Max-Planck-Institut für Mathematik
Bonn

Perron-Frobenius $\mathbb{R}$-trees for automorphisms of free groups

by

Martin Lustig

Max-Planck-Institut für Mathematik
Preprint Series 2020 (5)

Perron-Frobenius $\mathbb{R}$-trees for automorphisms of free groups

Martin Lustig

Max-Planck-Institut für Mathematik
Vivatsgasse 7
53111 Bonn
Germany

Aix Marseille Université
CNRS
Centrale Marseille
I2M UMR 7373
13453 Marseille
France

MPIM 20-5

PERRON-FROBENIUS $\mathbb{R}$-TREES FOR AUTOMORPHISMS OF FREE GROUPS

MARTIN LUSTIG

Abstract. Let $\varphi \in \mathrm{Out}(F_N)$ be an outer automorphism of a free group F_N of finite rank $N \geq 2$, and let $f : \Gamma \to \Gamma$ be an absolute train track representative of φ. Then any non-negative row eigenvector $\vec{v}^{\,*}$ with eigenvalue $\lambda > 1$ of the (non-negative) transition matrix $M(f)$ defines a projectively φ-invariant, expanding $\mathbb{R}$-tree $T^{\vec{v}^{\,*}}$ with isometric F_N-action, called a *Perron-Frobenius* $\mathbb{R}$*-tree*. Such a tree determines a point in the boundary $\partial \mathrm{cv}_N$ of the non-projectivized Outer space cv_N for F_N.

This fact is generalized here from the special case of absolute train track maps to *β-train-track maps*, which exist for every $\varphi \in \mathrm{Out}(F_N)$.

As main result of this paper we show the converse: Every $\mathbb{R}$-tree $T \in \partial \mathrm{cv}_N$, which is projectively fixed and expanded by φ, is a Perron-Frobenius tree: For any β-train-track representative f of φ there exists an eigenvector $\vec{v}^{\,*}$ as above with $T = T^{\vec{v}^{*}}$.

This gives a finite set of strong structural decomposition invariants for any outer automorphisms of F_N.

1. Introduction

In the mid 80's M. Culler and K Vogtmann [8] introduced for any integer $N \geq 2$ a space CV_N, now called *Outer space*, on which the group $\mathrm{Out}(F_N)$ of outer automorphisms of the free group F_N of finite rank $N \geq 2$ acts in a similar vein as the mapping class group acts on Teichmüller space. And just as in that case, the space CV_N has a natural compactification $\overline{\mathrm{CV}}_N$ to which the $\mathrm{Out}(F_N)$-action extends canonically. An element of $\overline{\mathrm{CV}}_N$ is a homothety class $[T]$ of very small $\mathbb{R}$-trees T with isometric F_N-action.

The space $\overline{\mathrm{CV}}_N$ and its $\mathrm{Out}(F_N)$-action (a right action !) have been studied extensively, but both still remain largely mysterious. Even for single automorphisms $\varphi \in \mathrm{Out}(F_N)$ the precise action of φ on $\overline{\mathrm{CV}}_N$ is only known in a rather special case:

If φ is an *iwip automorphism* (= irreducible with irreducible powers, also called *fully irreducible*), it has been shown in [14] that φ acts on $\overline{\mathrm{CV}}_N$ with North-South dynamics, with both fixed points $[T_+], [T_-]$ in the *Thurston boundary* $\partial \mathrm{CV}_N := \overline{\mathrm{CV}}_N \smallsetminus \mathrm{CV}_N$, one of them *expanding*, i.e. $T_+\varphi = \lambda T_+$ with $\lambda > 1$, and the other *contracting* (= expanding for φ^{-1}).

If φ is a Dehn twist automorphism, then by [5] the φ-action on CV_N is parabolic with fixed points all assembled in a particular simplex $\sigma_\varphi \subset \partial \mathrm{CV}_N$. This result has recently been generalized to quadratically growing automorphisms in [23]; however, in both cases the action on $\partial \mathrm{CV}_N$ is more complicated and not totally understood: For example, in general there are more fixed points in $\partial \mathrm{CV}_N$ than just the points in σ_φ.

Generalizing the above iwip case, let us assume that φ can be represented by a graph map $f : \Gamma \to \Gamma$ which has the train track property: for any edge e of Γ and any exponent $t \geq 1$ the path $f^t(e)$ is reduced. Then (see [16, 18]) every row eigenvector $\vec{v}^{\,*}$ with eigenvalue $\lambda > 1$ of the non-negative transition matrix $M(f)$ defines an $\mathbb{R}$-tree $T^{\vec{v}^{\,*}}$. This tree is contained in the boundary $\partial \mathrm{cv}_N$ of the unprojectivized version cv_N of Outer space CV_N, and it is projectively φ-invariant

2000 *Mathematics Subject Classification.* Primary 20F, Secondary 20E, 57M.

Key words and phrases. train track, $\mathbb{R}$-tree, free group automorphism, Perron-Frobenius.

and expanding, just as the tree T_+ above in the iwip case (which corresponds indeed to a subcase of the special case where $M(f)$ is primitive, see Remark 7.9 (2)).

Unfortunately the set of automorphisms which do admit absolute train track representative is not exhaustive, not even among hyperbolic $\varphi \in \mathrm{Out}(F_N)$ (see [20]). This is the reason why in the present paper we work with β-train-track maps, introduced first in [16] and studied further in [19]. In section 4.1 we review their fundamental properties which are used in this paper. β-Train-track maps are a universal tool, as any automorphism of F_N can be represented by such a map (see [19]). They have a wide range of applications, and are a corner stone in the solution of the conjugacy problem in $\mathrm{Out}(F_N)$ lined out in [16].

The following proposition is a streamlined version of Theorem 5.2 below, which explains in particular the precise way how $\vec{v}^*$ determines $T^{\vec{v}^*}$ and the associated translation length function $||\cdot||_{T^{\vec{v}^*}}$ on F_N.

Proposition 1.1. *Consider any automorphism $\varphi \in \mathrm{Out}(F_N)$ and any β-train-track representative $f : \mathcal{G} \to \mathcal{G}$ of φ. Then every non-negative row eigenvector $\vec{v}^*$ with eigenvalue $\lambda > 1$ of the transition matrix $M(f)$ defines an $\mathbb{R}$-tree $T^{\vec{v}^*} \in \partial \mathrm{cv}_N$, which is projectively invariant under φ and expanding:*

$$T^{\vec{v}^*}\varphi = \lambda T^{\vec{v}^*}$$

Since for non-polynomially growing φ the existence of an eigenvector $\vec{v}^*$ as in the above proposition is ensured by standard results from Perron-Frobenius theory of non-negative matrices, we call the tree $T^{\vec{v}^*}$ a *Perron-Frobenius $\mathbb{R}$-tree* (with respect to the β-train-track map f). It follows from our main result that this definition is actually independent from the particular β-train-track map used to construct $T^{\vec{v}^*}$. Indeed, we show (see section 6):

Theorem 1.2. *Let $\varphi \in \mathrm{Out}(F_N)$ be an outer automorphism, and let $T \in \overline{\mathrm{cv}}_N$ be a projectively φ-invariant expanding $\mathbb{R}$-tree.*

Then for any β-train-track map $f : \mathcal{G} \to \mathcal{G}$ which represents φ there exists a non-negative eigenvector $\vec{v}^*$ of $M(f)$, with eigenvalue $\lambda > 1$, such that T is a Perron-Frobenius tree with respect to f:

$$T = T^{\vec{v}^*}$$

Perron-Frobenius theory of non-negative matrices M has been studied intensively; for a quick survey of the general case (i.e. without assuming that M is primitive or irreducible) see Appendix A.3 of [1]. It is well known that any such M possesses (up to rescaling) only a finite number of extremal non-negative eigenvectors, and any other non-negative eigenvector of M is a convex combination of the latter. Hence for our purposes it is enough to consider the finitely many extremal eigenvectors $\vec{v}^*$ of $M(f)$, which can be readily computed from the given map f. This is explained in Remark 7.5 below.

For any eigenvector $\vec{v}^*$ as above we consider in section 7.3 the *zero-subgraph* $\mathcal{G}^{\vec{v}^*} \subset \mathcal{G}$, which consists of all edges where $\vec{v}^*$ has coordinate value 0, and the non-trivial connected components $\mathcal{G}_i$ of $\mathcal{G}^{\vec{v}^*}$. The given marking $\theta : \pi_1\mathcal{G} \to F_N$ defines an (up to conjugation) finite, partially ordered system $U(T^{\vec{v}^*})$ of finitely generated subgroups $\theta(\pi_1\mathcal{G}_i)$ of F_N. Alternatively, the subgroups of the system $U(T^{\vec{v}^*})$ are given by the non-trivial stabilizers of the branch points of $T^{\vec{v}^*}$, see Remark 7.11.

It follows from Theorem 1.2 that these subgroups are up to conjugation in F_N structural invariants of φ, i.e. they do not depend on the choice of the particular β-train-track representative f of φ. Denoting by $\mathcal{U}(\varphi)$ the union of the $U(T^{\vec{v}^*})$, for any eigenvector $\vec{v}^*$ as in Proposition 1.1, we obtain (see Corollary 7.13 for more detail):

Corollary 1.3. *Let $\varphi, \varphi', \psi \in \mathrm{Out}(F_N)$ be automorphism which satisfy $\varphi' = \psi \circ \varphi \circ \psi^{-1}$. Then one has:*

$$\mathcal{U}(\varphi') = \psi(\mathcal{U}(\varphi))$$

Since generators for any subgroup from the system $\mathcal{U}(\varphi)$ can be determined algorithmically from any β-train-track representative $f : \mathcal{G} \to \mathcal{G}$ of φ, the finite system $\mathcal{U}(\varphi)$ constitutes a powerful computable conjugacy invariant for any $\varphi \in \mathrm{Out}(F_N)$. Further structural information about φ derived from $\mathcal{U}(\varphi)$ as consequences of the above stated results are discussed in section 7.

History and Acknowledgements: This paper is an elaboration and extension of what has been presented in sections 1 and 2 of the MPI-preprint [16]. This work has started during the author's 2000-01 stay at the the Max-Planck Institut für Mathematik in Bonn, and was extended further during his renewed visit in the Fall of 2016.

2. Preliminaries

2.1. **Basics on graphs and graph maps.**

Throughout this paper a graph Γ means a topological space that consists of vertices and edges in the usual manner. Unless explicitly specified (for example if the graph in question is a covering space of another graph), we assume tacitly that Γ is finite (i.e. there are only finitely many vertices and edges). We also assume tacitly that Γ is connected, unless otherwise stated.

For every oriented edge e of Γ we denote by $\overline{e}$ the oppositely oriented edge, which gives $e \neq \overline{e}$ and $\overline{\overline{e}} = e$ for any edge e. We randomly pick for any set $\{e, \overline{e}\}$ one of its elements, and assemble all those chosen edges in the subset $\mathrm{Edges}^+(\Gamma)$ of the set $\mathrm{Edges}(\Gamma)$ of all edges of Γ.

An edge path $\gamma = e_1 e_2 \ldots e_s$ is said to have *combinatorial length* $|\gamma| = s$. The path γ is *trivial* if it has combinatorial length $|\gamma| = 0$. We say that γ is *reduced* if $e_{i+1} \neq \overline{e}_i$ for all $i = 1, \ldots, s-1$. If γ is not reduced, then one can perform an elementary reduction, which means erasing a subpath $e_i e_{i+1}$ with $e_{i+1} = \overline{e}_i$. For any edge path γ there exists a well define reduced path, denoted $[\gamma]$, which has the same endpoints as γ and can be obtained from γ by a finite sequence of elementary reductions. A loop $\widehat{\gamma}$ is reduced if the underlying path γ is reduced and cyclically reduced (i.e. $\gamma_1 \neq \overline{\gamma}_s$).

An edge path γ which has coinciding initial and terminal vertex $P \in \Gamma$ and can be reduced to the trivial edge path $[\gamma] = P$ performing iteratively elementary reductions, is called an *backtracking path*. It follows that an edge path γ is reduced if and only if it doesn't contain any non-trivial backtracking subpath.

A map $f : \Gamma \to \Gamma$ is a *graph map* if f sends vertices to vertices and edges to (not necessarily reduced) edge paths. An edge path γ is called *legal* (with respect to f) if all forward f-iterates of γ are reduced: one has

$$[f^t(\gamma)] = f^t(\gamma) \quad \text{for all} \quad t \geq 1\,.$$

The map f is said to have the *train track property* (or *f is a train track map*) if every edge, viewed as edge path of combinatorial length 1, is a legal path.

For any graph map $f : \Gamma \to \Gamma$ let $M(f) = (m_{e',e})_{e',e \in \mathrm{Edges}^+(\Gamma)}$ denote the *(geometric) transition matrix* for f, i.e. the entry $m_{e',e} \geq 0$ denotes the number of times the path $f(e)$ crosses over the edge e', disregarding the orientations.

An edge e of Γ is *polynomially growing* under iteration of a map $f : \Gamma \to \Gamma$ if the combinatorial length of the paths $f^t(e)$ is bounded above by a polynomial in $t \in \mathbb{N}$.

A graph Γ is *marked* if it is provided with an *F_N-marking* (or simply *marking*), i.e. a surjective homomorphism $\theta : \pi_1\Gamma \to F_N$, where F_N denotes the free group of finite rank $N \geq 0$. As is common, we consider two markings as equal if they only differ by an inner automorphism of F_N, which is why we do not need to specify a base point of Γ.

Throughout this paper we denote by $\widehat{\Gamma}$ the covering of Γ which is defined by $\ker\theta$; on this "Galois covering" the group F_N acts freely as group of deck transformations (= covering translations). If θ happens to be an isomorphism, then $\widehat{\Gamma}$ coincides with the universal covering $\widetilde{\Gamma}$ of Γ.

A graph map $f : \Gamma \to \Gamma$ of a marked graph Γ *represents an outer automorphism* $\varphi \in \text{Out}(F_N)$ if the given marking epimorphism θ and the endomorphism f_* induced by f on $\pi_1\Gamma$ satisfy:

$$\theta \circ f_* = \varphi \circ \theta$$

Here f_* is understood as "outer endomorphism", i.e. it is well defined only up to composition with inner automorphisms.

This terminology is sometimes extended to non-marked graphs Γ: In this case we say that a graph map $f : \Gamma \to \Gamma$ represents φ if there exists a marking $\theta : \pi_1\Gamma \twoheadrightarrow F_N$ so that f represents φ with respect to θ.

A lift $\widehat{f} : \widehat{\Gamma} \to \widehat{\Gamma}$ of f *represents an automorphism* $\Phi \in \text{Aut}(F_N)$ if for any element $w \in F_N$ one has:

$$\widehat{f}w = \Phi(w)\widehat{f} : \widehat{\Gamma} \to \widehat{\Gamma}$$

In this case it follows that f represents the automorphism $\varphi \in \text{Out}(F_N)$ which is the image of Φ under the quotient map $\text{Aut}(F_N) \to \text{Out}(F_N)$. Furthermore, any preimage Φ' of φ can be represented by a properly chosen lift $\widehat{f}'$ of f, and $\widehat{f}'$ is uniquely determined by Φ'.

The above introduced terminology will also be used in a more general context, i.e. for any space Z provided with an F_N-action: We say that $h : Z \to Z$ *represents* $\Phi \in \text{Aut}(F_N)$ if one has

$$\Phi(w)h = hw : Z \to Z \quad \text{for any} \quad w \in F_N\,. \tag{2.1}$$

For any second such space Z' and map $h' : Z' \to Z'$ which also represents Φ we say that an F_N-equivariant map $j : Z \to Z'$ *semi-commutes* with h and h' if one has:

$$jh = h'j \tag{2.2}$$

A further generalization of equality (2.1) which is occasionally useful occurs if $\Phi : G_1 \to G_2$ is any group homomorphism, and Z_1 and Z_2 are spaces with actions of G_1 and G_2 respectively. Then any map $h : Z_1 \to Z_2$ is said to *represent* Φ if one has

$$\Phi(g)h = hg : Z_1 \to Z_2 \quad \text{for any} \quad g \in G_1\,. \tag{2.3}$$

2.2. **Perron-Frobenius length functions.**

Definition 2.1. Let Γ be a graph with F_N-marking, and let $f : \Gamma \to \Gamma$ be a graph map that represents an automorphism $\varphi \in \text{Out}(F_N)$.

(1) Any map

$$L^+ : \text{Edges}^+(\Gamma) \to \mathbb{R}_{\geq 0}$$

can be completed to an *edge length function*

$$L : \text{Edges}(\Gamma) \to \mathbb{R}_{\geq 0}$$

by the convention $L(\overline{e}) = L(e)$ for any edge e of Γ.

(2) For any edge path or loop $\gamma = e_1 e_2 \ldots e_q$ in Γ the edge length function L defines the path length

$$L(\gamma) = \sum_{k=1}^{q} L(e_k)\,.$$

By lifting L equivariantly and setting

$$d^L(P, Q) = \inf\{L(\gamma) \mid \gamma \text{ connects } P \text{ to } Q\} \tag{2.4}$$

for any two vertices P and Q of $\widehat{\Gamma}$, one obtains a pseudo-metric d^L on the set of vertices of $\widehat{\Gamma}$, i.e. a non-negative continuous function on $\widehat{\Gamma} \times \widehat{\Gamma}$ that differs from a metric only in that distinct points may well have distance 0.

(3) If need be, we will distribute for any edge e of Γ the length $L(e)$ along e so that distance can be measured between any two points of Γ. More importantly, by lifting this extension of L equivariantly, the above distance function d^L extends to a pseudo-metric on all of $\widehat{\Gamma}$.
(4) For any constant $\lambda > 1$ which satisfies $L([f(e)]) \leq \lambda L(e)$ for any edge e of Γ, we define the *limit length* $L^\lambda_\infty(\gamma)$ to be the following length infimum:

$$L^\lambda_\infty(\gamma) = \inf_{t \in \mathbb{N}} \frac{1}{\lambda^t} L([f^t(\gamma)])$$

If this limit length needs to be distributed along the interior of the edges of Γ, we always use the convention that for any points x and y on e which are mapped to vertices $P = f^t(x)$ and $Q = f^t(y)$ the length of the segment $[x, y]$ on e is given by

$$L^\lambda_\infty([x, y]) := \frac{1}{\lambda^t} L^\lambda_\infty(\gamma')\,, \tag{2.5}$$

where γ' is the subpath of $f^t(e)$ which joins P to Q. It follows that on the closure $\mathcal{P}$ of the set of those points, which are eventually mapped by iterates of f to vertices, the distance function $d^{L^\lambda_\infty}$ from (2.4) is well defined, and through extending it by the constant function to the complement of $\mathcal{P}$, one obtains a well define pseudo-metric $d^{L^\lambda_\infty}$ on $\widehat{\Gamma}$.
(5) An edge length function

$$L : \text{Edges}(\Gamma) \to \mathbb{R}_{\geq 0}$$

is a *Perron-Frobenius length function* (abbreviated to *PF-length*[1]) if there exists a non-negative right eigenvector $\vec{v}^{\,*} = (v_e)_{e \in \text{Edges}^+(\Gamma)}$ of the transition matrix $M(f)$ with eigenvalue $\lambda > 1$, such that for any edge $e \in \text{Edges}^+(\Gamma)$ the coordinate equality

$$v_e = L(e)$$

is satisfied. In this case it is convenient to specify the above notation to $L^{\vec{v}^{\,*}} := L$, $L^{\vec{v}^{\,*}}_\infty := L^\lambda_\infty$ and $d^{\vec{v}^{\,*}}_\infty = d^{L^\lambda_\infty}$.

We observe directly from this definition:

Remark 2.2. Let $L = L^{\vec{v}^{\,*}}$ be a Perron-Frobenius length function on Γ as in Definition 2.1.
(1) For an arbitrary path or loop $\gamma = e_1 e_2 \ldots e_q$ in Γ we define a *length vector*

$$\vec{v}(\gamma) = (|\gamma|_e + |\gamma|_{\overline{e}})_{e \in \text{Edges}^+(\Gamma)}\,,$$

where $|\gamma|_e$ denotes the number of times that γ crosses over e (in forward direction only). Then the length vector of the unreduced path $f^t(\gamma)$ is given by the following matrix product:

$$\vec{v}(f^t(\gamma)) = M(f)^t\, \vec{v}(\gamma)$$

Furthermore, the PF-length $L(\gamma)$ is given as scalar product

$$L(\gamma) = \vec{v}^{\,*} \cdot \vec{v}(\gamma)\,,$$

so that the PF-length of the unreduced path $f^t(\gamma)$ is given by

$$L(f^t(\gamma)) = \vec{v}^{\,*} M(f)^t \vec{v}(\gamma) = \lambda^t L(\gamma)\,. \tag{2.6}$$

(2) Since reduction of paths always diminishes length, the sequence $\frac{1}{\lambda^t} L([f^t(\gamma)])$ is decreasing, so that one has

$$L^{\vec{v}^{\,*}}_\infty(\gamma) = \lim_{t \in \mathbb{N}} \frac{1}{\lambda^t} L([f^t(\gamma)]) \leq L(\gamma)\,. \tag{2.7}$$

[1] For convenience we allow ourselves to abbreviate on and off the two great mathematicians Perron and Frobenius to "PF" and much apologize for this !

In the special case that γ is legal, one has $[f^t(\gamma)] = f^t(\gamma)$ for any $t \geq 0$, and hence

$$L_\infty^{\vec{v}^*}(\gamma) = L(\gamma)\,. \tag{2.8}$$

(3) From equality (2.6) we deduce in particular that for any polynomially growing edge e one has $L(e) = 0$ and thus

$$L_\infty^{\vec{v}^*}(e) = 0\,.$$

(4) If f has the train track property, then there exists always a legal loop $\widehat{\gamma}$ in Γ with $L_\infty^{\vec{v}^*}(\widehat{\gamma}) > 0$. Indeed, since $\vec{v}^*$ is an eigenvector and thus $\vec{v}^* \neq \vec{0}$, it suffices to iterate some edge e with length $L(e) > 0$ until it is long enough to contain a suitable loop as subpath, which happens eventually due to the finiteness of Γ and our assumption $\lambda > 1$.

2.3. **$\mathbb{R}$-trees and Outer space.**

An $\mathbb{R}$-tree T is a non-empty path-connected 0-hyperbolic metric space. Equivalently, T is a non-empty metric space which has the property that for any two points $x, y \in T$ there is a unique isometric embedding of the interval $[0, d(x, y)] \subset \mathbb{R}$ into T such that the image is a segment $[x, y] \subset T$ which connects x to y.

All $\mathbb{R}$-trees in this paper are provided with an action of F_N by isometries (written as left-multiplication). Examples are given by the universal covering $\widetilde{\Gamma}$ of a graph Γ which is provided with a marking isomorphism $\theta : \pi_1\Gamma \to F_N$, and with a positive edge length function L that is lifted to $\widetilde{\Gamma}$ (see Definition 2.1 (3)).

Such *metric simplicial* $\mathbb{R}$-trees $T = \widetilde{\Gamma}$, with free F_N-action by deck transformations, are precisely the points of the *unprojectivized Outer space* cv_N for F_N, provided that T is *minimal*, i.e. there is no non-empty F_N-invariant proper subtree of T. The classical Outer space CV_N is obtained from cv_N by projectivization, i.e. T is replaced by the class $[T]$ of all $\mathbb{R}$-trees T' that are obtained from T by uniform rescaling with a factor $\lambda > 0$:

$$T' \in [T] \quad \Longleftrightarrow \quad T' = \lambda T$$

Every $\mathbb{R}$-tree T defines a *translation length function* $||\cdot||_T$ on F_N given by

$$||w||_T = \inf\{d(x, w \cdot x) \mid x \in T\}\,,$$

which is well known to depend only on the conjugacy class $[w]$ of w in F_N. Two $\mathbb{R}$-trees T_1 and T_2 have the same translation length function if and only if their minimal subtrees are F_N-equivariantly isometric (see [9]).

In particular, it follows that the issuing map

$$\mathrm{cv}_N \to \mathbb{R}_{\geq 0}^{F_N}\,, \;\; T \mapsto ||\cdot||_T$$

is an embedding, which is used to define the topology on cv_N. Every point in closure of the image of cv_N (but not in the image itself) is the translation length function of a well defined minimal $\mathbb{R}$-tree T with isometric F_N-action that is either not free or not simplicial (or neither). Such $\mathbb{R}$-trees T constitute the *Thurston boundary* $\partial\mathrm{cv}_N$ of cv_N, and projectivization gives the compactification $\overline{\mathrm{CV}}_N = \mathrm{CV}_N \cup \partial\mathrm{CV}_N$ of Outer space.

The spaces $\mathrm{cv}_N, \overline{\mathrm{cv}}_N = \mathrm{cv}_N \cup \partial\mathrm{cv}_N, \mathrm{CV}_N$ and $\overline{\mathrm{CV}}_N$ are equipped with a natural right action by $\mathrm{Out}(F_N)$, which is properly discontinuous on CV_N, so that the latter is rightfully considered a proper analogue of Teichmüller space, with $\mathrm{Out}(F_N)$ taking on the role of the mapping class group.

An $\mathbb{R}$-tree $T \in \overline{\mathrm{cv}}_N$ defines a fixed point

$$[T] = [T] \cdot \varphi \in \overline{\mathrm{CV}}_N$$

of $\varphi \in \mathrm{Out}(F_N)$ if $T \cdot \varphi = \lambda T$ for some $\lambda > 0$. Equivalently, one has

$$\Phi(w)H = Hw : T \to T \quad \text{for any} \quad w \in F_N\,, \tag{2.9}$$

for some homothety $H : T \to T$ with *stretching factor* λ and some lift $\Phi \in \text{Aut}(F_N)$ of φ (in which case H represents Φ). We say that $[T]$ is an *expanding* fixed point of φ if the homothety H expands T, i.e. if the stretching factor satisfies $\lambda > 1$. In this case the set of branch points is dense in T, so that T is not simplicial, and hence one has:

$$T \in \partial \text{cv}_N \tag{2.10}$$

Minimal non-simplicial $\mathbb{R}$-trees $T \in \partial \text{cv}_N$ are not complete as metric spaces. We denote by $\overline{T}$ the metric completion of any $\mathbb{R}$-tree T. Any homothety $H : T \to T$ extends canonically to a homothety $\overline{H} : \overline{T} \to \overline{T}$ with the same stretching factor as H. Clearly $\overline{H}$ represents the same $\Phi \in \text{Aut}(F_N)$ as does H.

An important property of $\mathbb{R}$-trees T comes from considering "arcs", i.e. segments $[x, y] \subset T$ which are not degenerated, i.e. with $x \neq y$. By the *stabilizer* in F_N of such an arc we always mean the point-wise stabilizer. The tree T (or the F_N-action on T) is called *small* if the stabilizer of any arc in T has rank ≤ 1, i.e. it is infinite cyclic or trivial. A further restriction, called *very small*, is to impose that the set of fixed points of any $w \in F_N \smallsetminus \{1\}$ does not contain any tripod, and that w doesn't invert any arc. It has been shown in [2, 5] that $\overline{\text{cv}}_N$ is precisely the set of all very small minimal $\mathbb{R}$-trees.

2.4. **Bounded backtracking.**

In [11] the following notion for isometric F_N-actions on $\mathbb{R}$-trees has been introduced and studied:

Definition 2.3. The action of F_N on the $\mathbb{R}$-tree T has *bounded backtracking* (for short: *T satisfies BBT*) if for every $Q \in T$ and every basis A of F_N there exists a constant $C \geq 0$ such that for any reduced words $v, w \in F_N$ with reduced product vw one has $d(vQ, [Q, vwQ]) \leq C$.

In this paper we will use the property BBT explicitly in the following equivalent form:

Lemma 2.4. *(1) The F_N-action on T satisfies BBT if and only if for any graph Γ with marking isomorphism $\theta : \pi_1\Gamma \xrightarrow{\cong} F_N$ and any F_N-equivariant map $i : \widetilde{\Gamma} \to T$ the following property is satisfied:*

- (*) *There is a constant $c \geq 0$ such that for any two vertices P and Q of the universal covering tree $\widetilde{\Gamma}$ the reduced edge path γ from P to Q is mapped by i into the c-neighborhood of the geodesic segment $[i(P), i(Q)] \subset T$.*

(2) The F_N-action on T satisfies BBT if property (*) *holds for some Γ and some i as in (1) above.* □

Remark 2.5. (1) It is easy to see that Lemma 2.4 stays valid if Γ is replaced by a graph Γ' which has as marking only a (not necessarily injective) epimorphism $\theta : \pi_1\Gamma' \twoheadrightarrow F_N$, the universal covering $\widetilde{\Gamma}$ is replaced by the F_N-covering $\widehat{\Gamma}'$, and the reduced path γ is replaced by any quasi-geodesic path γ' from P to Q with a priori fixed quasi-geodesy constants.
(2) If the F_N-action on the $\mathbb{R}$-tree T satisfies BBT, then so does the induced action of F_N on the metric completion $\overline{T}$ of T. This follows directly from Definition 2.3 or Lemma 2.4.

It has been shown in [11] that from a bounded cancellation lemma due to Bestvina-Feighn-Handel (see [3], Lemma 3.1 or [10], Lemma 2.2.4), the following can be derived:

Proposition 2.6. *Every $\mathbb{R}$-tree $T \in \overline{\text{cv}}_N$ satisfies BBT.* □

3. $\mathbb{R}$-trees obtained through train track iterations

In [11] a construction of $\mathbb{R}$-trees through iterations of a given train track map has been introduced, which plays a crucial role in this paper. We review in this section the main ingredients of this construction, in a slightly generalized form.

Let Z be a path-connected topological space with an action of F_N by homeomorphisms. Assume that the space Z is equipped with an F_N-invariant pseudo-metric d. Let $h : Z \to Z$ be a map which represents some automorphism $\Phi \in \mathrm{Aut}(F_N)$ (see equality (2.1)). Let $\lambda > 1$ be a constant which satisfies

$$d(h(x), h(y)) \leq \lambda d(x, y) \tag{3.1}$$

for any points $x, y \in Z$. For any integer $k \geq 0$ we define a new F_N-equivariant pseudo-metric on Z through

$$d_k(x, y) := \frac{d(h^k(x), h^k(y))}{\lambda^k}, \tag{3.2}$$

and from (3.1) we deduce $d_k(x, y) \leq d_{k+1}(x, y)$. Hence the d_k converge to an F_N-equivariant pseudo-metric

$$d_\infty(x, y) = \lim_{k \to \infty} d_k(x, y) \tag{3.3}$$

which satisfies

$$d_\infty(x, y) \leq d(x, y)$$

and

$$d_\infty(h(x), h(y)) = \lambda d_\infty(x, y)$$

for all $x, y \in Z$.

Let T_∞ be the canonical metric quotient space of Z with respect to the pseudo-metric d_∞, equipped with the F_N-action inherited from Z. Then the quotient map

$$i : Z \to T_\infty$$

is F_N-equivariant and continuous, and in particular T_∞ is path-connected. From the construction of the limit space T_∞ and the map i we observe directly that there is a homothety $H : T_\infty \to T_\infty$ with stretching factor λ which represents the same automorphism $\Phi \in \mathrm{Aut}(F_N)$ as does the map h. In particular i semi-commutes with h and H (see equality (2.2)).

Remark 3.1. It may well occur that the limit pseudo-metric d_∞ is the zero pseudo-metric throughout all of Z. This happens in particular if λ is chosen non-minimal with respect to inequality (3.1).

We now assume that the pseudo-metric d on Z is 0-hyperbolic. Since h is distance decreasing, the same follows for any of the intermediate pseudo-metrics d_k. Hence the limit metric d_∞ is also 0-hyperbolic, so that (see [11]) the limit space T_∞ is an $\mathbb{R}$-tree with an action of F_N by isometries and an F_N-equivariant map

$$i : Z \twoheadrightarrow T_\infty \tag{3.4}$$

which is surjective and *preserves the* d_∞*-pseudo-metric* in that

$$d(i(x), i(y)) = d_\infty(x, y)$$

for any $x, y \in Z$. The tree T_∞ is uniquely determined, up to F_N-equivariant isometry, by the space Z, the map h, the pseudo-metric d on Z, and by our choice of $\lambda > 1$, and we call it the *iteration tree* defined by those data. We thus have

$$\Phi(w)H = Hw : T_\infty \to T_\infty \tag{3.5}$$

for all $w \in F_N$, and

$$ih = Hi : Z \to T_\infty\,. \tag{3.6}$$

Remark 3.2. (1) With the above generality, of course, the iteration tree T_∞ may well consist of a single point only.

(2) The generality of the construction presented here also allows situations where the iteration tree T_∞ is non-zero, but not minimal. This can be achieved easily through attaching equivariantly additional edges e to Z along their terminal endpoint, with $h(e) = e \circ \gamma$ for some suitable path γ in Z.

To finish this section we want to consider a special case of the above construction which plays an important role in the subsequent sections:

Lemma 3.3. *Let Z, h and Φ be as above. Assume furthermore that Z is a topological tree, provided with a pseudo-metric d as above, which has the following properties, where Z^* denotes the minimal F_N-invariant subtree of Z (= the intersection of all F_N-invariant subtrees):*

(1) *There is a constant $C \geq 0$ such that every point in Z has distance at most C from Z^*.*
(2) *Every arc in Z has as stabilizer a subgroup of F_N which is of rank ≤ 1.*

Then the iteration tree T_∞ has trivial arc stabilizers, satisfies BBT, and is contained as subtree in the metric closure $\overline{T^*_\infty}$ of its minimal subtree T^*_∞:

$$T_\infty \subset \overline{T^*_\infty}$$

Proof. In order to see that property (2) is inherited by the $\mathbb{R}$-tree T_∞ we consider any segment $[x, y] \subset Z$ and any $w \in F_N$. We assume that $d_\infty(x, y) > 0$, and that there is no proper subsegment $[x', y'] \subsetneqq [x, y]$ with $d_\infty(x', y') = d_\infty(x, y)$ (or else we pass to such a subsegment). If w stabilizes the image segment $[i(x), i(y)]$, then one has $d_\infty(x, wx) = d_\infty(y, wy) = 0$, and since Z is topologically a tree, any two points on the unique embedded arc connecting x to wx, or on the similar arc connecting y to wy, have d_∞-distance 0. From the above choice of x and y we deduce that each of these two arcs intersects the segment $[x, y]$ only in their common endpoint. Thus $[wx, wy]$ contains $[x, y]$ as subsegment.

But w acts on Z as homeomorphism which preserves the d_∞-length, so that it preserves the "no proper subsegment" property used in the above choice of x and y. Hence we deduce $wx = x$ and $wy = y$. Thus the stabilizers of $[x, y]$ and of its i-image in T_∞ are identical. This shows that the tree T_∞ is small.

In the proof of Lemma 2.8 in [11] a general argument has been given, for any small $\mathbb{R}$-tree, to show that equality (3.5) implies that the $\mathbb{R}$-tree in question has trivial arc stabilizers. Hence we conclude that T_∞ is very small, so that (see subsection 2.3) the minimal subtree T^*_∞ of T_∞ satisfies

$$T^*_\infty \in \overline{\text{cv}}_N\,.$$

By Proposition 2.6 this implies that T^*_∞ has property BBT, and so has its metric closure $\overline{T^*_\infty}$ (see Remark 2.5 (2)).

Furthermore, hypothesis (1) implies that with respect to the pseudo-metric d_k from (3.3) any point from Z has distance $\frac{C}{\lambda^k}$ from Z^*, so that in the limit we have $T_\infty \subset \overline{T}^*_\infty$, and thus T_∞ also inherits BBT. □

4. β-train-track maps

4.1. Definition and basic properties.

Definition 4.1. Let Γ be a finite graph, and let $f : \Gamma \to \Gamma$ be a graph map. Let $\Gamma_0 \subset \Gamma$ be a possibly non-connected subgraph which is f-invariant:

$$f(\Gamma_0) \subset \Gamma_0$$

(1) An edge path γ in Γ is *reduced relative* Γ_0 if every non-trivial backtracking subpath of γ (see section 2.1) lies in Γ_0. In other words, γ maps to a reduced edge path if Γ is mapped to the quotient graph defined by contracting every connected component of Γ_0 to a single point.
(2) An edge path γ in Γ is *legal relative* Γ_0 if $f^t(\gamma)$ is reduced relative Γ_0, for any integer $t \geq 0$.

Definition 4.2. Let Γ, f and Γ_0 be as in Definition 4.1. Let $\Gamma' \subset \Gamma$ be a connected, possibly not f-invariant subgraph, and let $r : \Gamma \to \Gamma'$ be a graph map which is a retraction (i.e. $r(e) = e$ for any edge of Γ').

(1) An edge path γ in Γ is *strongly reduced relative* Γ_0 if both, γ and $r(\gamma)$, are reduced relative Γ_0.
(2) An edge path γ in Γ is *strongly legal (rel.* Γ_0) if $f^t(\gamma)$ is strongly reduced relative Γ_0 for any integer $t \geq 0$.
(3) The graph map $f : \Gamma \to \Gamma$ has the *strong train track property* (rel. Γ_0) if every edge e is strongly legal (rel. Γ_0).

If the *relative part* $\Gamma_0 \subset \Gamma$ has been specified clearly beforehand, then we allow ourselves to drop the parenthesis "(rel. Γ_0)". We will also generalize slightly the use of the terminology "rel. Γ_0", in that a property of any mathematical object in Γ is true "relative Γ_0" if the image object has this property when passing to the quotient of Γ obtained by contracting every connected component of Γ_0 to a single point. For example, two edge paths γ and γ' in Γ are *equal relative* Γ_0 if γ and γ' coincide up to subpaths entirely contained in Γ_0.

Definition 4.3. Let $\mathcal{G}$ be a finite connected graph and let $f : \mathcal{G} \to \mathcal{G}$ be a graph map. Define $X \subset \mathcal{G}$ to be the (possibly non-connected) subgraph which consists of all vertices of $\mathcal{G}$ and of all edges that grow polynomially under iteration of f.

Then f is called a *weak β-train-track map* if the following conditions are satisfied (where for all paths or loops "strongly legal" means "strongly legal rel. X"):

(1) There exists a connected, possibly not f-invariant subgraph $\Gamma \subset \mathcal{G}$, which contains all vertices of $\mathcal{G}$, and a retraction $r : \mathcal{G} \to \Gamma$, such that complement satisfies:
$$\mathcal{G} \smallsetminus \Gamma \ \subset \ X$$
(2) The map f has the strong train track property relative X.
(3) Strongly legal paths γ and γ' in $\mathcal{G}$ with $r(\gamma) = r(\gamma')$ are equal relative X.
(4) For every loop γ in $\mathcal{G}$ there is an exponent $t(\gamma) \geq 0$ and a strongly legal loop γ' such that $[rf^{t(\gamma)}(\gamma)] = r(\gamma')$.
(5) There is a finite filtration of $\mathcal{G}$, given by f-invariant and r-invariant subgraphs $\mathcal{G}_1 \subset \mathcal{G}_2 \subset \ldots \subset \mathcal{G}_s = \mathcal{G}$ which are not necessarily connected, pairwise distinct, and not contained in X. Relative to X there is no refinement of this filtration by decomposing some $\mathcal{G}_k$ into further f-invariant subgraphs. Setting $X_k = \mathcal{G}_k \cap X$ for $k \geq 1$ and $\mathcal{G}_0 = X \cap \Gamma$, one has:

$$r(X_k) \subset \mathcal{G}_{k-1} \qquad \text{for all} \quad k \geq 1 \tag{4.1}$$

Definition-Remark 4.4. Let $f : \mathcal{G} \to \mathcal{G}$ be a weak β-train-track map.

(1) We consider the "Galois covering" $\widehat{\mathcal{G}}$ associated to $\ker r \subset \pi_1\mathcal{G}$. It follows directly that the full lift $\widetilde{\Gamma} \subset \widehat{\mathcal{G}}$ of Γ is a simplicial tree, and that there is a unique lift $\widehat{r} : \widehat{\mathcal{G}} \to \widetilde{\Gamma}$ of r which is a retraction. We will use the notation $\widehat{X}$ to denote the full lift of X to $\widehat{\mathcal{G}}$, and a map $\widehat{f} : \widehat{\mathcal{G}} \to \widehat{\mathcal{G}}$ will always denote a lift of the map f.

It follows directly that any path γ in $\widehat{\mathcal{G}}$ is strongly legal relative $\widehat{X}$ with respect to $\widehat{f}$ if and only if γ is the lift of some path in $\mathcal{G}$ that is strongly legal relative X with respect to f. The same applies to "strongly reduced" etc.

(2) If the fundamental group of the subgraph $\Gamma \subset \mathcal{G}$ is identified with a free group F_N of finite rank $N \geq 1$ (which then acts as deck transformation group on $\widehat{\mathcal{G}}$ - a free action !), we refer to the covering space $\widehat{\mathcal{G}}$ also as the F_N*-covering* of $\mathcal{G}$.
(3) Sometimes the free group F_N is given a priori, and the identification $\pi_1\Gamma \cong F_N$ is given by a *marking* of $\mathcal{G}$, i.e. a epimorphism $\theta : \pi_1\mathcal{G} \twoheadrightarrow F_N$ which satisfies the condition $\theta = \theta \circ r_*$ and restricts to an isomorphism $\theta|_{\pi_1\Gamma}$. (We don't specify a base point of $\mathcal{G}$ as θ only matters up to inner automorphisms of F_N.)

As in subsection 2.1 we say that a marked β-train-track map $f : \mathcal{G} \to \mathcal{G}$ *represents* an automorphism $\varphi \in \text{Out}(F_N)$ if the given marking satisfies $\varphi\theta = \theta f_*$. If $\mathcal{G}$ is unmarked, then f *represents* φ if there is a marking θ of $\mathcal{G}$ such that f represents φ with respect to θ.

From the above definitions we observe directly the following:

Lemma 4.5. *For any strongly reduced edge path γ in $\widehat{\mathcal{G}}$ every vertex on γ is also contained in the reduced path $\widehat{r}(\gamma)$.* □

Remark 4.6. The author would like to point out that weak β-train-track maps as define above are a generalization of stronger β-train-track maps that have been considered and studied before, in various dialects (see [16, 19]). These "strong" β-train-track maps have a lot more specific properties, like for example (for the terminology used below see for instance [21]):

1. The composed map $r \circ f|_\Gamma : \Gamma \to \Gamma$ is a classical relative train track map as in [4].
2. Every *auxiliary edge* e, i.e. e belongs to $\mathcal{G} \smallsetminus \Gamma$, has the property that there is a periodic INP-path $\eta = \gamma \circ \gamma'$ in $\mathcal{G}$ with strongly legal branches γ, γ' and same endpoints as e, such that $r(e) = r(\eta)$.
3. The r-image of any strongly legal path in $\mathcal{G}$ is pseudo-legal (=a legal concatenation of legal and periodic INP subpaths) in Γ.

However, the precise definition of a strong β-train-track map is a lot more involved, so that we have preferred to present here the "light" version given in Definition 4.3, which only states the properties needed in this paper. Since the stronger version will not appear any further in this paper, we refer from now on to weak β-train-track maps simply as *β-train-track maps*.

Proposition 4.7. *Let $f : \mathcal{G} \to \mathcal{G}$ be a β-train-track map which represents $\varphi \in \text{Out}(F_N)$ with respect to some marking θ of $\mathcal{G}$. Then the following holds:*

1. *Every strongly reduced path γ in $\widehat{\mathcal{G}}$ is relative $\widehat{X}$ equal to a quasi-geodesic, with respect to the combinatorial metric on $\widehat{\mathcal{G}}$ (for quasi-geodesy constants fixed independently of γ).*
2. *For every conjugacy class $[w] \subset F_N$ there is a strongly legal loop $\gamma(w)$ in $\mathcal{G}$ and an exponent $t(w) \geq 0$ such that $\gamma(w)$ represents $\varphi^{t(w)}([w])$ with respect to the given marking $\theta : \pi_1\mathcal{G} \twoheadrightarrow F_N$.*
   The loop $\gamma(w)$ is uniquely determined relative X by $\varphi^{t(w)}([w])$.

Proof. (1) Since we are allowed to modify γ within $\widehat{X}$ we can assume that γ is reduced (meant here in an absolute sense and not rel. X), and that γ doesn't cross over any edge e from $\widehat{X} - \widetilde{\Gamma}$ with $\widehat{r}(e) \subset \widehat{X}$ (or else we replace e by $\widehat{r}(e)$). It follows from Lemma 4.5 that there is an upper bound to the combinatorial length of any backtracking subpath in $\widehat{r}(\gamma)$. Thus γ fellow travels the geodesic $[\widehat{r}(\gamma)]$ in the tree $\widetilde{\Gamma}$, and since the embedding $\widetilde{\Gamma} \subset \widehat{\mathcal{G}}$ is a quasi-isometry, it follows that γ is a quasi-geodesic in $\widehat{\mathcal{G}}$.
(2) This is a direct consequence of properties (3) and (4) of Definition 4.3. □

The prime examples of β-train-track maps are absolute train track maps $f : \Gamma \to \Gamma$ without periodic INP-paths. Here one simply sets $\mathcal{G} = \Gamma$ and $r = \text{id}_\mathcal{G}$ and notices that in this case the notions of "legal" and "strongly legal" edge paths agree.

However, as is well known, not every $\varphi \in \mathrm{Out}(F_N)$ can be represented by an absolute train track map, not even if one assumes that φ is hyperbolic (see [20]). But this is remedied by the following:

Theorem 4.8 ([19]). *For every automorphism $\varphi \in \mathrm{Out}(F_N)$ there is a β-train-track map that represents φ.*

4.2. **β-train-tracks and $\mathbb{R}$-trees.**

In this subsection we will draw some first connections between β-train-track maps and isometric F_N-actions on $\mathbb{R}$-trees. Throughout this subsection we will work with the following data, which are natural assumptions, in light of what has been presented in section 3:

Hypothesis 4.9. Let $\varphi \in \mathrm{Out}(F_N)$ and let $\Phi \in \mathrm{Aut}(F_N)$ be a lift of φ. Let $f : \mathcal{G} \to \mathcal{G}$ be a β-train-track map which represents φ, and let $\widehat{f} : \widehat{\mathcal{G}} \to \widehat{\mathcal{G}}$ be the lift of f to the F_N-covering $\widehat{\mathcal{G}}$ of $\mathcal{G}$ which represents Φ. Assume that the vertex set of $\widehat{\mathcal{G}}$ is equipped with an F_N-invariant pseudo-metric d.

Let T be an $\mathbb{R}$-tree with isometric F_N-action (where T is possibly non-minimal or trivial), and let $H : T \to T$ be a homothety with stretching factor $\lambda > 1$ which represents Φ:

$$\Phi(w)H = Hw : T \to T \qquad \text{for all} \qquad w \in F_N \tag{4.2}$$

Let $i : \widehat{\mathcal{G}} \to T$ be an F_N-equivariant map which is distance decreasing:

$$d(i(P), i(Q)) \leq d(P, Q) \quad \text{for any vertices} \quad P, Q \quad \text{of} \quad \widehat{\mathcal{G}} \tag{4.3}$$

Furthermore assume that $i, \widehat{f}$ and H satisfy:

$$i\widehat{f}(P) = Hi(P) \quad \text{for any vertex} \quad P \quad \text{of} \quad \widehat{\mathcal{G}} \tag{4.4}$$

Lemma 4.10. *Assume that all data from Hypothesis 4.9 are given. Then for any edge e of $\widehat{\mathcal{G}}$, which under iteration of $\widehat{f}$ grows polynomially, the endpoints P and Q of e are mapped by i to the same point $i(P) = i(Q) \in T$.*

Proof. If one has $i(P) \neq i(Q)$, then by the assumption $\lambda > 1$ the distance $d(H^t i(P), H^t i(Q))$ grows exponentially in t. From equalities (4.4) and (4.3) we deduce that the same holds for the distance $d(\widehat{f}^t(P), \widehat{f}^t(Q))$. But this contradicts the assumption that e grows polynomially. □

From Lemma 4.10 and from the definition of $X \subset \mathcal{G}$ in Definition 4.3 we obtain directly that for any *vertex space* X_v of $\widehat{\mathcal{G}}$, i.e. any connected component of the full lift $\widehat{X} \subset \widehat{\mathcal{G}}$ of the relative part $X \subset \mathcal{G}$, the i-image must be a single point $i(X_v) \in T$. We say that an edge path γ in $\widehat{\mathcal{G}}$ *traverses* a vertex space $X_v \subset \widehat{\mathcal{G}}$ if γ crosses over at least one vertex of X_v.

Proposition 4.11. *Assume that Hypothesis 4.9 is given, and assume that T satisfies BBT (see subsection 2.4). Let γ be a strongly legal edge path in $\widehat{\mathcal{G}}$ with terminal vertices spaces X_P and X_Q.*

Then any vertex space X_R traversed by γ is mapped by i to a point $i(X_R)$ on the geodesic segment $[i(X_P), i(X_Q)] \subset T$. Furthermore, the order of such traversed vertices spaces X_j along the path γ is preserved by the map i.

Proof. We recall from Proposition 4.7 (1) that any strongly legal path γ in $\widehat{\mathcal{G}}$ is a quasi-geodesic relative to the full lift $\widehat{X} \subset \widehat{\mathcal{G}}$ of the relative part X. In other words: there is a quasi-geodesic γ' in $\widehat{\mathcal{G}}$, with respect to the combinatorial metric in $\widehat{\mathcal{G}}$, which differs from γ only in subpaths entirely contained in $\widehat{X}$.

But since we derived above that i maps any vertex space X_v to a single point in T, it follows $i(\gamma) = i(\gamma')$. Hence we derive from the BBT assumption on T and from Lemma 2.4 and Remark 2.5 (1) that there is a constant $c \geq 0$, which does not depend on the choice of γ, such that i maps the strongly legal path γ into the c-neighborhood of the geodesic segment $[i(X_P), i(X_Q)]$.

Now any of the paths $\widehat{f}^t(\gamma)$ with $t \geq 0$ is again strongly legal, and for any vertex space X_R traversed by γ the image vertex space $\widehat{f}^t(X_R)$ is again traversed by the reduced edge path $\widehat{f}^t(\gamma)$, so that the point $i(\widehat{f}^t(X_R))$ must be contained in the c-neighborhood of the geodesic segment $[i(\widehat{f}^t(X_P)), i(\widehat{f}^t(X_Q))]$, for any $t \geq 0$. Let $c' \geq 0$ be the distance of $i(X_R)$ to $[i(\widehat{f}^t(X_P)), i(\widehat{f}^t(X_Q))]$.

Since we know from condition (4.4) of Hypothesis 4.9 that $i(\widehat{f}^t(X_R)) = H^t(i(X_R))$ and that $[i(\widehat{f}^t(X_P)), i(\widehat{f}^t(X_Q))] = [H^t(i(X_P)), H^t(i(X_Q))]$, and since H is a homothety with stretching factor $\lambda > 1$, it follows that the distance between $H^t(i(X_R))$ and $[H^t(i(X_P)), H^t(i(X_Q))]$ is equal to $\lambda^t c'$, so that we can conclude from the previous paragraph and from the assumption $\lambda > 1$ of Hypothesis 4.9 that $c' = 0$. This means that $i(X_R)$ must actually lie on the geodesic segment $[i(X_P), i(X_Q)] \subset T$.

In order to show that i preserves the order of the vertex spaces X_j traversed by γ it suffices to consider the sub-edge-path γ' of γ bounded by any two such vertex spaces X_1 and X_2: The previously derived result shows that any further such vertex space X_3 traversed γ' is mapped by i to a point on the segment $[i(X_1), i(X_2)] \subset T$. □

Remark 4.12. We would like to point out a subtlety in the above proof: The fact that the given strongly legal path γ (or the derived edge path γ') is a quasi-geodesic with respect to the combinatorial metric in $\widehat{\mathcal{G}}$ does *not* imply that γ (or γ') is a quasi-geodesic with respect to the given distance function d.

In fact, this discrepancy is the reason why in the approach pursued in this paper one needs to keep track of the subgraph $\Gamma \subset \mathcal{G}$ in Definition 4.3.

On the other hand, however, one can deduce the fact that γ is a quasi-geodesic with respect to d from the above Proposition 4.11, provided that no edge e of d-length > 0 is contracted by the map i to a single point.

Remark 4.13. When dealing with β-train-track maps $f : \mathcal{G} \to \mathcal{G}$, the philosophy is always to consider everything relative to the polynomially growing part $X \subset \mathcal{G}$. Equivalently, when passing to the F_N-covering $\widehat{\mathcal{G}}$, one actually works (at least morally) in the quotient space of $\widehat{\mathcal{G}}$ defined by contracting every vertex space X_v to a single point.

However, when applying the previous Proposition 4.11 it is sometimes useful to notice that its statement can easily improved to the following statement, which is *not* meant relative $\widehat{X}$:

Any strongly legal path γ in $\widehat{\mathcal{G}}$ is mapped by $i : \widehat{\mathcal{G}} \to T$ in a vertex-order preserving fashion to a segment in T.

This follows directly from the fact that every vertex space X_v is mapped by i to a single point $i(X_v)$ in T, so that every subpath of γ entirely contained in X_v is mapped by i in a "vertex-order preserving fashion" to this point $i(X_v)$.

5. Perron-Frobenius trees

Throughout this section we work with a β-train-track map $f : \mathcal{G} \to \mathcal{G}$ which represents a given automorphism $\varphi \in \mathrm{Out}(F_N)$. We consider a *Perron-Frobenius eigenvector $\vec{v}^*$ for f* with eigenvalue $\lambda > 1$, by which we mean a non-negative row eigenvector $\vec{v}^*$ of the transition matrix $M(f)$ with eigenvalue λ. As laid out in subsection 2.2 the vector $\vec{v}^*$ determines an associated edge length function

$$L^{\vec{v}^*} : \mathrm{Edges}(\mathcal{G}) \to \mathbb{R}_{\geq 0}\,,$$

which in turn defines the limit length $L^{\vec{v}^*}_\infty$ (see Definition 2.1 (5)).

Recall from Remark 2.2 (3) that for any edge e from the polynomially growing part $X \subset \mathcal{G}$ one necessarily has $L^{\vec{v}^*}(e) = L^{\vec{v}^*}_\infty(e) = 0$

In particular (see equalities (2.7) and (2.8)), we will use below that for any strongly legal path γ one has the equalities

$$L^{\vec{v}^*}_{\infty}(\gamma) = \lim_{t\in\mathbb{N}} \frac{1}{\lambda^t} L^{\vec{v}^*}([f^t(\gamma)]) = \lim_{t\in\mathbb{N}} \frac{1}{\lambda^t} L^{\vec{v}^*}(f^t(\gamma)) = L^{\vec{v}^*}(\gamma)\,. \tag{5.1}$$

We now recall from Proposition 4.7 (2) that for every conjugacy class $[w] \subset F_N$ there is an exponent $t(w) \geq 0$ and a strongly legal loop $\gamma(w)$ which represents the image conjugacy class $\varphi^{t(w)}([w])$. Here $\gamma(w)$ is uniquely determined, relative to the polynomially growing part, by $\varphi^{t(w)}([w])$.

It hence follows from (5.1) that any PF-length $L^{\vec{v}^*}$ on $\mathcal{G}$ induces a length function

$$L^{\vec{v}^*}_{F_N} : F_N \to \mathbb{R}_{\geq 0}$$

which is well defined by

$$L^{\vec{v}^*}_{F_N}(w) := \frac{1}{\lambda^{t(w)}} L^{\vec{v}^*}_{\infty}(\gamma(w)) = \frac{1}{\lambda^t} L^{\vec{v}^*}(f^{t-t(w)}(\gamma(w))) \tag{5.2}$$

for any $t \geq t(w)$.

Remark 5.1. There is always some element $w \in F_N$ with $L^{\vec{v}^*}_{F_N}(w) > 0$. This follows as in Remark 2.2 (4), with "legal" replaced by "strongly legal".

The goal of this section is to prove the following, which is a strong version of Proposition 1.1:

Theorem 5.2. *Consider any $\varphi \in \mathrm{Out}(F_N)$, and let $f : \mathcal{G} \to \mathcal{G}$ be a β-train-track map which represents φ via some marking θ of $\mathcal{G}$.*

Let $\vec{v}^* \neq \vec{0}$ be a Perron-Frobenius eigenvector for f with eigenvalue $\lambda > 1$, let $L^{\vec{v}^*}$ be the associated edge length function on $\mathcal{G}$, and let $L^{\vec{v}^*}_{F_N}$ be the derived length function on F_N as in (5.2).

Then there exists an $\mathbb{R}$-tree T with the following properties:

(1) The tree T is φ-invariant with stretching factor λ. Furthermore T has trivial arc stabilizers and satisfies BBT. The tree T is contained in the metric completion $\overline{T^*}$ of its minimal subtree T^*.

(2) There exists a surjective F_N-invariant map $i = i^{\vec{v}^*}\colon \widehat{\mathcal{G}} \twoheadrightarrow T$ which maps every strongly legal path γ to a geodesic segment of length $L^{\vec{v}^*}(\gamma)$ in T.

(3) The map i semi-commutes with the lift $\widehat{f}$ of f to $\widehat{\mathcal{G}}$ and the homothety $H : T \to T$, if both represent the same lift $\Phi \in \mathrm{Aut}(F_N)$ of φ (compare (2.2) and (2.9):

$$Hi = i\widehat{f} \tag{5.3}$$

(4) The translation length function associated to T satisfies

$$||w||_T = L^{\vec{v}^*}_{F_N}(w)$$

for all $w \in F_N$. In particular, the minimal subtree T^* of T is non-trivial.

Proof. We prove the statement by induction over the number s of subgraphs $\mathcal{G}_1 \subset \mathcal{G}_2 \subset \ldots \subset \mathcal{G}_s = \mathcal{G}$ of the β-train-track map $f : \mathcal{G} \to \mathcal{G}$, provided by property (5) of Definition 4.3. For the purpose of the induction we need to work with a small modification of our claim: We include the possibility that $\vec{v}^*$ is the zero vector, and hence T is degenerated to the trivial tree T_0 which consists of just a single point. Hence we denote - just for the length of this proof - by $\overline{\mathrm{cv}}^*_N$ the set of $\mathbb{R}$-trees T for which the minimal subtree T^* belongs to $\overline{\mathrm{cv}}_N$ or is equal to T_0, and which satisfy $T \subset \overline{T^*}$.

By setting $\mathcal{G}' := \mathcal{G}_{s-1}$ we can decompose $\mathcal{G}$ into an f-invariant, not necessarily connected subgraph $\mathcal{G}' \subset \mathcal{G}$ and a "top stratum" $\mathcal{H} := \mathcal{G} \smallsetminus \mathcal{G}'$, such that $\mathcal{H} \cap \Gamma$ determines in the transition matrix $M(f)$ a diagonal block $M_{\mathcal{H}}$ which is irreducible (including the exceptional case of the (1×1)-zero-matrix, which may occur in the inductive procedure). The irreducibility of $M_{\mathcal{H}}$ is a direct consequence of the condition in property (5) of Definition 4.3, that relative X there is no proper f-invariant refinement of the given strata-decomposition.

Then $\vec{v}^*$ determines on every f-invariant connected component $\mathcal{G}'_k$ of $\mathcal{G}'$ an eigenvector $\vec{v}^*_k$ with eigenvalue λ (including the possibility $\vec{v}^*_k = \vec{0}$), so that by induction we can assume that for $\mathcal{G}'_k$ there is a PF-tree $T_k := T^{\vec{v}^*_k}$ which has all properties as stated in the claim (including the possibility $T_k = T_0$).

If the component $\mathcal{G}'_k$ of $\mathcal{G}'$ is not f-invariant but only f-periodic, then we replace f by a suitable positive power, proceed as above, and pass back to the corresponding root of this power, which then permutes the components of $\mathcal{G}'$ in question. Correspondingly, the associated PF-trees are permuted by the resulting homotheties $H_k : T_k \to T_{k'}$ with stretching factor λ.

If a component $\mathcal{G}'_k$ of $\mathcal{G}'$ is not f-periodic, we fix the smallest exponent $t_k \geq 1$ such that $f^{t_k}(\mathcal{G}'_k)$ is contained in some f-periodic component $\mathcal{G}'_{k'}$ of $\mathcal{G}'$. Since f induces φ which is an automorphism, and since $\mathcal{G}' = \mathcal{G}_{s-1}$ is r-invariant (see Definition 4.3 (5)), it follows that $\theta(\pi_1 \mathcal{G}'_k) = \theta\, r_*(\pi_1 \mathcal{G}'_k)$ is trivial (see Definition-Remark 4.4 (3)). Thus every lift $\widehat{\mathcal{G}}'_h$ of $\mathcal{G}'_k$ to $\widehat{\mathcal{G}}$ is a homeomorphic copy of $\mathcal{G}'_k$. Such a connected component $\widehat{\mathcal{G}}'_h$ of the full lift $\widehat{\mathcal{G}}' \subset \widehat{\mathcal{G}}$ of $\mathcal{G}'$ will be called *inessential*; any other connected component $\widehat{\mathcal{G}}'_j$ of $\widehat{\mathcal{G}}'$ will be called *essential*.

We now build a "mixed" $\mathbb{R}$-tree $T_{\mathcal{H}}$ in several steps: We first consider the F_N-covering $\widehat{\mathcal{G}}$ of $\mathcal{G}$ and a lift $\widehat{f} : \widehat{\mathcal{G}} \to \widehat{\mathcal{G}}$ which represents some lift $\Phi \in \mathrm{Aut}(F_N)$ of the given outer automorphism φ. We now replace in $\widehat{\mathcal{G}}$ every essential connected component $\widehat{\mathcal{G}}'_j$ of $\widehat{\mathcal{G}}'$ by a copy T_j of the corresponding $\mathbb{R}$-tree T_k, where "corresponding" means that $\widehat{\mathcal{G}}'_j$ is a covering of the connected component $\mathcal{G}'_k$. Let $i_j : \widehat{\mathcal{G}}'_j \to T_j$ be a copy of the map $i_k^{\vec{v}^*_k} : \widehat{\mathcal{G}}'_k \to T_k$ provided by induction through our claim.

For any inessential connected component $\widehat{\mathcal{G}}'_h$ of $\widehat{\mathcal{G}}'$, which is a copy of some non-periodic component $\mathcal{G}'_k$ of $\mathcal{G}'$, we consider the essential component $\widehat{\mathcal{G}}'_j$ which contains $\widehat{f}^{t_k}(\widehat{\mathcal{G}}'_h)$. We set $T'_h := i_j(\widehat{f}^{t_k}(\widehat{\mathcal{G}}'_h)) \subset T_j$ and observe that T'_h is a finite metric tree. In order to built $T_{\mathcal{H}}$ from $\widehat{\mathcal{G}}$, any inessential $\widehat{\mathcal{G}}'_h$ is replaced by the scaled-down finite metric tree $T_h = \lambda^{-t_k} T'_h$. We define a map $i_h : \widehat{\mathcal{G}}'_h \to T_h$ as composition of the map $i_j \widehat{f}^{t_k}$ with the scaling-down homothety $T'_h \to T_h$.

To finish the construction of $T_{\mathcal{H}}$ at this state, any edge e' of $\widehat{\mathcal{G}} \smallsetminus \widehat{\mathcal{G}}'$ is replaced in $T_{\mathcal{H}}$ by a copy, called e, where the endpoints of e are attached to the above trees T_j and T_h according to the following rule: If an endpoint of e' is attached to some vertex P of an essential component $\widehat{\mathcal{G}}'_j \subset \widehat{\mathcal{G}}'$, then we attach the corresponding endpoint of e to $i_j(P)$. If an endpoint of e' is attached to some vertex Q in an inessential component $\widehat{\mathcal{G}}'_h \subset \widehat{\mathcal{G}}'$, then we attach the corresponding endpoint of e to $i_h(Q)$.

This gives an F_N-equivariant map $i_{\mathcal{H}} : \widehat{\mathcal{G}} \to T_{\mathcal{H}}$, which is a simplicial isomorphism on $\widehat{\mathcal{G}} \smallsetminus \widehat{\mathcal{G}}'$ and restricts on each component $\widehat{\mathcal{G}}'_j$ or $\widehat{\mathcal{G}}'_h$ of $\widehat{\mathcal{G}}' \subset \widehat{\mathcal{G}}$ to the map i_j or i_h respectively.

Note that at this state of the construction $T_{\mathcal{H}}$ is not necessarily a tree, since there may be edges in X that do not belong to $\mathcal{G}'$. However, from the inclusion (4.1), applied to the top stratum of $\mathcal{G}$, we know that for any edge e' from $\widehat{X} - \widehat{\mathcal{G}}'$ the endpoints P and Q both belong to the same component $\widehat{\mathcal{G}}'_j$ or $\widehat{\mathcal{G}}'_h$, and from Lemma 4.10 we deduce $i_{\mathcal{H}}(P) = i_{\mathcal{H}}(Q)$. Hence we can modify $T_{\mathcal{H}}$ and $i_{\mathcal{H}}$ in that, rather than attaching an edge e corresponding to e' along its endpoints at $i_{\mathcal{H}}(P) = i_{\mathcal{H}}(Q)$, we map all of e' by the map $i_{\mathcal{H}}$ to the point $i_{\mathcal{H}}(P) = i_{\mathcal{H}}(Q)$. After this modification the "mixed" space $T_{\mathcal{H}}$ is indeed topologically a tree, since the edges from $\widehat{\mathcal{G}} \smallsetminus \widehat{\mathcal{G}}'$ outside $\widehat{X}$ belong to the simplicial tree $\widetilde{\Gamma} \subset \widehat{\mathcal{G}}$ (see Definition-Remark 4.4 (1)). By construction, the F_N-action on $\widehat{\mathcal{G}}$ induces an F_N-action on $T_{\mathcal{H}}$, which makes the quotient map $i_{\mathcal{H}}$ F_N-equivariant.

We can now define an F_N-equivariant map $f_{\mathcal{H}} : T_{\mathcal{H}} \to T_{\mathcal{H}}$ in the following manner:

(a) Any edge $e = i_{\mathcal{H}}(e')$ with e' in $\widehat{\mathcal{G}} \smallsetminus \widehat{\mathcal{G}}'$ is mapped to the path $f_{\mathcal{H}}(e) := i_{\mathcal{H}}(\widehat{f}(e'))$. More precisely, any point x on e, which is the $i_{\mathcal{H}}$-image of some well defined point x' on e', is mapped by $f_{\mathcal{H}}$ to $i_{\mathcal{H}}(\widehat{f}(x'))$.

(b) For any essential connected component $\widehat{\mathcal{G}}'_j$ we know from the f-invariance of $\mathcal{G}'$ that there is some essential connected component $\widehat{\mathcal{G}}'_{j'}$ with $\widehat{f}(\widehat{\mathcal{G}}'_j) \subset \widehat{\mathcal{G}}'_{j'}$. We define $f_{\mathcal{H}}$ on $T_j = i_{\mathcal{H}}(\widehat{\mathcal{G}}'_j)$ to be the homothety $w_j H_j : T_j \to T_{j'}$, where H_j is a copy of the corresponding homothety $H_k : T_k \to T_{k'}$ provided above by our induction hypothesis. The element $w_j \in F_N$ is a correction term needed because of the possible difference between the restriction to $\pi_1 \mathcal{G}_k$ of the above automorphism $\Phi \in \mathrm{Aut}(F_N)$ (which represents the chosen lift $\widehat{f}$ of f) on one hand, and on the other hand the isomorphism Φ_k which is represented by H_k (compare (2.3)).

(c) For any inessential component $\widehat{\mathcal{G}}'_h$ let $\widehat{\mathcal{G}}'_{h'}$ be the (essential or inessential) component which contains $\widehat{f}(\widehat{\mathcal{G}}'_h)$. Then our definition of $T_h = i_{\mathcal{H}}(\widehat{\mathcal{G}}'_h)$ gives straight forward a map $H_h : T_h \to T_{h'}$ for which the co-restriction $T_h \to H_h(T_h)$ is a homothety with stretching factor λ. This map defines $f_{\mathcal{H}}$ on T_h.

Combining these three definitions together gives a map $f_{\mathcal{H}} : T_{\mathcal{H}} \to T_{\mathcal{H}}$, which by construction also represents Φ, and which satisfies the equality

$$f_{\mathcal{H}} i_{\mathcal{H}} = i_{\mathcal{H}} \widehat{f}\,.$$

We now provide every edge $e = i_{\mathcal{H}}(e')$ of the "top stratum" $i_{\mathcal{H}}(\widehat{\mathcal{G}} \smallsetminus \widehat{\mathcal{G}}')$ of $T_{\mathcal{H}}$ with the length $L^{\vec{v}^*}(e')$ defined by the eigenvector $\vec{v}^*$ and lifted from $\mathcal{G}$ to $\widehat{\mathcal{G}}$. Since e' is strongly legal (by Definition 4.3 (2)), we have from (5.1) the equality $L^{\vec{v}^*}(e') = L^{\vec{v}^*}_{\infty}(e')$. We use the convention from Definition 2.1 (4) and (5) to distribute this length along e', and transfer it via $i_{\mathcal{H}}$ to e. In this way we obtain an F_N-equivariant pseudo-metric $d_{\mathcal{H}}$ on the tree $T_{\mathcal{H}}$ which extends the given $\mathbb{R}$-tree metrics on any of the subtrees T_j or T_h.

Since T_0 is topologically a tree, the pseudo-metric $d_{\mathcal{H}}$ is automatically 0-hyperbolic. Furthermore it satisfies by construction

$$d(f_{\mathcal{H}}(x), f_{\mathcal{H}}(y)) \leq \lambda d(x, y)$$

for any points x, y both contained in an edge e from the top stratum, or both contained in any of the metric trees T_j or T_h. Hence, using again that $T_{\mathcal{H}}$ is topologically a tree, this inequality holds also for arbitrary two points x and y of $T_{\mathcal{H}}$.

Thus $f_{\mathcal{H}}$ and $T_{\mathcal{H}}$ satisfy all conditions used in section 3 to define (for $Z = T_{\mathcal{H}}$ and $h = f_{\mathcal{H}}$) "intermediate" 0-hyperbolic pseudo-metrics d_k on $T_{\mathcal{H}}$, which all agree on any of the subtrees T_j or T_h with the given $\mathbb{R}$-tree's metric. The pseudo-metrics d_k converge to a 0-hyperbolic limit pseudo-metric d_∞, which gives as associated canonical metric quotient space an $\mathbb{R}$-tree T with induced action of F_N by isometries. From the construction in section 3 we also obtain a surjective, F_N-equivariant and distance-decreasing map $i'_{\mathcal{H}} : T_{\mathcal{H}} \to T$, which restricts to an isometric map on any of the subtrees T_j or T_h. The construction in section 3 yields furthermore a homothety $H : T \to T$ which restricts on any T_j to the homothety $w_j H_j$, and on any T_h to the homothety H_h. From equality (3.6) we furthermore obtain

$$H i'_{\mathcal{H}} = i'_{\mathcal{H}} f_{\mathcal{H}}\,,$$

so that the composition with the the above derived equality $f_{\mathcal{H}} i_{\mathcal{H}} = i_{\mathcal{H}} \widehat{f}$ gives the claimed equality (5.3), for the composed map $i := i'_{\mathcal{H}} i_{\mathcal{H}}$. This proves part (3) of the claimed statement.

We now use that by induction claim (1) holds for the subtrees T_j, and that the F_N-action on $\widehat{\mathcal{G}}$ is free, in order to deduce that $T_{\mathcal{H}}$ and $f_{\mathcal{H}}$ satisfy all conditions of Lemma 3.3. This proves part (1) of the claimed statement, and shows in particular that T belongs to $\overline{\mathrm{cv}}^*_N$.

In order to prove part (2) we consider any strongly legal path γ in $\widehat{\mathcal{G}}$. Since we know from part (1) that T satisfies BBT, we can apply Proposition 4.11 to obtain that $i(\gamma)$ is a segment in T, and that the map i preserves along this segment the order of vertices and hence the order in which the maximal subpaths γ_ℓ of γ occur that are contained in $\widehat{\mathcal{G}}'$. For any such subpath γ_ℓ we know from our induction procedure that the segments $i_{\mathcal{H}}(\gamma_\ell)$ and thus also $i(\gamma_\ell)$ have the same length as $L^{\vec{v}^*}(\gamma_\ell)$. For any top-stratum edge e on γ, i.e. e belongs to $\widehat{\mathcal{G}} \smallsetminus \widehat{\mathcal{G}}'$, we know from the above definition of the pseudo-metric $d_{\mathcal{H}}$ that the endpoints of $i_{\mathcal{H}}(e)$ in $T_{\mathcal{H}}$ have distance equal to $L^{\vec{v}^*}(e)$. Hence the map $i_{\mathcal{H}}$ preserves along γ the pseudo-distance defined by $L^{\vec{v}^*}$, and the same is true for all iterates $\widehat{f}^t(\gamma)$ of γ. It follows that the pseudo-metric iteration procedure on $T_{\mathcal{H}}$ to define the limit metric space T is described along γ by the length function iteration of $L^{\vec{v}^*}$ given by the equality (5.1). This proves statement (2) of our claim.

Since by the defining equality (5.2) the length function $L^{\vec{v}^*}_{F_N}$ on the elements (or rather "conjugacy classes") of F_N reduces to considering strongly legal loops only, we obtain now part (4) of the claim as direct consequence of the already proved part (2). The fact that T^* is non-trivial is a consequence of the observation stated as Remark 5.1. This completes the proof. ◻

Definition-Remark 5.3. (1) The "clumsiness" in the above proof which comes from including non-minimal trees in the iteration process is unfortunately forced on us: For example, it may well occur that in some f-invariant subgraph $\mathcal{G}_k$ of $\mathcal{G}$ the top stratum consists of a single edge e only, which is attached to the rest of $\mathcal{G}_k$ only at one of its endpoints, say P, and is mapped by f to a path $f(e) = e \circ \gamma$. If, through a proper choice of γ, this image path "picks up" some non-zero limit length from lower strata, then it can happen that in the limit the endpoint of e other than P will be contained in the completion of the minimal subtree T^*_k, but not in T^*_k itself. On the other hand, this "other endpoint" of e may be needed in order to attach edges from the top stratum of $\mathcal{G}$, and hence its lifts to $\widehat{\mathcal{G}}$ as well as their $i_{\mathcal{H}}$-images are essential in the construction of the mixed tree $T_{\mathcal{H}}$.
(2) The classical construction of the φ-invariant tree T^* via absolute train track maps $f : \Gamma \to \Gamma$ rather than β-train-track maps can be reinterpreted (see [18]) as defining $T^* \in \partial \text{cv}_N$ as limit of the sequence of rescaled φ-iterates $\frac{1}{\lambda^t} T \cdot \varphi^t$ in unprojectivized Outer space cv_N, which makes the arguments for the above proof (for this particular case) a lot easier. However, a similar approach for the general case must fail, independently of which kind of train track maps one may try to put at work. The reason is that for some eigenvectors $\vec{v}^*$ the projective tree class $[T^*]$ is simply not the limit point for any φ-orbit in Outer space CV_N, and not even for any φ-orbit of projectivized simplicial trees in ∂CV_N.
(3) Instead, the inductive procedure presented in the above proof, amounts to combining in the construction of the "mixed tree" $T_{\mathcal{H}}$ a simplicial component for the top stratum with iteratively obtained $\mathbb{R}$-trees for the connected components $\mathcal{G}_k$ of the lower strata. On the bottom strata the classical argument applies, as there the restriction of f are absolute train track maps (relative X). It follows by induction that each of the T^*_k are very small trees, and hence the mixed tree $T_{\mathcal{H}}$, after passing to its canonical metric quotient tree, possesses a minimal subtree, denoted $T^*_{\mathcal{H}}$, which belongs to $\overline{\text{cv}}_N$. We thus obtain the minimal subtree T^* of the limit tree T as forward limit point of the φ-orbit in $\overline{\text{CV}}_N$ of the projective class $[T^*_{\mathcal{H}}]$, which gives an alternative argument why T^* does belong to $\overline{\text{cv}}_N$. In fact, by (2.10) we actually know $T^* \in \partial \text{cv}_N$.
(4) To summarize this discussion, we introduce the terminology *Perron-Frobenius tree* for the minimal subtree T^* of the tree T obtained in Theorem 5.2, and denote it from now on by $T^{\vec{v}^*}$. We have thus shown:

Corollary 5.4. *Any Perron-Frobenius eigenvector $\vec{v}^*$ with eigenvalue $\lambda > 1$ for a β-train-track map $f : \mathcal{G} \to \mathcal{G}$ defines a Perron-Frobenius tree*

$$T^{\vec{v}^*} \in \partial \text{cv}_N \,.$$

If f represents $\varphi \in \mathrm{Out}(F_N)$, then $T^{\vec{v}^*}$ is projectively fixed by φ, and expanded by φ with factor λ. □

This proves Proposition 1.1 from the Introduction.

6. Proof of Theorem 1.2

6.1. **A criterion for Perron-Frobenius trees.**

In this subsection we give a criterion for projectively φ-invariant $\mathbb{R}$-trees to be Perron-Frobenius.

Proposition 6.1. *Let $f : \mathcal{G} \to \mathcal{G}$ be an β-train-track map which represents an automorphism $\varphi \in \mathrm{Out}(F_N)$, and let $\widehat{f} : \widehat{\mathcal{G}} \to \widehat{\mathcal{G}}$ be a lift of f to the F_N-covering $\widehat{\mathcal{G}}$ of $\mathcal{G}$ which represents some preimage $\Phi \in \mathrm{Aut}(F_N)$ of φ.*

Let T be a (not necessarily minimal) very small $\mathbb{R}$-tree with isometric F_N-action, which is equipped with a homothety $H : T \to T$ with stretching factor $\lambda > 1$ that also represents Φ. We assume in addition that T is contained in the metric closure $\overline{T^*}$ of its minimal subtree T^*.

Assume furthermore that for the vertex set $V(\widehat{\mathcal{G}})$ of $\widehat{\mathcal{G}}$ there exists an F_N-equivariant map $i : V(\widehat{\mathcal{G}}) \to T$ which satisfies:

$$i\widehat{f}(P) = Hi(P) \quad \text{for any vertex} \quad P \quad \text{of} \quad \widehat{\mathcal{G}} \tag{6.1}$$

Then there exists a Perron-Frobenius eigenvector $\vec{v}^*$ for f with eigenvalue λ, such that the minimal F_N-invariant subtree $T^* \subset T$ is F_N-equivariantly isometric to the Perron-Frobenius tree $T^{\vec{v}^*}$.

Proof. We define a length $L(\widehat{e})$ for every edge $\widehat{e}$ in $\widehat{\mathcal{G}}$, say with initial and terminal vertices P and Q, through

$$L(\widehat{e}) := d(i(P), i(Q))\,.$$

As i is assumed to be F_N-equivariant, this induces a well defined edge length function $L : \mathrm{Edges}(\mathcal{G}) \to \mathbb{R}_{\geq 0}$ on $\mathcal{G}$. As explained in Definition 2.1 (1) and (2) this defines for any edge path $\gamma = e_1 \ldots e_r$ in $\mathcal{G}$ (and similarly in $\widehat{\mathcal{G}}$) a path length $L(\gamma) = \sum_{k=1}^{r} L(e_k)$, and consequently a pseudo-metric d^L on the vertex set of $\widehat{\mathcal{G}}$.

It follows from the triangular inequality that the pseudo-metric d^L satisfies

$$d(i(P), i(Q)) \leq d^L(P, Q)$$

for any vertices P, Q of $\widehat{\mathcal{G}}$, so that all conditions of Hypothesis 4.9 are satisfied. In addition T satisfies BBT, since it is assumed that $T \subset \overline{T^*}$, and T (and thus T^*) is very small, so that Proposition 2.6 and Remark 2.5 (2) apply.

Hence we can use Proposition 4.11 and Remark 4.13 in order to deduce that for any edge $\widehat{e}$ of $\widehat{\mathcal{G}}$ all vertices of the strongly legal path $\widehat{f}(\widehat{e})$ are mapped by i in an order-preserving fashion to the geodesic segment $[i\widehat{f}(P), i\widehat{f}(Q)]$ in T, where P and Q are the endpoints of $\widehat{e}$. Hence we deduce from the above definition of the length function L that the length of the segment $[i\widehat{f}(P), i\widehat{f}(Q)]$ is equal to $L(\widehat{f}(\widehat{e}))$.

On the other hand, the length of $[i(P), i(Q)]$ is equal to $L(\widehat{e})$, and from hypothesis (6.1) we know $[i\widehat{f}(P), i\widehat{f}(Q)] = H([i(P), i(Q)])$. Since H is a homothety with stretching factor λ, we deduce that $L(\widehat{f}(\widehat{e})) = \lambda L(\widehat{e})$.

We now pass to the image edge e of $\widehat{e}$ in $\mathcal{G}$ and obtain

$$\lambda L(e) = L(f(e)) = \sum_{e' \in \mathrm{Edges}^+(\mathcal{G})} m_{e',e}\, L(e')\,, \tag{6.2}$$

where the $m_{e',e}$ are the coefficients of the transition matrix $M(f)$ as given in subsection 2.1. But (6.2) is the defining equality for the coefficients of a row eigenvector of a matrix with coefficients $m_{e',e}$, so that we conclude that the length function L on the edges of $\mathcal{G}$ (and of $\widehat{\mathcal{G}}$) is given by a Perron-Frobenius row eigenvector $\vec{v}^*$ of $M(f)$, with eigenvalue equal to the given stretching factor $\lambda > 1$ of H. This gives:

$$L = L^{\vec{v}^*}$$

It has been shown above that any strongly legal edge path γ is mapped by i to a segment $i(\gamma) \subset T$, and that the length of this segment is equal to $L(\gamma)$. Hence for any $w \in F_N$ we derive from the definition of $L^{\vec{v}^*}_{F_N}(w)$ in (5.2) that the translation length $||w||_T$ in T satisfies:

$$||w||_T = L^{\vec{v}^*}_{F_N}(w)$$

Since from property (4) of Theorem 5.2 we know that $L^{\vec{v}^*}_{F_N}(w) = ||w||_{T^{\vec{v}^*}}$, it follows directly (see subsection 2.3) that the minimal subtree T^* of T must be F_N-equivariantly isometric to the Perron-Frobenius tree $T^{\vec{v}^*}$. □

6.2. **Fixed points of homotheties.**

The purpose of this subsection is to construct a map i as in Proposition 6.1. We crucially use the dynamics of H to define this map i.

Throughout this section we assume that T is an $\mathbb{R}$-tree with isometric F_N-action, and that $H : T \to T$ is a homothety with stretching factor $\lambda > 1$ that represents some automorphism $\Phi \in \mathrm{Aut}(F_N)$. For the proof of Proposition 6.5 below one does not need to assume that the tree T is minimal, nor that T satisfies BBT.

We denote by $\overline{T}$ the metric completion of the tree T, and by $\overline{H} : \overline{T} \to \overline{T}$ the well defined continuous extension of H to $\overline{T}$. Clearly $\overline{H}$ is also a homothety with stretching factor λ which represents Φ.

Furthermore, let $f : \mathcal{G} \to \mathcal{G}$ be an β-train-track map that represents the outer automorphism $\varphi \in \mathrm{Out}(F_N)$ defined by Φ. As before, we denote by $\widehat{f} : \widehat{\mathcal{G}} \to \widehat{\mathcal{G}}$ the lift of f to the F_N-covering $\widehat{\mathcal{G}}$ of $\mathcal{G}$ which represents Φ.

We first assemble some useful facts:

Fact 6.2. *Let $\overline{H} : \overline{T} \to \overline{T}$ be as above, and consider any element $w \in F_N$ and any integer $m \geq 1$. It follows that $w\overline{H}^m : \overline{T} \to \overline{T}$ has a fixed point*

$$Q(w, m) \in \overline{T}$$

which is uniquely determined by w and m, as the stretching factor of the homothety $w\overline{H}^m$ is equal to $\lambda^m > 1$.

This is a direct consequence of the fact that $\overline{T}$ is metrically complete and that H^{-1} is contracting.

Fact 6.3. *Since $\overline{H} : \overline{T} \to \overline{T}$ represents $\Phi \in \mathrm{Aut}(F_N)$, for every $w \in F_N$ one has $\overline{H}w = \Phi(w)\overline{H}$ and hence (through replacing w by $\Phi^{-1}(w)$ and through conjugation by $\overline{H}^{-1}$) also $\overline{H}^{-1}w = \Phi^{-1}(w)\overline{H}^{-1}$. Iterated applications of these formulae gives $\overline{H}^m w = \overline{H}^{m-1}\Phi(w)\overline{H} = \overline{H}^{m-2}\Phi^2(w)\overline{H}^2 = \ldots$ and hence*

$$\overline{H}^m w = \Phi^m(w)\overline{H}^m : \overline{T} \to \overline{T} \tag{6.3}$$

for any $m \in \mathbb{Z}$.

Similarly, since $\widehat{f} : \widehat{\mathcal{G}} \to \widehat{\mathcal{G}}$ also represents Φ, iterated application of the formula $\widehat{f}w = \Phi(w)\widehat{f}$ for arbitrary $w \in F_N$ gives

$$\widehat{f}^m w = \Phi(w)^m \widehat{f}^m : \widehat{\mathcal{G}} \to \widehat{\mathcal{G}}\,, \tag{6.4}$$

but here we have to assume $m \geq 0$, since f^{-1} is not well defined.

Fact 6.4. *The set $V(\mathcal{G})$ of vertices of $\mathcal{G}$ is finite. Let $V_0 \subset V(\mathcal{G})$ be the set of f-periodic vertices. Thus for some suitable integer $m \geq 1$ we obtain $f^m(P) = P$ for every $P \in V_0$ and $f^m(P) \in V_0$ for every $P \in V(\mathcal{G}) \smallsetminus V_0$.*

Recall from Definition-Remark 4.4 (2) that the action of F_N on $\widehat{\mathcal{G}}$ by covering translations is free. It follows that for every vertex P in the full lift $\widehat{V}_0 \subset \widehat{\mathcal{G}}$ of V_0 there is precisely one element $w_P \in F_N$ which satisfies (for m fixed from now on to be as in Fact 6.4):

$$w_P \widehat{f}^m(P) = P \tag{6.5}$$

We use Fact 6.2 to define a map $i : V(\widehat{\mathcal{G}}) \to \overline{T}$, given for any $P \in \widehat{V}_0$ through

$$i(P) := Q(w_P, m) \tag{6.6}$$

and for all other vertices $P \in V(\widehat{\mathcal{G}}) - \widehat{V}_0$ with $P_0 = \widehat{f}^m(P)$ through

$$i(P) = \overline{H}^{-m} i \widehat{f}^m(P) = \overline{H}^{-m}(Q(w_{P_0}, m))\,. \tag{6.7}$$

Proposition 6.5. *The above defined map $i : V(\widehat{\mathcal{G}}) \to \overline{T}$ satisfies*

(a) *$vi = iv$ for all $v \in F_N$, and*

(b) *$i\widehat{f} = \overline{H}i$.*

Proof. (I) We first consider the case $P \in \widehat{V}_0$, and let $P' = vP$. In order to prove (a) we have to show that $Q(w_{P'}, m) = vQ(w_P, m)$, for $m \geq 1$ as in Fact 6.4:

From the above equations (6.4) and (6.5) one derives $\widehat{f}^m(vP) = \Phi^m(v)\widehat{f}^m(P) = \Phi^m(v)w_P^{-1}P = \Phi^m(v)w_P^{-1}v^{-1}(vP)$. This shows $w_{P'} = vw_P\Phi^m(v^{-1})$. On the other hand, for $Q = Q(w_P, m)$ the definition of $Q(w_P, m)$ gives $w_P\overline{H}^m(Q) = Q$, and thus we derive from (6.3) that $\overline{H}^m(vQ) = \Phi^m(v)\overline{H}^m(Q) = \Phi^m(v)w_P^{-1}Q = \Phi^m(v)w_P^{-1}v^{-1}(vQ)$. This proves $Q(vw_P\Phi^m(v^{-1}), m) = vQ$ and hence $Q(w_{P'}, m) = vQ(w_P, m)$.

In order to prove (b) for the case $P \in \widehat{V}_0$ we have to show that $\overline{H}(Q(w_P, m)) = Q(w_{\widehat{f}(P)}, m)$:

From (6.5) and (6.4) it follows $\widehat{f}(P) = \widehat{f}(w_P\widehat{f}^m(P)) = \Phi(w_P)\widehat{f}^{m+1}(P) = \Phi(w_P)\ \widehat{f}^m(\widehat{f}(P))$ and hence $w_{\widehat{f}(P)} = \Phi(w_P)$. But then the definition of $Q(w_P, m)$ in Fact 6.2 together with (6.3) gives $w_{\widehat{f}(P)}\overline{H}^m(\overline{H}(Q(w_P, m))) = \Phi(w_P)\overline{H}(\overline{H}^m(Q(w_P, m))) = \Phi(w_P)\overline{H}(w_P^{-1}Q(w_P, m)) = \Phi(w_P)\ \Phi(w_P^{-1})\overline{H}(Q(w_P, m)) = \overline{H}(Q(w_P, m))$. This implies $\overline{H}(Q(w_P, m)) = Q(w_{\widehat{f}(P)}, m)$.

(II) Assume now $P \in V(\widehat{\mathcal{G}}) \smallsetminus \widehat{V}_0$, and let $P_0 = \widehat{f}^m(P) \in \widehat{V}_0$. In order to show (a) we use (in this order) the equalities (6.7), (6.4), the above derived equation (a) for P_0, (6.3) and finally again (6.7), in order to compute $i(vP) = \overline{H}^{-m}i\widehat{f}^m(vP) = \overline{H}^{-m}i\Phi^m(v)\widehat{f}^m(P) = \overline{H}^{-m}i(\Phi^m(v)P_0) = \overline{H}^{-m}\Phi^m(v)i(P_0) = v\overline{H}^{-m}i\widehat{f}^m(P) = vi(P)$.

In order to show (b) for P we use (b) for P_0 to compute $\overline{H}i(P) = \overline{H}\overline{H}^{-m}i\widehat{f}^m(P) = \overline{H}^{-m}\overline{H}i(P_0) = \overline{H}^{-m}i\widehat{f}(P_0) = \overline{H}^{-m}i\widehat{f}^{m+1}(P) = i\widehat{f}(P)$. The last equation, in case $\widehat{f}(P) \in V(\widehat{\mathcal{G}}) \smallsetminus \widehat{V}_0$, follows from the definition of the map i in (6.7), and in case $\widehat{f}(P) \in \widehat{V}_0$ it follows from equation (b) as proved above for vertices in $\widehat{V}_0$. ☐

Now all arguments are ready to prove the main result of this paper:

Proof of Theorem 1.2. We first observe that from the assumption $T \in \overline{\mathrm{cv}}_N$ it follows that T is minimal and that T satisfies BBT. From the assumption that T is φ-invariant and expanding it follows (see subsection 2.3) that there is a homothety $H : T \to T$ with stretching factor $\lambda > 1$

which represents some preimage $\Phi \in \mathrm{Aut}(F_N)$ of φ. Let $\widehat{f} : \widehat{\mathcal{G}} \to \widehat{\mathcal{G}}$ be the lift of f which also represents Φ.

We now replace T by its metric completion $\overline{T}$ and H by the canonical extension to $\overline{H}$. We then apply Proposition 6.5 to obtain an F_N-equivariant map $i : V(\widehat{\mathcal{G}}) \to \overline{T}$ which satisfies $i\widehat{f} = \overline{H}i$. Thus Proposition 6.1 shows that there is an eigenvector $\vec{v}^*$ for f with eigenvalue λ such that for the corresponding Perron-Frobenius tree $T^{\vec{v}^*}$ there is F_N-equivariant isometric embedding $j : T^{\vec{v}^*} \to \overline{T}$. But as T is the minimal subtree of $\overline{T}$, it follows that j induces an F_N-equivariant isometry onto T, so that we have:

$$T = T^{\vec{v}^*} \in \overline{\mathrm{cv}}_N$$

$\square$

7. Comments and outlook

7.1. The fixed point set $\mathfrak{T}(\varphi)$ and its linear structure.

The proof of our main result shows actually something slightly stronger than what is stated in Proposition 1.1 and in Theorem 1.2. In order to explain this, we define

$$\mathfrak{T}(\varphi) \subset \overline{\mathrm{cv}}_N$$

to be the set of all very small minimal $\mathbb{R}$-trees which are expandingly fixed by the automorphism $\varphi \in \mathrm{Out}(F_N)$. We will denote by $\mathbb{P}\mathfrak{T}(\varphi) \subset \overline{\mathrm{CV}}_N$ the image of $\mathfrak{T}(\varphi)$ under the projectivization, i.e. the set of all $[T] \in \overline{\mathrm{CV}}_N$ with $T\varphi = \lambda T$ for some $\lambda > 1$.

We also need to consider, for any β-train-track map $f : \mathcal{G} \to \mathcal{G}$, the non-negative cone $C(f) := \mathbb{R}_{\geq 0}^{\mathrm{Edges}^+(\mathcal{G})}$ and the subset $\mathrm{PF}(f) \subset C(f)$ which consists of all Perron-Frobenius eigenvectors $\vec{v}^*$ for f with eigenvalue $\lambda > 1$. Projectivization gives $\mathbb{P}\mathrm{PF}(f) \subset \mathbb{P}C(f) \subset \mathbb{P}\mathbb{R}_{\geq 0}^{\mathrm{Edges}^+(\mathcal{G})}$.

Corollary 7.1. *For any β-train-track representative $f : \mathcal{G} \to \mathcal{G}$ of any $\varphi \in \mathrm{Out}(F_N)$ the association*

$$\vec{v}^* \mapsto T^{\vec{v}^*}$$

defines a bijective linear map

$$\mathfrak{b} : \mathrm{PF}(f) \to \mathfrak{T}(\varphi)$$

which quotients to the bijection

$$\mathfrak{B} : \mathbb{P}\mathrm{PF}(f) \to \mathbb{P}\mathfrak{T}(\varphi)$$

Proof. The existence and well-definedness of $\mathfrak{b}$ and $\mathfrak{B}$ has been proved in Theorem 5.2, and the injectivity as well as the linearity is a direct consequence of its statement (4). The surjectivity is the content of Theorem 1.2. $\square$

We'd like to add a word of caution here: From the above proposition we deduce that β-train-track representatives f_+ of φ and f_- of φ^{-1} serve to determine precisely all expanding and all contracting fixed points of any automorphism $\varphi \in \mathrm{Out}(F_N)$ in the boundary $\partial\mathrm{CV}_N$ of Outer space. However, in general the automorphism φ may possesses further fixed points $[T] \in \partial\mathrm{CV}_N$ which are neither expanding nor contracting: Such "neutral fixed points" can even exist for automorphisms that admit absolute train track representatives with primitive transition matrix, but which are not iwip.

Remark 7.2. A particularly easy way to construct such an automorphism is given by taking an absolute train track representative $f_0 : \Gamma_0 \to \Gamma_0$ of some iwip automorphism $\varphi_0 \in \mathrm{Out}(F_{N-1})$ and glueing together two distinct vertices of Γ_0 that are both fixed by f_0. This gives an absolute train track map $f_1 : \Gamma_1 \to \Gamma_1$ which then represents some "outer" endomorphism φ_1 of F_N; the latter is easily seen to be an automorphism $\varphi_1 \in \mathrm{Out}(F_N)$.

A useful description of the "neutral" fixed point for φ_1 is obtained by going back to Γ_0 and adding an extra edge e, which connects the two fixed points in question and is itself fixed by f. The graph-of-groups decomposition $\mathfrak{G}$ of F_N, obtained from this augmented graph through contracting all edges but e to a single vertex and providing this vertex with $\pi_1\Gamma_0$ as vertex group, describes precisely (through passing to the associated Bass-Serre tree $T_{\mathfrak{G}}$) the "neutral" fixed point $[T_{\mathfrak{G}}]$ of φ_1 in $\partial \mathrm{CV}_N$. The above "1-edge/1-vertex" graph which underlies $\mathfrak{G}$ shows up again below as "decomposition graph" in Remark 7.9.

7.2. Non-ergometric, dusted and other decomposable trees in $\mathfrak{T}(\varphi)$ and its convex hull.

In this subsection we will use Corollary 7.1 to analyze the set $\mathfrak{T}(\varphi)$ and its elements further. We first recall the embedding

$$\overline{\mathrm{cv}}_N \to \mathbb{R}_{\geq 0}^{F_N}\,, \quad T \mapsto ||\cdot||_T$$

from subsection 2.3 and note that for arbitrary trees $T_1, T_2 \in \overline{\mathrm{cv}}_N$ a non-negative linear combination of the two translation length functions $||\cdot||_{T_1}$ and $||\cdot||_{T_2}$ in $\mathbb{R}_{\geq 0}^{F_N}$ will in general not define an $\mathbb{R}$-tree length function. An exception is given if T_1 and T_2 are based on the same topological tree, or equivalently, if T_1 and T_2 determine the same dual algebraic lamination (see [6] for details). In this case we denote the resulting tree by

$$T = \lambda_1 T_1 + \lambda_2 T_2\,,$$

with λ_1 and λ_2 given through $||\cdot||_T = \lambda_1||\cdot||_{T_1} + \lambda_2||\cdot||_{T_2}$.

Definition 7.3 (see [6])**.** (1) A tree $T \in \overline{\mathrm{cv}}_N$ is called *ergometric* if for any linear combination

$$T = \lambda_1 T_1 + \lambda_2 T_2 \tag{7.1}$$

with $T_1, T_2 \in \overline{\mathrm{cv}}_N$ and $\lambda_1, \lambda_2 > 0$ the three trees T, T_1 and T_2 are equal up to rescaling: for some $\lambda > \lambda_1$ one has $T = \lambda T_1 = \frac{\lambda\,\lambda_2}{\lambda - \lambda_1} T_2$.
(2) In the complementary case, i.e. if T can be written as positive linear combination as in (7.1) with $[T_1] \neq [T_2]$, we say that T is *non-ergometric.*

Since arbitrarily chosen trees $T_1, T_2 \in \overline{\mathrm{cv}}_N$ do in general not admit convex combinations in $\overline{\mathrm{cv}}_N$, non-ergonometric trees have to be considered as exceptional. This gives a certain interest to the linearity of the map $\mathfrak{b}$ from Corollary 7.1, as it is fairly easy to construct examples of β-train-track maps $f : \mathcal{G} \to \mathcal{G}$ that admit such convex combinations based on points in the subset $\mathrm{PF}(f)$ of the non-negative cone $\mathcal{C}(f)$.

Remark 7.4. Non-ergometric trees T show up in the context of the set $\mathfrak{T}(\varphi)$ as a composition of the following two basic types:

(a) The tree T is the non-trivial convex combination $\lambda T^{\vec{v}_1^*} + (1-\lambda) T^{\vec{v}_2^*}$ of two PF-trees $T^{\vec{v}_1^*}$ and $T^{\vec{v}_2^*}$ which have both the same stretching factor. In this case T as well as the whole line segment $[T^{\vec{v}_1^*}, T^{\vec{v}_2^*}] \subset \overline{\mathrm{cv}}_N$ belong to $\mathfrak{T}(\varphi)$.

(b) There is a tree $T^{\vec{v}^*}$ in $\mathfrak{T}(\varphi)$ which is "dusted" (= "pointillé"), see [13]. In this case there is a second PF-tree $T^{\vec{v}_0^*} \in \mathfrak{T}(\varphi)$, which is simpler in that it is not dusted, and $T^{\vec{v}_0^*}$ has strictly smaller stretching factor than $T^{\vec{v}^*}$. The whole segment $[T^{\vec{v}^*}, T^{\vec{v}_0^*}]$ belongs to $\overline{\mathrm{cv}}_N$ and is fixed by φ, but not pointwise: under iteration of φ all points of this segment except for the non-dusted endpoint $T^{\vec{v}_0^*}$ converge to the dusted endpoint $T^{\vec{v}^*}$. Thus only the endpoints of the segment belong to $\mathfrak{T}(\varphi)$. Every tree T in the interior of the segment $[T^{\vec{v}^*}, T^{\vec{v}_0^*}] \subset \overline{\mathrm{cv}}_N$ is non-ergometric, and also dusted. The two endpoints $T^{\vec{v}^*}$ and $T^{\vec{v}_0^*}$ are both ergometric (as they are "principal", see below), so that $T^{\vec{v}^*}$ is an example for the remarkable case of a dusted but ergometric $\mathbb{R}$-tree.

The general case is of course more complex, and the issuing dynamics of φ on the "convex hull" of $\mathfrak{T}(\varphi)$ has some unexpected, rather interesting features.

A more complete understanding of the set $\mathfrak{T}(\varphi)$ is obtained when passing via Corollary 7.1 (1) to the subset $\text{PF}(f)$ of the non-negative cone $\mathcal{C}(f)$. For this purpose one has to consider the transition matrix $M(f)$ and apply some basic Perron-Frobenius theory (see [1]) or [22]).

Remark 7.5. (1) Let M be a non-negative integer square matrix. After a suitable permutation of the coordinates there is a partition of the coordinates in *canonical coordinate blocks* B_k, such that the issuing subdivision of M into matrix blocks gives a lower triangular block matrix, with the property that each diagonal matrix block $M_{k,k}$ is irreducible.

Here *irreducible* means that $M_{k,k}$ is either the 1×1 zero-matrix, or a (non-splittable) permutation matrix, or else it is *expanding*, i.e. its PF-eigenvalue (= the spectral radius) satisfies $\lambda_k > 1$, in which case we also say the the corresponding coordinate block B_k is expanding. Furthermore, an expanding irreducible block $M_{k,k}$, has always a positive power f^p which is a block diagonal matrix where all coordinates of any diagonal block are positive. If there is only one such diagonal block, then $M_{k,k}$ is called *primitive.*

The lower triangular block structure of M defines a natural partial order $B_h \preceq B_k$ among the coordinate blocks, which is generated by all non-zero off-diagonal blocks $M_{k,h}$ of M.

A coordinate block B_k is *distinguished* if and only if $\lambda_k > \lambda_h$ for all blocks $B_h \neq B_k$ with $B_h \preceq B_k$.

For each distinguished block B_k there is a special eigenvector $\vec{v}_k$ of M which is called *principal* if $M_{k,k}$ is primitive, and *barycentric principal* otherwise. A vector $\vec{v} \neq \vec{0}$ is a Perron-Frobenius eigenvector of M (i.e. all coordinates of $\vec{v}$ are non-negative) if and only if $\vec{v}$ is a non-negative linear combination of such (barycentric) principal $\vec{v}_k$ which all have the same eigenvalue $\lambda \geq 0$.

In Appendix A.3 of [1] some more details are assembled in a fairly accessible survey presentation.

(2) We will term a block B_k *polynomial* if no block B_h with $B_h \preceq B_k$ is expanding. For our purposes it is convenient to assemble all coordinates that belong to a polynomial block into a single block B_0. Of course, the corresponding diagonal block $M_{0,0}$ will in general not be irreducible. By definition the block B_0 is minimal with respect to the partial order (also denoted by $\preceq$) that is induced by $\preceq$ on the "quotient partition" obtained from assembling all polynomial blocks into the single block B_0.

We now apply this terminology to the transition matrix $M(f)$ of any β-train-track map $f : \mathcal{G} \to \mathcal{G}$:

Lemma 7.6. *(1) Let $\mathcal{G}_0 \subset \mathcal{G}_1 \subset \ldots \subset \mathcal{G}_s = \mathcal{G}$ be the filtration of f given by Definition 4.3 (5). Then there is a 1-1 correspondence between the strata $\mathcal{G}_k \smallsetminus \mathcal{G}_{k-1}$ from this filtration on the one hand, and the non-polynomial canonical coordinate blocks B_k of $M(f)$ on the other.*

More concretely, for any $k \geq 1$ the coordinates of B_k are given precisely by the edges that belong to $\mathcal{G}_k - (\mathcal{G}_{k-1} \cup X)$.

(2) The total order $\mathcal{G}_0 \subset \mathcal{G}_1 \subset \ldots \subset \mathcal{G}_s = \mathcal{G}$, when translated via the above 1 - 1 correspondence into a total order of the coordinate blocks of $M(f)$, is a refinement of the natural partial order $\preceq$ on the coordinate blocks defined in Remark 7.5 (1) and (2).

Proof. This is a direct consequence of the facts and conventions from Remark 7.5, together with the "no refinement" condition in part (5) of Definition 4.3. □

Since in the context of Perron-Frobenius $\mathbb{R}$-trees we have to work with row eigenvectors of $M(f)$, we need to first transpose the matrix and then apply the facts and the terminology assembled in Remark 7.5. This means that for $M(f)$ we have to replace the partial order $\preceq$ by its reversion $\succeq$, which leads of course to very different distinguished coordinate blocks.

It is often convenient to raise φ to a positive power φ^p such that for every expanding distinguished block B_k of $M(f^p) = M(f)^p$ the corresponding diagonal block $M_{k,k}$ is primitive. In this case the expanding distinguished blocks B_k (with respect to the reversed order $\succeq$) are in 1-1 correspondence with the principal eigenvectors $\vec{v}_k^*$ of $M(f^p)$ with eigenvalue $\lambda > 1$. We call the corresponding trees $T^{\vec{v}_k^*}$ *principal Perron-Frobenius trees* for φ. From the following proposition it follows that, up to uniform rescaling, these $T^{\vec{v}_k^*}$ depend only on φ and not on the particular β-train-track representative f.

Proposition 7.7. *For every automorphism $\varphi \in \mathrm{Out}(F_N)$ and any β-train-track representative $f : \mathcal{G} \to \mathcal{G}$ of φ there is an exponent $p \geq 1$ such that the following holds:*

(1) There are finitely many principal PF-trees $T^{\vec{v}_k^*}$ in $\mathfrak{T}(\varphi^p)$, and every other $T \in \mathfrak{T}(\varphi^p)$ is a non-negative linear combination of the latter.

More precisely, a positive linear combination of some principal PF-trees $T^{\vec{v}_k^*}$ belongs to $\mathfrak{T}(\varphi^p)$ if and only if all those $T^{\vec{v}_k^*}$ have the same stretching factor $\lambda > 1$.

(2) The automorphism φ permutes the projectivized principal PF-trees $[T^{\vec{v}_k^*}] \in \mathbb{P}\mathfrak{T}(\varphi)$ for φ^p and fixes for each permutation-orbit its barycentric convex combination. Every $T \in \mathfrak{T}(\varphi)$ is a non-negative linear combination of those "barycentric principal PF-trees", and any positive linear combination of some of the latter belongs to $\mathfrak{T}(\varphi^p)$ if and only if these barycentric principal PF-trees have the same stretching factor.

Proof. This is a direct consequence of Corollary 7.1, as the equivalent statement holds for the set $\mathrm{PF}(f)$. Indeed, the exponent $p \geq 1$ has been exhibited in Definition-Remark 11.8 and Remark 11.9 of [1]; in Proposition 11.11 of [1] it has been shown that for every distinguished coordinate block there exists "principal eigenvector" with the properties stated in part (1) of our claim, and in Definition-Remark 11.13 "barycentric principal eigenvectors" with the properties from part (2) have been exhibited. Since the eigenvectors in question are row eigenvectors, we have to transpose $M(f)$ before applying the above named quotes. ☐

From this proposition we see that the principal PF-trees determine extremal points of $\mathbb{P}\mathfrak{T}(\varphi^p)$. From the finite dimensionality of $\overline{\mathrm{cv}}_N$ (see [12]) and general facts from ergodicity theory one can hence deduce that the principal PF-trees are ergometric, while statement (1) of Proposition 7.7 shows that all other points of $\mathbb{P}\mathfrak{T}(\varphi^p)$ determine non-ergometric $\mathbb{R}$-trees. Going back to Remark 7.4, we see that these non-ergometric PF-trees are all of the type (a) described there. However, the non-ergometric trees from type (b) in the convex hull of $\mathbb{P}\mathfrak{T}(\varphi^p)$ can also be seen directly from $\mathrm{PF}(f)$ and the linear bijection $\mathfrak{b}$, namely as follows:

Using again the partial order $\succeq$ for the coordinate blocks of $M(f)$, or rather the reversed order $\preceq$ for the transpose ${}^tM(f)$, one sees directly (see [1], Proposition 11.11 (3a)) that for every non-minimal distinguished expanding block B_k of ${}^tM(f^p)$ there is a minimal expanding block B_h (or possibly several ones) such that the corresponding (barycentric) principal eigenvectors satisfy $\lambda\vec{v}_h^* < \vec{v}_k^*$ (meant coordinate-wise) for some $\lambda > 0$. From property (4) of Theorem 5.2 one can then derive that the tree $T^{\vec{v}_k^*}$ is dusted, while the "minimal-matrix-block" PF-tree $T^{\vec{v}_h^*}$ isn't, and that any positive linear combination of $T^{\vec{v}_k^*}$ and $T^{\vec{v}_h^*}$ defines a non-ergometrix $\mathbb{R}$-tree in $\overline{\mathrm{cv}}_N$ (albeit not in $\mathfrak{T}(\varphi)$!).

"Dustedness" is only one of the two basic phenomena when trying to decompose an $\mathbb{R}$-tree as combination of simpler $\mathbb{R}$-trees (see [13]). In fact, one should not think that the non-dusted principal PF-trees $T^{\vec{v}_h^*}$ are indecomposable, as the other basic decomposition principal may still apply: in general a minimal-matrix-block PF-tree $T^{\vec{v}_h^*}$ as above may well be a "graph-of-actions", which can be seen as follows:

Definition-Remark 7.8. (1) From Remark 7.5 (1) it follows that every coordinate block which is minimal among the expanding coordinate blocks of ${}^tM(f)$ is automatically distinguished. It follows that for any coordinate block B_h which is maximal among the expanding coordinate blocks of $M(f)$ there is always an associated principal (or barycentric principal) PF-tree $T^{\vec{v}_h^*}$.

(2) Any such "maximal" B_h defines an *augmented expanding top stratum* B_h^* of f, which consists precisely of those edges of $\mathcal{G}$ that belong to coordinate blocks which are bigger or equal to B_h with respect to the natural partial order $\preceq$ (and are hence not expanding, if distinct from B_h, since B_h is maximal among the expanding coordinate blocks).

(3) Since they play an important role, we will use the term *expanding top tree* for the principal or barycentric principal PF-tree $T^{\vec{v}_h^*}$ that correspond to the above considered "maximal" B_h. The tree $T^{\vec{v}_h^*}$ is always non-dusted, and it is ergometric if and only if it is principal and not just barycentric principal.

It is hence quite possible that the stratum of f, which corresponds to an expanding block B_h that defines some given expanding top tree $T^{\vec{v}_h^*}$, is not the top stratum of f, in that there are higher-up strata which define periodic (or zero) diagonal blocks B_k of the matrix $M(f)$ with $B_h \preceq B_k$. If $\mathcal{G}$ has no vertices of valence 1, then such strata always define a refined decomposition of $T^{\vec{v}_h^*}$ as a graph of actions. However, the converse conclusion will in general fail, so that one needs to be very careful here:

Remark 7.9. (1) Let $f : \mathcal{G} \to \mathcal{G}$ be a β-train-track representative of φ, and assume that the top stratum of $\mathcal{G}$ (with respect to the total order on the strata from Definition 4.3 (5)) defines an irreducible diagonal block B_s of $M(f)$ which is expanding. Let $T^{\vec{v}_s^*}$ be the expanding top tree which is associated to B_s. Let us furthermore assume that $f = f^p$ for the exponent $p \geq 1$ from Proposition 7.7, i.e. "principal" and "barycentric principal" coincide, so that $T^{\vec{v}_s^*}$ is a principal PF-tree for f.

Even then, the tree $T^{\vec{v}_s^*}$ will in general *not* be indecomposable: An example is given by the absolute train track map f_1 described in Remark 7.2, where $M(f_1)$ consists only of a single expanding diagonal block, while from the construction of f_1 we see that f_1 decomposes as graph of actions: the decomposition graph is a loop with a single vertex, and the vertex tree is given by the (up to uniform rescaling unique) PF-tree of the iwip automorphism φ_0 defined by the original train track map f_0.

(2) The above automorphism φ_1 a typical example for a class of automorphisms that are well known among the experts for iwip automorphisms: Every iwip automorphism $\varphi \in \mathrm{Out}(F_N)$ possesses an absolute train track representative $f : \Gamma \to \Gamma$ without polynomially growing edges and with transition matrix $M(f)$ that is primitive.

On the other hand, an arbitrary graph map $f : \Gamma \to \Gamma$, with primitive transition matrix $M(f)$ and without polynomially growing edges, does in general *not* represent an iwip automorphism: In order to conclude that φ is iwip the map f also needs to satisfy a "connectedness condition" on the Whitehead graph of each vertex, which turns out (see section 8 of [7]) to be equivalent to the condition that the PF-tree $T^{\vec{v}^*}$ determined by the (up to scaling unique) PF-eigenvector $\vec{v}^*$ for f is indecomposable.

The reason for this apparent "incoherence" is that the expanding diagonal blocks of $M(f)$ are in a strong, well-defined sense (see subsection 7.3) structural invariants of φ, while the appearance or non-appearance of possible non-expanding diagonal blocks is accidental and depends on the particular β-train-track map chosen to represent φ.

Remark 7.10. It turns out that for many algorithmic purposes the indecomposability of the expanding top tree (or of the more general principal PF-trees) is not relevant. For example, the solution of the conjugacy problem for iwip automorphisms $\varphi \in \mathrm{Out}(F_N)$ presented in [18], which is crucially based on the use of the unique expanding φ-invariant $\mathbb{R}$-tree, extends fairly directly to the more general situation given by any of the above defined augmented expanding top strata of f, together with its unique associated expanding top tree.

This leads (see [17]) to a "relative solution" of the conjugacy problem for arbitrary $\varphi, \varphi' \in \mathrm{Out}(F_N)$ that are not polynomially growing, as for such automorphisms there is at least one augmented expanding top stratum in any β-train-track representatives $f : \mathcal{G} \to \mathcal{G}$ and $f' : \mathcal{G}' \to \mathcal{G}'$ of φ and φ' respectively, and hence at least one expanding top tree $T^{\vec{v}^*}$.

In fact, the technique from [18] allows one to deduce from the assumption of a conjugating automorphism ψ the existence of a conjugating graph map $h : \mathcal{G} \to \mathcal{G}'$ with upper bounds for the top-stratum-length of the h-image for any edge of $\mathcal{G}$, which can be derived algorithmically from f and f'. In addition we deduce a bijection from the subgraphs of $\mathcal{G}$ to those of $\mathcal{G}'$ which are non-trivial connected components of the complement of the augmented expanding top strata (as they represent the non-trivial branch point stabilizers of $T^{\vec{v}^*}$, see Remark 7.11 (3)). On these subgraphs the conjugacy of the restrictions of φ and φ' has yet to be decided, but the rank is now strictly smaller than N, so that one can put in place an induction procedure.

The relative finiteness of the centralizer of φ shown in [18] translates here into finitely many possibilities of how the attaching data (see [17]) from the augmented expanding top stratum are mapped to β-train-track representatives of the branch point stabilizers of $T^{\vec{v}^*}$, modulo powers of the restricted automorphisms. In the long run this will lead to a normal form for arbitrary outer automorphisms of F_N.

7.3. The set $\mathcal{U}(\varphi)$ of characteristic subgroups.

In this subsection we will explain the relevance of the set $\mathbb{P}\mathfrak{T}(\varphi)$ of projectively invariant expanding $\mathbb{R}$-trees for a structural understanding of any given automorphism $\varphi \in \mathrm{Out}(F_N)$. The fundamental approach here is that, by definition, the sets $\mathfrak{T}(\varphi)$ and $\mathbb{P}\mathfrak{T}(\varphi)$ are invariants of φ: If $\varphi' \in \mathrm{Out}(F_N)$ is conjugate to φ, say through $\varphi' = \psi^{-1}\,\varphi\,\psi$ for some $\psi \in \mathrm{Out}(F_N)$, then one has:

$$\mathfrak{T}(\varphi') = \mathfrak{T}(\varphi)\psi \quad \text{and} \quad \mathbb{P}\mathfrak{T}(\varphi') = \mathbb{P}\mathfrak{T}(\varphi)\psi$$

On the other hand, a given automorphism φ can have many distinct β-train-track representatives $f : \mathcal{G} \to \mathcal{G}$, and it is a priori not at all clear to what extend for example the strata structure of such (or other types of) train track maps is an invariant of φ. Hence the maps $\mathfrak{b}$ and $\mathfrak{B}$ from Proposition 7.1 play a fundamental role to deduce from $\mathfrak{T}(\varphi)$ and $\mathbb{P}\mathfrak{T}(\varphi)$ data of f that are structural invariants of φ. For this purpose we first note:

Remark 7.11. (1) Let $T \in \overline{\mathrm{cv}}_N$ be an $\mathbb{R}$-tree with trivial arc stabilizers. Then every point $x_i \in T$ with non-trivial stabilizer $\mathrm{stab}(x_i) \subset F_N$ is a branch point, and stabilizers of distinct points have trivial intersection. It is well known (see [12]) that there is only a finite number of F_N-orbits of points x_i with $\mathrm{stab}(x_i) \neq \{1\}$ (indeed less then $2N - 1$), and that each $\mathrm{stab}(x_i)$ has finite rank (indeed bounded above by $N - 1$). Since for points x_i in the same F_N-orbit the stabilizers $\mathrm{stab}(x_i)$ are conjugate in F_N, such a tree T determines a finite set $\mathcal{U}(T)$ of conjugacy classes $[U_i]$ of finitely generated non-trivial subgroups $U_i := \mathrm{stab}(x_i)$ in F_N. For more detail about such trees and its point stabilizers see [12].

(2) If we now pass to the special case of a PF-tree $T^{\vec{v}^*}$, then each conjugacy class $[U_i] \subset \mathcal{U}(T^{\vec{v}^*})$ is represented by a subgraph $\mathcal{G}_i$ of $\mathcal{G}$. Indeed, the vector $\vec{v}^* = (v_e)_{e \in \mathrm{Edges}^+(\mathcal{G})}$ determines a *zero-subgraph* $\mathcal{G}^{\vec{v}^*} \subset \mathcal{G}$, which consists of all edges e with $\vec{v}^*$-coordinate $v_e = 0$. It follows directly from Theorem 5.2 (4) that every connected component $\mathcal{G}_i$ of $\mathcal{G}^{\vec{v}^*}$, when lifted to a connected component of $\widehat{\mathcal{G}}$, is mapped by $i^{\vec{v}^*}$ to the F_N-orbit of some point $x_i \in T^{\vec{v}^*}$, and that $\mathrm{stab}(x_i)$ is conjugate in F_N to the *zero-subgroup* $\theta(\pi_1 \mathcal{G}_i)$. If $\theta(\pi_1 \mathcal{G}_i) = \{1\}$ we say that the component $\mathcal{G}_i$ is *trivial*; otherwise $\mathcal{G}_i$ is called *non-trivial*, and in this case x_i is a branch point of T.

(3) This defines a bijection from the set of non-trivial connected components of the zero-subgraph $\mathcal{G}^{\vec{v}^*}$ to the set $\mathcal{U}(T^{\vec{v}^*})$. Since $T^{\vec{v}^*}$ is projectively φ-invariant, it follows that φ preserves $\mathcal{U}(T^{\vec{v}^*})$

by permuting the conjugacy classes of the zero-subgroups $U_i = \theta(\pi_1 \mathcal{G}_i)$. Correspondingly, the zero-subgraph $\mathcal{G}^{\vec{v}^*}$ is f-invariant, and f permutes its non-trivial connected components.

In light of this remark, we now obtain the following invariant:

Definition-Remark 7.12. (1) For any $\varphi \in \mathrm{Out}(F_N)$ we define the set

$$\mathcal{U}(\varphi) = \cup\{\mathcal{U}(T) \mid T \in \mathfrak{T}(\varphi)\}$$

of conjugacy classes of *characteristic subgroups* U_i of φ.
(2) For any β-train-track representative $f : \mathcal{G} \to \mathcal{G}$ we have

$$\mathcal{U}(\varphi) = \cup\{\mathcal{U}(T^{\vec{v}^*}) \mid \vec{v}^* \in \mathrm{PF}(f)\}$$

(3) Every conjugacy class $[U_i]$ in $\mathcal{U}(\varphi)$ is represented by a *characteristic zero-subgraph* $\mathcal{G}_i$ of $\mathcal{G}$ via $U_i = \theta(\pi_1 \mathcal{G}_i)$, where $\mathcal{G}_i$ is a non-trivial connected component of the zero-subgraph $\mathcal{G}^{\vec{v}^*}$ of some Perron-Frobenius eigenvector $\vec{v}^*$ for f:

$$\mathcal{U}(\varphi) = \{[U_i] \mid U_i = \theta(\pi_1(\mathcal{G}_i)) \neq \{1\}, \mathcal{G}_i \in \pi_0(\mathcal{G}^{\vec{v}^*}), \vec{v}^* \in \mathrm{PF}(f)\}$$

Since $\mathcal{G}$ has only finitely many subgraphs $\mathcal{G}'$, the set of those which are invariant under f (or under powers of f) can be readily computed. Not any such invariant subgraph $\mathcal{G}'$, however, is a connected component of the zero-subgraph $\mathcal{G}^{\vec{v}^*}$ of some Perron-Frobenius eigenvector $\vec{v}^*$. In order to distinguish those, one has to consider the transition matrix $M(f)$ and apply the canonical block decomposition together with the natural partial order $\preceq$ from Remark 7.5. From the expanding distinguished coordinate blocks of ${}^tM(f)$ and the associated finitely many principal eigenvectors we obtain hence directly the computable conjugacy invariants exhibited in Corollary 1.3, stated here in a more explicit form:

Corollary 7.13. *For any automorphisms $\varphi \in \mathrm{Out}(F_N)$ we consider the set $\mathcal{U}(\varphi)$ of conjugacy classes $[U_i]$ characteristic subgroups U_i.*
(1) The set $\mathcal{U}(\varphi)$ is finite, and for each of its elements $[U_i]$ the rank is finite:

$$\mathrm{rk\,U_i \leq N - 1}$$

(2) The set $\mathcal{U}(\varphi)$ is a conjugacy invariant of φ: For any $\psi \in \mathrm{Out}(F_N$ one has:

$$\mathcal{U}(\psi \circ \varphi \circ \psi^{-1}) = \psi \mathcal{U}(\varphi)$$

(3) The set $\mathcal{U}(\varphi)$ can be determined algorithmically: A finite generating set of a representative subgroup U_i, for each element of $\mathcal{U}(\varphi)$, can be derived (essentially by hand) from any β-train-track representative $f : \mathcal{G} \to \mathcal{G}$ of φ.
(4) The inclusion in F_N induces a natural partial order on $\mathcal{U}(\varphi)$, such that the following holds:
The elements of the maximal classes $[U_j]$ are precisely the point stabilizers of the finitely many expanding top trees $T^{\vec{v}^*} \in \mathfrak{T}(\varphi)$ from Definition-Remark 7.8 (3): The zero-subgraph $\mathcal{G}^{\vec{v}^*}$ of such a PF-eigenvector $\vec{v}^*$ is characterized by the following properties:
(1) $\mathcal{G}^{\vec{v}^*}$ *contains all but one of the expanding strata of the β-train-track map $f : \mathcal{G} \to \mathcal{G}$,*
(2) $\mathcal{G}^{\vec{v}^*}$ *is f-invariant, and*
(3) $\mathcal{G}^{\vec{v}^*}$ *is maximal with respect to the properties (1) and (2).* □

In order to draw the connection with the previous subsection more closely, we'd like to note here that for any expanding top tree $T^{\vec{v}^*}$ the complement of the zero-subgraph $\mathcal{G}^{\vec{v}^*}$ in $\mathcal{G}$ gives precisely the augmented expanding top stratum associated to the principal eigenvector $\vec{v}^*$.

Remark 7.14. This last observation is the starting point of a structural decomposition of $\mathcal{G}$ into f-invariant, partially ordered *expanding coarse strata* $\mathcal{H}_j$ by an iterative procedure "from the top", i.e. starting with f and $\mathcal{G}$. Here each stratum $\mathcal{H}_j$ is an augmented expanding top stratum of the restriction of f (or of a positive power of f) to some connected component of the zero-subspace of some expanding top tree previously obtained in our iterative procedure.

The φ-invariant collection $\mathcal{U}^*(\varphi)$ of conjugacy classes of subgroups that results from this expanding-coarse-strata decomposition of $\mathcal{G}$ is slightly richer (but also more complicated to describe) than the collection $\mathcal{U}(\varphi)$ from Corollary 7.13, since every expanding coordinate block B_k for $M(f)$ (i.e. not just the distinguished ones) gives rise to a subgroup in this collection: the expanding coarse stratum $\mathcal{H}_k$ associated to the expanding block B_k consists of all blocks B_h such that B_k is maximal among the set of all expanding blocks $B_{k'}$ which satisfy $B_{k'} \preceq B_h$.

Each expanding coarse stratum $\mathcal{H}_k$ defines a possibly non-connected *zero-subgraph* $\mathcal{G}^{\mathcal{H}_k} \subset \mathcal{G}$, which consists of all edges that belong to any stratum $B_h \neq B_k$ with $B_h \preceq B_k$. Here each polynomial coordinate block B_h has to be treated individually (i.e. without assembling all of them into B_0 as in Remark 7.5 (2)). Each non-trivial connected component $\mathcal{G}_j^*$ of $\mathcal{G}^{\mathcal{H}_k}$ gives rise to an element $[U_j^*] = [\theta(\pi_1\mathcal{G}_j)] \in \mathcal{U}^*(\varphi)$.

The importance of Corollary 7.13 (or of the expanding-coarse-strata decomposition from Remark 7.14) is emphasized by the following warning:

Remark 7.15. (1) While the decomposition of F_N into the characteristic subgroups from $\mathcal{U}(\varphi)$ or from $\mathcal{U}^*(\varphi)$, derived from the expanding distinguished coordinate blocks B_k or the coarse expanding strate $\mathcal{H}_j$ respectively, are structural invariants of φ, the analogous statement is not true for the totally ordered strata decomposition of $\mathcal{G}$ from property (5) of Definition 4.3.

(2) Of course, the strata-decomposition of any a classical relative train track representative of φ has no claim to structural φ-invariance either, as they are indeed structurally weaker than β-train-track representatives. In fact, it is easy to construct examples of relative train track representatives of the same automorphism which differ in the number of strata and also in their order, when compared by checking the associated eigenvalue.

(3) To the best of our knowledge the same is true for any of the improved and improved-improved versions of relative train tracks that presently exist in the literature.

(4) The ever growing fan club of free factor systems, free factor complexes, etc may want to notice that the above exhibited structurally invariant characteristic subgroups from $\mathcal{U}(\varphi)$ or $\mathcal{U}^*(\varphi)$ are finitely generated and of infinite index, but in general *not* free factors of F_N.

(5) The only known structural invariants that are comparable to the characteristic subgroups from $\mathcal{U}(\varphi)$ or $\mathcal{U}^*(\varphi)$ are the expanding algebraic laminations, or perhaps more telling, the associated sets of attracting fixed currents that can be derived for example from an absolute train track representative of φ (see [21]). It is possible (but not at all obvious) to derive *carrier subgroups* for such laminations (see [15]), and there is indeed a strong relationship to the invariant subgroups $U_i^* = \theta(\pi_1\mathcal{G}_i^*)$ named above, but the resulting decompositions are not identical:

For example, consider the automorphism φ_1 and the absolute train track map $f_1 : \Gamma_1 \to \Gamma_1$ from Remark 7.2, derived from an absolute train track representative $f_0 : \mathcal{G}_0 \to \mathcal{G}_0$ of an iwip automorphism φ_0 through identifying two distinct vertices fixed by f_0. The characteristic subgroup set $\mathcal{U}(\varphi_1)$ from Corollary 7.13 is empty: There is only a single stratum which contains all of the graph Γ_1 (see Remark 7.9). The carrier subgroup of the (unique) expanding lamination, however, would not be the full group $\pi_1\Gamma_1$ but only the fundamental group $\pi_1\Gamma_0 \subsetneqq \pi_1\Gamma_1$ of the graph Γ_0 from before doing the above two-vertex identification.

References

[1] N. Bédaride, A. Hilion and M. Lustig, *Graph towers, laminations and their invariant measures.* J. London Math. Soc. (2) **95** (2020), 1–61

[2] M. Bestvina and M. Feighn, *Outer limits.* Preprint 1993 http://andromeda.rutgers.edu/ feighn/papers/outer.pdf

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

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

[5] M. Cohen and M. Lustig, *Very small group actions on R-trees and Dehn twist automorphisms.* Topology **34** (1995), 575–617

[6] T. Coulbois, A. Hilion, and M. Lustig, *Non-unique ergodicity, observers' topology and the dual algebraic lamination for* $\mathbb{R}$*-trees.* Illinois J. of Math. **51** (2007), 897–911

[7] T. Coulbois and M. Lustig, *Long turns, INP's and indices for free group automorphisms.* Illinois J. Math. **59** (2015), 1087–1109

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

[9] M. Culler and J. W. Morgan, *Group actions on R-trees.* Proc. Lond. Math. Soc. (3) **55** (1987), 571–604.

[10] W. Dicks and E. Ventura, *The group fixed by a family of injective endomorphisms of a free group.* Contemporary Mathematics **195**, Providence, RI: AMS, 1996

[11] D. Gaboriau, A. Jaeger, G. Levitt and M. Lustig, *An index for counting fixed points of automorphisms of free groups.* Duke Math. J. **93** (1998), 425–452

[12] D. Gaboriau and G. Levitt, *The rank of actions on* $\mathbb{R}$*-trees.* Ann. Sci. École Norm. Sup. (4) **28** (1995), 549–570

[13] V. Guirardel, *Actions of finitely generated groups on R-trees.* Ann. Inst. Fourier (Grenoble) 58 (2008), 159–211

[14] G. Levitt and M. Lustig, *Irreducible automorphisms of* F_n *have North-South dynamics on compactified outer space.* J. Inst. Math. Jussieu **2** (2003), 59–72

[15] M. Lustig, *Automorphismen von freien Gruppen.* Habilitationsschrift, Ruhr Universität Bochum, 1992

[16] M. Lustig, *Structure and conjugacy for automorphisms of free groups I.* Max-Planck Institut für Mathematik Preprint Series 2000, No 130

[17] M. Lustig, *Structure and conjugacy for automorphisms of free groups II.* Max-Planck Institut für Mathematik Preprint Series 2001, No 4

[18] M. Lustig, *Conjugacy and centralizers for iwip automorphisms of free groups.* In "Geometric Group Theory : Geneva and Barcelona Conferences", eds. G. Arzhantseva, L. Bartholdi, J. Burillo and E. Ventura, Trends in Mathematics, Birkhäuser 2007, 197–224

[19] M. Lustig, Appendix to *The mapping-torus of a free group automorphism is hyperbolic relative to the canonical subgroups of polynomial growth* (by F. Gautero and M. Lustig). arXiv:0707.0822

[20] M. Lustig, *Straight paths in* $\mathbb{R}$*-trees and train track maps for hyperbolic automorphisms of* F_N*.* Preliminary preprint 2017

[21] M. Lustig and C. Uyanik, *North-South dynamics of hyperbolic free group automorphisms on the space of currents.* J. Topol. Anal. **11** (2019), 427–466

[22] M. Lustig and C. Uyanik, *PerronFrobenius theory and frequency convergence for reducible substitutions.* Discrete Contin. Dyn. Syst. **37** (2017), 355–385

[23] M. Lustig and K. Ye, *Normal form and parabolic orbits for quadratically growing automorphisms of free groups.* arXiv:1705.04110v2

[24] K. Vogtmann, *Automorphisms of Free Groups and Outer Space.* Geometriae Dedicata **94** (2002), 1–31

Aix Marseille Université, CNRS, Centrale Marseille, I2M UMR 7373, 13453 Marseille, France
E-mail address: Martin.Lustig@univ-amu.fr