 1}$ whose translate $y \cdot \iota(K)$ covers a nondegenerate subinterval $[p, p']$. So $[p, p'] = y \cdot \iota([b, b']_K)$ for some $b, b' \in K$ and $y(b) = a$ (Lemma A.2). Minimality of y and $[p, p'] \cap \iota(K) = \{p\}$ imply $[b, b']_K \cap \text{dom}(y_1) = \{b\}$. Then $b \in S$.

Deduce that the F_K -orbit of any branch point in T_K intersects $\iota(S)$. Similarly, the F_K -orbit of any direction at a branch point contains an ‘end’ of $\iota(K \setminus S)$; therefore, the number of G_p -orbits of directions at p is at most the number of ends of $\iota(K) \setminus (\iota(S) \cap (F_K \cdot p))$. \square

• Geometric actions

Fix a basis $\{x_1, \dots, x_n\}$ for the free group F . Let T be a real pretree with a rigid F -action. For any finite real pretree $K \subset T$, we get a system $\mathcal{K} = (K, \{x_i\})$ by letting $x_i \in F$ partially act on K as the partial pretree-automorphism given by restricting the action of x_i on T to the domain $(x_i^{-1} \cdot K) \cap K$. The free groups F_K and F will be identified as they have the ‘same’ basis $\{x_i\}$.

Choose a finite real pretree $K \subset T$ so that the system \mathcal{K} is nontrivial, in which case it will be a rigid system. By Theorem A.1 and Lemma A.4, the rigid system \mathcal{K} has an associated real pretree, denoted T_K , with a rigid F -action. By Lemma A.2, the rigid F -action on T_K is small if the rigid F -action on T was small.

Now suppose the F -action on (a nondegenerate) T is minimal and pick a basepoint $p_0 \in T$. By minimality of the rigid F -action on T , the point p_0 is in the T -axis for some loxodromic element $x \in F$. For any integer m , define K_m to be the convex hull of the partial orbit $\{y \cdot p_0 : |y| \leq m\}$, and let T_m be the real pretree associated to $(K_m, \{x_i\})$. Choose any $m \geq |x|$; then the interval $[x^{-1} \cdot p_0, p_0] \subset K_m$ is contained in the domain of $x: K_m \rightarrow K_m$ (as the composition of partial pretree-automorphisms $x_i: K_m \rightarrow K_m$). Since $x([x^{-1} \cdot p_0, p_0]) = [p_0, x \cdot p_0]$, we deduce x is T_m -loxodromic and its axis in T_m contains $\iota_m(p_0)$. Consequently, the rigid F -action on T_m is minimal.

Using the universal property in Theorem A.1, the chain $K_1 \subset K_2 \subset \dots \subset T$ gives us a direct system of equivariant morphisms $T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T$. The real pretree T is a strong limit of this direct system: each closed interval $I \subset T_1$ has an image under the morphism $T_1 \rightarrow T_m$ (for large enough m) that is embedded into T by the morphism $T_m \rightarrow T$. For the proof, pick a closed interval $I \subset T_1$; its image under the morphism $j_1: T_1 \rightarrow T$ is a finite real pretree, which is contained in K_m for large m ; therefore, $\iota_m(j_1(I)) \subset \iota_m(K_m)$, the image of I under $T_1 \rightarrow T_m$, is embedded into T (with image $j_1(I)$) by $T_m \rightarrow T$ as the latter restricts to a pretree-embedding of $\iota_m(K_m)$.

Lemma A.6 (cf. [8, Proposition II.1]). *Let T be a real pretree with a minimal rigid F -action. The following conditions are equivalent:*

1. *T is equivariantly pretree-isomorphic to T_K for some rigid system $\mathcal{K} = (K, \{x_i\})$; and*
2. *T can only be a strong limit in a trivial way: if T is a strong limit of a direct system of equivariant morphisms $T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T$ with minimal rigid F -actions on T_j , then $T_m \rightarrow T$ is an equivariant pretree-isomorphism for some $m \geq 1$.*

A real pretree T with a minimal rigid F -action is geometric if these conditions hold. By the second characterization, ‘geometricity’ will be independent of the choice of basis for F .

Sketch of proof. (2 \implies 1): By the preceding discussion, T is a strong limit of a direct system of equivariant morphisms $T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T$, and the rigid F -action on T_j is minimal for large enough j . Since we are assuming strong limits are trivial, we have $T_m \rightarrow T$ is an equivariant pretree-isomorphism for some large m . Recall that T_m is associated to a rigid system given by a finite real pretree $K_m \subset T$, and we are done.

(1 \implies 2): Let $\rho: T_{\mathcal{K}} \rightarrow T$ be an equivariant pretree-isomorphism for some rigid system $\mathcal{K} = (K, \{x_i\})$. In particular, T satisfies the universal property in Theorem A.1 with respect to the pretree-embedding $\rho \circ \iota: K \rightarrow T$. We need to show that an arbitrary strong limit $T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T$ of minimal rigid actions is trivial. As a strong limit of minimal actions, we can lift the union of $\rho(\iota(K)) \subset T$ and its translates $x_i \cdot \rho(\iota(K)) \subset T$ ($1 \leq i \leq n$) to a pretree-isomorphic copy in T_m for some large enough m . By the universal property, there is a unique equivariant morphism $T \rightarrow T_m$. So the equivariant morphism $T_m \rightarrow T$ (from the strong limit) must be a pretree-isomorphism, and hence, the strong limit is trivial. \square

• Counting branch points

We are now ready to prove the index inequality.

Theorem A.7 (cf. [8, Theorem III.2]). *For a minimal small rigid F -action on a real pretree T ,*

1. $i(F \setminus T) = c(F) - 1$ if T is geometric, and
2. $i(F \setminus T) < c(F) - 1$ otherwise.

Proof of Part 1. Suppose $\mathcal{K} = (K, \{x_i: A_i \rightarrow B_i\})$ is a rigid system and the rigid F -action on the associated real pretree $T_{\mathcal{K}}$ is small and minimal. For a finite real pretree H , the valence of $h \in H$ – denoted $\nu_H(h)$ – is the number of directions at h . The valences satisfy the identity $\sum_{h \in H} (\nu_H(h) - 2) = -2$.

Fix a point $p \in T_{\mathcal{K}}$. Then $p = x \cdot \iota(a)$ for some $x \in F$ and $a \in K$ by Theorem A.1(3). Define $V_p := \{b \in K : y \cdot \iota(a) = \iota(b) \text{ for some } y \in F\}$. The ‘Cayley graph’ \mathcal{O}_p is the oriented labelled graph with vertex set V_p and an oriented labelled edge $b \xrightarrow{x_i} c$ if $x_i(b) = c$. The Cayley graph is connected (Lemma A.2).

Each oriented labelled edge has a weight $w(b \xrightarrow{x_i} c) := \nu_{A_i}(b)$. The ‘blow-up’ \mathcal{O}'_p is constructed by replacing 1) each vertex $b \in V_p$ with $\nu_K(b)$ -many vertices corresponding to the K -directions at b and 2) each edge $b \xrightarrow{x_i} c$ with $w(b \xrightarrow{x_i} c)$ -many oriented labelled edges in the ‘obvious way’. Let $\pi: \mathcal{O}'_p \rightarrow \mathcal{O}_p$ be the natural finite-to-one projection.

Lemma A.8 (cf. [8, Lemma III.5]). *Let $G_p \leq F$ be the stabilizer of $p \in T_{\mathcal{K}}$.*

1. $\pi_1(\mathcal{O}_p) \cong G_p$ (natural group isomorphism).
2. $\pi_0(\mathcal{O}'_p) \cong G_p \setminus \pi_0(T_{\mathcal{K}} \setminus \{p\})$ (natural set-bijection).
3. for $\mathcal{O}_d \in \pi_0(\mathcal{O}'_p)$ corresponding to a $T_{\mathcal{K}}$ -direction d at p , $\pi_1(\mathcal{O}_d) \cong \text{Stab}_{G_p}(d)$.

Proof sketch for lemma. There is a natural injective homomorphism $\pi_1(\mathcal{O}_p, a) \rightarrow G_{\iota(a)}$ given by ‘reading the oriented labels’. This homomorphism is surjective (Lemma A.2) – as $G_{\iota(a)} = x^{-1}G_p x$, this concludes part 1.

Distinct components of \mathcal{O}'_p correspond to distinct F -orbits of $T_{\mathcal{K}}$ -directions at translates of p by Lemma A.2. So there is a natural injective set-function $\pi_0(\mathcal{O}'_p) \rightarrow G_p \setminus \pi_0(T_{\mathcal{K}} \setminus \{p\})$. The proof of Lemma A.5 shows that this function is surjective – this concludes part 2.

As in the first part, we have an injective homomorphism $\pi_1(\mathcal{O}_d, d_b) \rightarrow \text{Stab}_{G_{\iota(b)}}(d_b)$, where $b \in V_p$; in particular, $G_{\iota(b)} = x^{-1}G_p x$ and $\text{Stab}_{G_{\iota(b)}}(d_b)$ is conjugate in F to a stabilizer in G_p of a $T_{\mathcal{K}}$ -direction at p . This natural homomorphism is surjective (Lemma A.2), and we are done. \square

Recall that $S \subset K$ consists of the vertices of K , A_i and B_i ; it is a finite set. Let $\mathcal{G} \subset \mathcal{O}_p$ be a finite connected subgraph containing all vertices in S and all edges of weight $\neq 2$. Set $\mathcal{G}' := \pi^{-1}(\mathcal{G}) \subset \mathcal{O}'_p$; this is a finite subgraph as well since π was finite-to-one.

By Lemma A.5 and Lemma A.8(2-3), \mathcal{O}'_p has finitely many components, each with rank 0 or 1 (small rigid action). Enlarge \mathcal{G} while maintaining finiteness, and assume \mathcal{G}' supports the fundamental group of each component of \mathcal{O}'_p .

For any finite connected graph, we have the Euler characteristic identity

$$1 - \text{rank}(\pi_1(\cdot)) = \#V(\cdot) - \#E(\cdot),$$

where $V(\cdot)$ and $E(\cdot)$ denote the vertex and edge sets, respectively. Sum over the finitely many components of $\mathcal{G}'_j \subset \mathcal{G}'$; we get

$$\sum_j \left(1 - \text{rank}(\pi_1(\mathcal{G}'_j)) \right) = \sum_{b \in V(\mathcal{G})} v_K(b) - \sum_{e \in E(\mathcal{G})} w(e).$$

Subtract the double of the Euler characteristic identity applied to \mathcal{G} :

$$\begin{aligned} 2 \cdot \text{rank}(\pi_1(\mathcal{G})) - 2 + \sum_j \left(1 - \text{rank}(\pi_1(\mathcal{G}'_j)) \right) &= \sum_{b \in V(\mathcal{G})} (v_K(b) - 2) - \sum_{e \in E(\mathcal{G})} (w(e) - 2) \\ &= \sum_{b \in V_p} (v_K(b) - 2) - \sum_{e \in E(\mathcal{O}_p)} (w(e) - 2). \end{aligned} \quad (\text{A.1})$$

The last equality follows from the fact \mathcal{G} was assumed to contain all vertices and edges with nonzero summand contributions to the right-hand side. Note that the right-hand side is finite and independent of \mathcal{G} . The summands on the left-hand side are nonnegative; therefore, $\pi_1(\mathcal{G})$ has uniformly bounded rank, and $\pi_1(\mathcal{O}_p)$ is finitely generated.

Assume \mathcal{G} supports the fundamental group of \mathcal{O}_p . Then each component of \mathcal{O}'_p has at most one component of \mathcal{G}' . By Lemma A.8, Equation (A.1) can be rewritten as

$$i[p] = \sum_{b \in V_p} (v_K(b) - 2) - \sum_{i=1}^n \sum_{b \in V_p \cap A_i} (v_{A_i}(b) - 2).$$

Summing over all F -orbits $[p] \in F \setminus T_K$ and applying the valence identity in the first paragraph of the proof, we get

$$\begin{aligned} i(F \setminus T_K) &= \sum_{b \in K} (v_K(b) - 2) - \sum_{i=1}^n \sum_{b \in A_i} (v_{A_i}(b) - 2) \\ &= -2 + 2n \\ &= c(F) - 1. \end{aligned}$$

This concludes the geometric case of the theorem. □

Proof of Part 2. Suppose T is a nontrivial strong limit of a direct system of equivariant morphisms $T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T$, where each T_i is geometric.

Fix a branch point $p \in T$, then choose $g_1, \dots, g_r \in G_p$ that freely generate a free factor of G_p and directions d_1, \dots, d_s at p that have trivial stabilizers and are in distinct G_p -orbits. As a strong limit, we can lift the point p to $p' \in T_m$ and the directions d_j ($1 \leq j \leq s$) to d'_j in T_m so that p' is fixed by all g_k ($1 \leq k \leq r$) for large enough $m \geq 1$. By the equivariance of the morphism $T_m \rightarrow T$, the T_m -directions d'_j have trivial stabilizers and $G_{p'} \leq G_p$ contains the free factor generated by g_1, \dots, g_r . So $2r - 2 + s \leq i[p']$. From Part (1), $i[p']$ is uniformly bounded by $i(F \setminus T_m) = c(F) - 1$. Thus, G_p

is finitely generated, there are finitely many G_p -orbits of directions with trivial stabilizers at p , and $i[p] \leq i[p']$.

Similarly, any finite set of branch points $p_1, \dots \in T$ in distinct F -orbits can be lifted to branch points $p'_1, \dots \in T_m$ with $i[p_j] \leq i[p'_j]$ for large enough m . By equivariance, the branch points p'_1, \dots are in distinct F -orbits. So we get $i(F \setminus T) \leq c(F) - 1$ from Part (1).

To show the inequality $i(F \setminus T) < c(F) - 1$, we assume the equality $i(F \setminus T) = c(F) - 1$ and deduce a contradiction. Suppose each T_m is associated to a finite real pretree $K_m \subset T$ that is the convex hull of some partial F -orbit of a branch point $p_0 \in T$.

For large $m \geq 1$, K_m intersects each F -orbit $[p]$ of points in T with $i[p] > 0$, and $p' \in \iota_m(K_m \cap [p]) \subset T_m$ has index $i[p'] = i[p]$. Conversely, as $i(F \setminus T_m) = i(F \setminus T)$, every F -orbit $[p']$ of points in T_m with $i[p'] > 0$ intersects $\iota(K_m \cap [p])$ for some F -orbit of points in T with $i[p] = i[p']$. Fix a large enough $m \geq 1$.

We finally use the hypothesis that T is a nontrivial strong limit of $T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T$ (i.e., the equivariant morphism $j_m: T_m \rightarrow T$ is not injective). Thus, two distinct directions d_1, d_2 at some point $q' \in T_m$ are mapped to the same direction at $q := j_m(q') \in T$.

If $i[q'] > 0$, then $i[q] = i[q']$ implies d_1, d_2 have nontrivial stabilizers. Their image in T also has a nontrivial cyclic stabilizer and so the union of their stabilizers generates a cyclic subgroup; therefore, some nondegenerate arc in $d_1 \cup \{q'\} \cup d_2$ intersecting both directions is fixed by a nontrivial element $y \in F$. But Corollary A.3 states that j_m is injective on the fixed-point set of a nontrivial element – a contradiction.

If q' is a branch point with $i[q'] = 0$, then again, d_1, d_2 have nontrivial stabilizers and the argument is the same. If q' is not a branch point, then some nondegenerate arc in $d_1 \cup \{q'\} \cup d_2$ intersecting both directions is contained some translate $y \cdot \iota_m(K_m)$ – this is where we use the assumption K_m was the convex hull of some branch points. We have another contradiction as j_m is injective on $\iota_m(K_m)$, and we are done. \square

B. Solutions to some exercises

We now state elementary facts about pretrees whose proofs are left to the reader. Let p, q, r be arbitrary points in a pretree $(T, [\cdot, \cdot])$.

- (subinterval) If $q \in [p, r]$, then $[p, q] \subset [p, r]$.
- (arc-overlaps) If $[p, q] \cap [p, r]$ is not a singleton, then it is an arc.
- (pseudocenters) If $m, n \in [p, q] \cap [q, r] \cap [p, r]$, then $m = n$.

Checking pretree axioms are satisfied in Proposition 4.3. Let $p^*, q^*, r^* \in T^*$.

• Symmetric (axiom 1):

Suppose $\pi(p^*) = \pi(q^*)$. If $\pi(p^*) \notin H \cdot \mathcal{V}$, then $p^* = q^*$, and the axiom holds trivially. If $\pi(p^*) \in F \cdot \mathcal{V}$, then $p^* = \iota(y, p)$ and $q^* = \iota(y, q)$ for some $y \in H$, $v \in \mathcal{V}$ and $p, q \in \widehat{T}_v$. The axiom follows from symmetry of $[\cdot, \cdot]_v$.

We may now assume $\pi(p^*) \neq \pi(q^*)$. By definition, the π -preimages of the open interval $(\pi(p^*), \pi(q^*))$ in $[p^*, q^*]^*$ and $[q^*, p^*]^*$ are the same subset of T^* as they have the same *unordered* \mathcal{V} -turns in $[\pi(p^*), \pi(q^*)]$: use symmetry of $[\cdot, \cdot]$ and $[\cdot, \cdot]_v$ ($v \in \mathcal{V}$). Finally, the π -preimages of $\pi(p^*)$ in $[p^*, q^*]^*$ and $[q^*, p^*]^*$ are the same: use symmetry of $[\cdot, \cdot]_v$ if $\pi(p^*) \in H \cdot \mathcal{V}$; use a similar argument for $\pi(q^*)$. So $[\cdot, \cdot]^*$ is symmetric.

Set $P = [\pi(p^*), \pi(r^*)] \cap [\pi(p^*), \pi(q^*)]$ and $R = [\pi(p^*), \pi(r^*)] \cap [\pi(q^*), \pi(r^*)]$. By thinness of $[\cdot, \cdot]$, $[\pi(p^*), \pi(r^*)] = P \cup R$. By arc-overlaps for $[\cdot, \cdot]$, P and R are arcs (if nondegenerate). As $(T, [\cdot, \cdot])$ is real, closed intervals in T are complete and P, R are intervals.

Suppose $M = P \cap R$ is empty; then (without loss of generality) $P = [\pi(p^*), m]$ and $R = [m, \pi(r^*)]$ for some $m \in [\pi(p^*), \pi(r^*)]$. So $[\pi(p^*), \pi(q^*)] \cap [\pi(q^*), \pi(r^*)] = [\pi(q^*), m]$ and $[\pi(p^*), \pi(q^*)] = [\pi(p^*), m] \cup (m, \pi(q^*))$ – absurd – therefore, M is not empty.

In fact, $M = \{m\}$ – pseudocenters for $[\cdot, \cdot]$. Thus, $P = [\pi(p^*), m]$ and $R = [m, \pi(r^*)]$.

- Thin (axiom 2):

Suppose $\pi(p^*) = \pi(r^*)$. If $\pi(p^*) \notin H \cdot \mathcal{V}$, then $p^* = r^*$, and the axiom holds trivially. Otherwise, $[p^*, r^*]^* = \iota(y, [p, r]_v)$. The π -preimages of $\pi(p^*)$ in $[p^*, q^*]^*$ and $[q^*, r^*]^*$ are $\iota(y, [p, q]_v)$ and $\iota(y, [q, r]_v)$ respectively, and the axiom follows from thinness of $[\cdot, \cdot]_v$.

Now assume $\pi(p^*) \neq \pi(r^*)$. By definition of $[\cdot, \cdot]^*$, the π -preimages of $[\pi(p^*), m]$ in $[p^*, r^*]^*$ and $[p^*, q^*]^*$ are the same subset of T^* as they depend only on the \mathcal{V} -turns in P . The same goes for $(m, \pi(q^*))$. Thus,

$$[p^*, r^*]^* \setminus \pi^{-1}(m) \subset ([p^*, q^*]^* \cup [q^*, r^*]^*) \setminus \pi^{-1}(m).$$

To get the desired conclusion, we need only check that

$$[p^*, r^*]^* \cap \pi^{-1}(m) \subset ([p^*, q^*]^* \cup [q^*, r^*]^*) \cap \pi^{-1}(m).$$

If $m \notin (F \cdot \mathcal{V})$, then $\pi^{-1}(m)$ is a singleton, and we are done. Otherwise, the π -preimages of m in $[p^*, r^*]^*$, $[p^*, q^*]^*$ and $[q^*, r^*]^*$ are $\iota(y, [p, r]_v)$, $\iota(y, [p, q]_v)$, and $\iota(y, [q, r]_v)$, respectively. The axiom follows from thinness of $[\cdot, \cdot]_v$.

- Linear (axiom 3):

Suppose $r^* \in [p^*, q^*]^*$ and $q^* \in [p^*, r^*]$. Then $\pi(q^*) = \pi(r^*)$ by linearity of $[\cdot, \cdot]$. If $\pi(q^*) \notin H \cdot \mathcal{V}$, then $q^* = r^*$, and we are done. Otherwise, $q^* = \iota(y, q)$ and $r^* = \iota(y, r)$. The π -preimages of $\pi(q^*)$ in $[p^*, q^*]^*$ and $[p^*, r^*]^*$ are $\iota(y, [p, q]_v)$ and $\iota(y, [p, r]_v)$, respectively, and the axiom follows from linearity of $[\cdot, \cdot]_v$.

Characterizing when f^* is a pretree-automorphism in Theorem 4.4.

Recall that the set-bijection $f^*: T^* \rightarrow T^*$ is a pretree-automorphism exactly when $f^*([p^*, q^*]^*) = [f^*(p^*), f^*(q^*)]^*$ for all $p^*, q^* \in T^*$. Let $\pi: T^* \rightarrow T$ be the collapse map from the construction of T^* . By definition of f^* , $\pi \circ f^* = f \circ \pi$. As π restricts to a pretree-isomorphism $\iota(T \setminus (F \cdot \mathcal{V})) \rightarrow T \setminus (F \cdot \mathcal{V})$, we only need to check that

$$f^*([p^*, q^*]^* \cap \iota(F \times \mathcal{T}_V)) = [f^*(p^*), f^*(q^*)]^* \cap \iota(F \times \mathcal{T}_V).$$

Moreover, if $\pi(p^*) = \pi(q^*) = x \cdot v$ for some $x \in F$ and $v \in \mathcal{V}$ (i.e., $p^* = \iota(x, p)$ and $q^* = \iota(x, q)$), then

$$\begin{aligned} f^*([p^*, q^*]^*) &= f^*(\iota(x, [p, q]_v)) \\ &= \iota(\psi(x)x_v, \hat{h}_v([p, q]_v)) \\ &= [\iota(\psi(x)x_v, \hat{h}_v(p)), \iota(\psi(x)x_v, \hat{h}_v(q))]^* = [f^*(p^*), f^*(q^*)]^*. \end{aligned}$$

The first interesting case is when both $\pi(p^*) = p, \pi(q^*) = q$ are not in $F \cdot \mathcal{V}$ (i.e. $p^* = \iota(p)$ and $q^* = \iota(q)$). We may also assume $p \neq q$ to ignore a degenerate case. Since we are considering the intersection with $\iota(F \times \mathcal{T}_V)$, assume $y \cdot v \in [p, q]$. Let $ys_- \cdot d_-$ and $ys_+ \cdot d_+$ (for some $s_{\pm} \in G_v, d_{\pm} \in \mathcal{D}_v$) be the two directions at $y \cdot v$ determined by the interval $[p, q]$. So

$$\begin{aligned} f^*([p^*, q^*]^* \cap \iota(y, T_v)) &= f^*(\iota(y, [s_- \cdot p_-, s_+ \cdot p_+]_v)) \\ &= \iota(\psi(y)x_v, [\psi_v(s_-) \cdot \hat{h}_v(p_-), \psi_v(s_+) \cdot \hat{h}_v(p_+)]_{\beta \cdot v}), \end{aligned}$$

where p_{\pm} are the attaching points in \vec{p} chosen for direction d_{\pm} , respectively.

Applying f to $y \cdot v \in [p, q]$ gives us $\psi(y)x_v \cdot \beta(v) \in [f(p), f(q)]$; recall that $f(v) = x_v \cdot \beta(v)$ by definition of $\beta \in \text{Sym}(\mathcal{V})$. The two directions at $\psi(y)x_v \cdot \beta(v)$ determined by the interval $[f(p), f(q)]$ are $\psi(ys_{\pm})x_v s_{\partial \cdot d_{\pm}} \cdot \partial(d_{\pm})$ (for some $s_{\partial \cdot d_{\pm}} \in G_{\beta \cdot v}$), respectively, by definition of $\partial \in \text{Sym}(\bigcup_{v \in \mathcal{V}} \mathcal{D}_v)$. Since $\psi(ys_{\pm})x_v = \psi(y)x_v \psi_v(s_{\pm})$ and $\psi_v(s_{\pm}) \in G_{\beta \cdot v}$, we have

$$[f^*(p^*), f^*(q^*)]^* \cap \iota(\psi(y)x_v, T_{\beta \cdot v}) = \iota(\psi(y)x_v, [\psi_v(s_-)s_{\partial \cdot d_-} \cdot q_-, \psi_v(s_+)s_{\partial \cdot d_+} \cdot q_+]_{\beta \cdot v}),$$

where q_{\pm} are the attaching points in \vec{p} chosen for directions $\partial(d_{\pm})$, respectively.

So $f^*([p^*, q^*]^* \cap \iota(y, T_v)) = [f^*(p^*), f^*(q^*)]^* \cap \iota(\psi(y)x_v, T_{\beta \cdot v})$ if and only if

$$\hat{h}_v(p_-) = s_{\partial \cdot d_-} \cdot q_- \quad \text{and} \quad \hat{h}_v(p_+) = s_{\partial \cdot d_+} \cdot q_+.$$

The remaining cases are handled almost exactly the same. Put together, this proves that the set-bijection f^* is a pretree-automorphism if and only if

$$\hat{h}_v(p_d) = s_{\partial \cdot d} \cdot p_{\partial \cdot d} \quad \text{for all } v \in \mathcal{V} \text{ and } d \in \mathcal{D}_v.$$

Acknowledgments. I would like to thank Ursula Hamenstädt for suggestions that helped me clarify the blow-up construction, and the referee, for improving my exposition. I am grateful to the Max Planck Institute for Mathematics for their continued support.

Competing interest. The author has no competing interest to declare.

References

- [1] M. Bestvina, M. Feighn and M. Handel, ‘Laminations, trees, and irreducible automorphisms of free groups’, *Geom. Funct. Anal.* **7**(2) (1997), 215–244.
- [2] M. Bestvina, M. Feighn and M. Handel, ‘The Tits alternative for (F_n) . I. Dynamics of exponentially-growing automorphisms’, *Ann. of Math.* (2) **151**(2) (2000), 517–623.
- [3] M. Bestvina and M. Handel, ‘Train tracks and automorphisms of free groups’, *Ann. of Math.* (2) **135**(1) (1992), 1–51.
- [4] J. S. Birman, A. Lubotzky and J. McCarthy, ‘Abelian and solvable subgroups of the mapping class groups’, *Duke Math. J.* **50**(4) (1983), 1107–1120.
- [5] B. H. Bowditch, ‘Treelike structures arising from continua and convergence groups’, *Mem. Amer. Math. Soc.* **139**(662) (1999), viii+86.
- [6] M. Culler and J. W. Morgan, ‘Group actions on R-trees’, *Proc. London Math. Soc.* (3) **55**(3) (1987), 571–604.
- [7] E. Field, ‘Trees, dendrites and the Cannon-Thurston map’, *Algebr. Geom. Topol.* **20**(6) (2020), 3083–3126.
- [8] D. Gabai and G. Levitt, ‘The rank of actions on R-trees’, *Ann. Sci. École Norm. Sup.* (4) **28**(5) (1995), 549–570.
- [9] D. Gabai, G. Levitt and M. Lustig, ‘A dendrological proof of the Scott conjecture for automorphisms of free groups’, *Proc. Edinburgh Math. Soc.* (2) **41**(2) (1998), 325–332.
- [10] I. Kapovich and M. Lustig, ‘Cannon-Thurston fibers for iwip automorphisms of F_N ’, *J. Lond. Math. Soc.* (2) **91**(1) (2015), 203–224.
- [11] G. Levitt, ‘Graphs of actions on R-trees’, *Comment. Math. Helv.* **69**(1) (1994), 28–38.
- [12] G. Levitt, ‘Non-nesting actions on real trees’, *Bull. London Math. Soc.* **30**(1) (1998), 46–54.
- [13] G. Levitt, ‘Counting growth types of automorphisms of free groups’, *Geom. Funct. Anal.* **19**(4) (2009), 1119–1146.
- [14] G. Levitt and M. Lustig, ‘Periodic ends, growth rates, Hölder dynamics for automorphisms of free groups’, *Comment. Math. Helv.* **75**(3) (2000), 415–429.
- [15] G. Levitt and M. Lustig, ‘Irreducible automorphisms of F_n have north-south dynamics on compactified outer space’, *J. Inst. Math. Jussieu* **2**(1) (2003), 59–72.
- [16] R. A. Lyman, ‘Train track maps on graphs of groups’, *Groups Geom. Dyn.* **16**(4) (2022), 1389–1422.
- [17] M. Mitra, ‘Ending laminations for hyperbolic group extensions’, *Geom. Funct. Anal.* **7**(2) (1997), 379–402.
- [18] J. W. Morgan and P. B. Shalen, ‘Free actions of surface groups on R-trees’, *Topology* **30**(2) (1991), 143–154.
- [19] J. P. Mutanguha, ‘Limit trees for free group automorphisms: universality’, Preprint, 2022, arXiv e-prints.
- [20] J.-P. 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.
- [21] W. P. Thurston, ‘On the geometry and dynamics of diffeomorphisms of surfaces’, *Bull. Amer. Math. Soc. (N.S.)* **19**(2) (1988), 417–431.
- [22] C. Uyanik, ‘Hyperbolic extensions of free groups from atoroidal ping-pong’, *Algebr. Geom. Topol.* **19**(3) (2019), 1385–1411.