

The mapping torus group of a free group automorphism is hyperbolic relative to the canonical subgroups of polynomial growth

F. Gautero, M. Lustig

October 26, 2008

Abstract

We prove that the mapping torus group $\mathbb{F}_n \rtimes_{\alpha} \mathbb{Z}$ of any automorphism α of a free group \mathbb{F}_n of finite rank $n \geq 2$ is weakly hyperbolic relative to the canonical (up to conjugation) family $\mathcal{H}(\alpha)$ of subgroups of \mathbb{F}_n which consists of (and contains representatives of all) conjugacy classes that grow polynomially under iteration of α . Furthermore, we show that $\mathbb{F}_n \rtimes_{\alpha} \mathbb{Z}$ is strongly hyperbolic relative to the mapping torus of the family $\mathcal{H}(\alpha)$. As an application, we use a result of Drutu-Sapir to deduce that $\mathbb{F}_n \rtimes_{\alpha} \mathbb{Z}$ has Rapić Decay.

1 Introduction

Let \mathbb{F}_n be a (non-abelian) free group of finite rank $n \geq 2$, and let α be any automorphism of \mathbb{F}_n . It is well known (see [3] and [27]) that elements $w \in \mathbb{F}_n$ grow either at least exponentially or at most polynomially, under iteration of α . This terminology is slightly misleading, as in fact it is the translation length $\|w\|_{\mathcal{A}}$ of w on the Cayley tree of \mathbb{F}_n with respect to some basis \mathcal{A} that is being considered, which is the same as the word length in $\mathcal{A}^{\pm 1}$ of any cyclically reduced $w' \in \mathbb{F}_n$ conjugate to w .

There is a canonical collection of finitely many conjugacy classes of finitely generated subgroups H_1, \dots, H_r in \mathbb{F}_n which consist entirely of elements of polynomial growth, and which has furthermore the property that every polynomially growing element $w \in \mathbb{F}_n$ is conjugate to an element $w' \in \mathbb{F}_n$ that belongs to some of the H_i . In other words, the set of all polynomially growing elements of \mathbb{F}_n is identical with the union of all conjugates of the H_i . For more details see §3 below.

This *characteristic family* $\mathcal{H}(\alpha) = (H_1, \dots, H_r)$ is α -invariant up to conjugation, and in the mapping torus group

$$\mathbb{F}_n \rtimes_{\alpha} \mathbb{Z} = \langle x_1, \dots, x_n, t \mid tx_it^{-1} = \alpha(x_i) \text{ for all } i = 1, \dots, n \rangle$$

one can consider induced mapping torus subgroups $H_i^\alpha = H_i \rtimes_{\alpha^{m_i}} \mathbb{Z}$, where $m_i \geq 1$ is the smallest exponent such that $\alpha^{m_i}(H_i)$ is conjugate to H_i .

There is a canonical family \mathcal{H}_α of such mapping torus subgroups, which is uniquely determined, up to conjugation in $\mathbb{F}_n \rtimes_\alpha \mathbb{Z}$, by the characteristic family $\mathcal{H}(\alpha)$ (see Definition 2.8).

Theorem 1.1. *Let $\alpha \in \text{Aut}(\mathbb{F}_n)$, let $\mathcal{H}(\alpha) = (H_1, \dots, H_r)$ be the characteristic family of subgroups of polynomial α -growth, and let \mathcal{H}_α be its mapping torus. Then:*

- (1) $\mathbb{F}_n \rtimes_\alpha \mathbb{Z}$ is weakly hyperbolic relative to $\mathcal{H}(\alpha)$.
- (2) $\mathbb{F}_n \rtimes_\alpha \mathbb{Z}$ is strongly hyperbolic relative to \mathcal{H}_α .

Here a group G is called *weakly hyperbolic* relative to a family of subgroups H_i if the Cayley graph of G , with every left coset of any of the H_i coned off, is a δ -hyperbolic space (compare Definition 2.2). We say that G is *strongly hyperbolic* relative to (H_1, \dots, H_r) if in addition this coned off Cayley graph is *fine*, compare Definition 2.1. The concept of relatively hyperbolic groups originates from Gromov’s seminal work [22]. It has been fundamentally shaped by Farb [15] and Bowditch [8], and it has since then been placed into the core of geometric group theory in its most present form, by work of several authors, see for example [35], [9] and [34]. The relevant facts about relative hyperbolicity are recalled in §2 below.

As a consequence of our main theorem we derive the following corollary, using earlier results of Jolissaint [25] and Drutu-Sapir [14].

Corollary 1.2. *For every $\alpha \in \text{Aut}(\mathbb{F}_n)$ the mapping torus group $\mathbb{F}_n \rtimes_\alpha \mathbb{Z}$ satisfies the Rapid Decay property.*

The proof of this corollary, as well as definitions and background about the Rapid Decay property, are given below in §9.

Another consequence of our main theorem, pointed out to us by M. Bridson, is an alternative (and perhaps conceptually simpler) proof of the following recent result:

Theorem 1.3 (Bridson-Groves). *For every $\alpha \in \text{Aut}(\mathbb{F}_n)$ the mapping torus group $\mathbb{F}_n \rtimes_\alpha \mathbb{Z}$ satisfies a quadratic isoperimetric inequality.*

The proof of this result is given in a sequence of three long papers [4] [5] [6], where a non-trivial amount of technical machinery is developed. However, a first step is much easier: The special case of the above theorem where all of \mathbb{F}_n has polynomial α -growth (compare also [32]). It is shown by Farb [15] that, if a group G is strongly hyperbolic relative to a finite family of subgroups which all satisfy a quadratic isoperimetric inequality, then G itself satisfies a quadratic isoperimetric inequality. Thus, the special case of Bridson-Groves’ result, together with our Theorem 1.1, gives the full strength of Theorem 1.3.

This paper has several “predecessors”: The absolute case, where the characteristic family $\mathcal{H}(\alpha)$ is empty, has been proved by combined work of Bestvina-Feighn [2] (see also [19]) and

Brinkmann [7]. In [20] the case of geometric automorphisms of \mathbb{F}_n (i.e. automorphisms induced by surface homeomorphisms) has been treated. The methods developed there and in [19] have been further extended in [21] to give a general combination theorem for relatively hyperbolic groups (see also [33]). This combination theorem is a cornerstone in the proof of our main result stated above; it is quoted in the form needed here as Theorem 2.9.

The other main ingredient in the proof of Theorem 1.1 are β -train track representatives for free group automorphisms as developed by the second author (see Appendix), presented here in §4 and §5 below. These train track representatives combine several advantages of earlier such train track representatives, although they are to some extent simpler, except that their universal covering is not a tree.

The bulk of the work in this paper (§6 and §7) is devoted to make up for this technical disadvantage: We introduce and analyze *normalized paths* in β -train tracks, and we show that they can be viewed as proper analogues of geodesic segments in a tree.

In particular, we prove that in the universal covering of a β -train track

- (1) any two vertices are connected by a unique normalized path, and
- (2) normalized paths are quasi-geodesics (with respect to both, the absolute and the relative metric, see §7).

Normalized paths are useful in other contexts as well. In this paper they constitute the main tool needed to prove the following proposition.

The precise definition of a *relatively hyperbolic* automorphism is given below in Definition 2.6.

Proposition 1.4. *Every automorphism $\alpha \in \text{Aut}(\mathbb{F}_n)$ is hyperbolic relative to the characteristic family $\mathcal{H}(\alpha)$ of subgroups of polynomial α -growth.*

Acknowledgments. The first author would like to thank Martin Bridson for helpful and encouraging remarks. The second author would like to point out that some of the work presented here has also been inspired by his collaboration with Gilbert Levitt. Further thanks go to Université P. Cézanne and Université de Provence at Aix-Marseille and to the CIRM at Luminy for having supported during the residential session “Groups 007” in February 2007 a 4 week stay of the first author at Marseille.

2 Relative hyperbolicity

Let Γ be a connected, possibly infinite graph. We assume that every edge e of Γ has been given a length $L(e) > 0$. This makes Γ into a metric space. If Γ is locally finite, or if the edge lengths are chosen from a finite subset of \mathbb{R} , then Γ is furthermore a *geodesic* space, i.e. any two points are connected by a path that has as length precisely the distance between its endpoints.

Definition 2.1. A graph Γ is called *fine* if for every integer $n \in \mathbb{N}$ any edge e of Γ is contained in only finitely many circuits of length less or equal to n . Here a *circuit* is a closed edge path that passes at most once over any vertex of Γ .

Let G be a finitely generated group and let $S \subset G$ be a finite generating system. We denote by $\Gamma_S(G)$ the Cayley graph of G with respect to S . We define for every edge e the edge length to be $L(e) = 1$.

Let $\mathcal{H} = (H_1, \dots, H_r)$ be a finite family of subgroups of G , where in the context of this paper the H_i are usually finitely generated.

Definition 2.2. The \mathcal{H} -coned Cayley graph, denoted by $\Gamma_S^{\mathcal{H}}(G)$, is the graph obtained from $\Gamma_S(G)$ as follows:

1. We add an *exceptional vertex* $v(gH_i)$, for each coset gH_i of any of the H_i .
2. We add an edge of length $\frac{1}{2}$ connecting any vertex g of $\Gamma_S(G)$ to any of the exceptional vertices $v(gH_i)$.

We denote by $|\cdot|_{S, \mathcal{H}}$ the minimal word length on G , with respect to the (possibly infinite) generating system given by the finite set S together with the union of all the subgroups in \mathcal{H} . It follows directly from the definition of the above lengths that for any two non-exceptional vertices $g, h \in \Gamma_S^{\mathcal{H}}(G)$ the distance is given by:

$$d(g, h) = |g^{-1}h|_{S, \mathcal{H}}$$

Definition 2.3. Let G be a group with a finite generating system $S \subset G$, and let $\mathcal{H} = (H_1, \dots, H_r)$ be a finite family of finitely generated subgroups H_i of G .

- (1) The group G is *weakly hyperbolic relatively to \mathcal{H}* if the \mathcal{H} -coned Cayley graph $\Gamma_S^{\mathcal{H}}(G)$ is δ -hyperbolic, for some $\delta \geq 0$.
- (2) The group G is *strongly hyperbolic relatively to \mathcal{H}* if the graph $\Gamma_S^{\mathcal{H}}(G)$ is δ -hyperbolic and fine.

It is easy to see that these definitions are independent of the choice of the finite generating system S .

Definition 2.4. A finite family $\mathcal{H} = (H_1, \dots, H_r)$ of subgroups of a group G is called *malnormal* if:

- (a) for any $i \in \{1, \dots, r\}$ the subgroup H_i is malnormal in G (i.e. $g^{-1}H_i g \cap H_i = \{1\}$ for any $g \in G \setminus H_i$), and
- (b) for any $i, j \in \{1, \dots, r\}$ with $i \neq j$, and for any $g \in G$, one has $g^{-1}H_i g \cap H_j = \{1\}$.

This definition is stable with respect to permutation of the H_i , or replacing some H_i by a conjugate.

However, we would like to alert the reader that, contrary to many concepts used in geometric group theory, malnormality of a subgroup family $\mathcal{H} = (H_1, \dots, H_r)$ of a group G is not stable with respect to the usual modifications of \mathcal{H} that do not change the geometry of G relative to \mathcal{H} up to quasi-isometry. Such modifications are, for example, (i) the replacement of some H_i by a subgroup of finite index, or (ii) the addition of a new subgroup H_{r+1} to the family which is conjugate to a subgroup of some of the “old” H_i , etc. Malnormality, as can easily be seen, is sensible with respect to such changes: For example the infinite cyclic group \mathbb{Z} contains itself as malnormal subgroup, while the finite index subgroup $2\mathbb{Z} \subset \mathbb{Z}$ is not malnormal. Similarly, we verify directly that with respect to the standard generating system $S = \{1\}$ the coned off Cayley graph $\Gamma_S^{2\mathbb{Z}}$ is not fine. This underlines the well known but often not clearly expressed fact that the notion of strong relative hyperbolicity (i.e. “ δ -hyperbolic + fine”) is not invariant under quasi-isometry of the coned off Cayley graphs (compare also [13]), contrary to the otherwise less useful notion of weak relative hyperbolicity.

The following lemma holds for any hyperbolic group G , compare [8]. In the case used here, where $G = \mathbb{F}_n$ is a free group, the proof is indeed an exercise.

Lemma 2.5. *Let G be a hyperbolic group, and let $\mathcal{H} = (H_1, \dots, H_r)$ be a finite family of finitely generated subgroups.*

- (1) *If the family \mathcal{H} consists of quasi-convex subgroups, then G is weakly hyperbolic relative to \mathcal{H} .*
- (2) *If the family \mathcal{H} is quasi-convex and malnormal, then G is strongly hyperbolic relative to \mathcal{H} . \square*

For any $\alpha \in \text{Aut}(G)$, for any group G , a family of subgroups $\mathcal{H} = (H_1, \dots, H_r)$ is called α -invariant up to conjugation if there is a permutation σ of $\{1, \dots, r\}$ as well as elements $h_1, \dots, h_r \in G$ such that $\alpha(H_k) = h_k H_{\sigma(k)} h_k^{-1}$ for each $k \in \{1, \dots, r\}$.

The following notion has been proposed by Gromov [22] in the absolute case (i.e. all H_i are trivial) and generalized subsequently in [21].

Definition 2.6. Let G be a group generated by a finite subset S , and let \mathcal{H} be a finite family of subgroups of G . An automorphism α of G is *hyperbolic relative to \mathcal{H}* , if \mathcal{H} is α -invariant up to conjugation and if there exist constants $\lambda > 1, M \geq 0$ and $N \geq 1$ such that for any $w \in G$ with $|w|_{S, \mathcal{H}} \geq M$ one has:

$$\lambda |w|_{S, \mathcal{H}} \leq \max\{|\alpha^N(w)|_{S, \mathcal{H}}, |\alpha^{-N}(w)|_{S, \mathcal{H}}\}$$

The concept of a relatively hyperbolic automorphism is a fairly “stable” one, as shown by the following remark:

Remark 2.7. Let G, S, \mathcal{H} and α be as in Definition 2.6. The following statements can be derived directly from this definition.

- (a) The condition stated in Definition 2.6 is independent of the particular choice of the finite generating system S .
- (b) The automorphism α is hyperbolic relative to \mathcal{H} if and only if α^m is hyperbolic relative to \mathcal{H} , for any integer $m \geq 1$.
- (c) The automorphism α is hyperbolic relative to \mathcal{H} if and only if $\alpha' = \iota_v \circ \alpha$ is hyperbolic relative to \mathcal{H} , for any inner automorphisms $\iota_v : \mathbb{F}_n \rightarrow \mathbb{F}_n, w \mapsto v w v^{-1}$.

Every automorphism α of any group G defines a semi-direct product

$$G_\alpha = G \rtimes_\alpha \mathbb{Z} = G * \langle t \rangle / \langle\langle t g t^{-1} = \alpha(g) \text{ for all } g \in G \rangle\rangle$$

which is called the *mapping torus group* of α . In our case, where $G = \mathbb{F}_n$, one has

$$G_\alpha = \mathbb{F}_n \rtimes_\alpha \mathbb{Z} = \langle x_1, \dots, x_n, t \mid t x_i t^{-1} = \alpha(x_i) \text{ for all } i = 1, \dots, n \rangle$$

It is well known and easy to see that this group depends, up to isomorphisms which leave the subgroup $G \subset G_\alpha$ elementwise fixed, only on the outer automorphism defined by α .

Let $\mathcal{H} = (H_1, \dots, H_r)$ be a finite family of subgroups of G which is α -invariant up to conjugacy. For each H_i in \mathcal{H} let $m_i \geq 1$ be the smallest integer such that $\alpha^{m_i}(H_i)$ is conjugate in G to H_i , and let h_i be the conjugator: $\alpha^{m_i}(H_i) = h_i H_i h_i^{-1}$. We define the *induced mapping torus subgroup*:

$$H_i^\alpha = \langle H_i, h_i^{-1} t^{m_i} \rangle \subset G_\alpha$$

It is not hard to show that two subgroups H_i and H_j of G are, up to conjugation in G , in the same α -orbit if and only if the two induced mapping torus subgroups H_i^α and H_j^α are conjugate in the mapping torus subgroup G_α . (Note also that in a topological realization of G_α , for example as a fibered 3-manifold, the induced fibered submanifolds, over an invariant collection of disjoint subspaces with fundamental groups H_i , correspond precisely to the conjugacy classes of the H_i^α .)

Definition 2.8. Let $\mathcal{H} = (H_1, \dots, H_r)$ be a finite family of subgroups of G which is α -invariant up to conjugacy. A family of induced mapping torus subgroups

$$\mathcal{H}_\alpha = (H_1^\alpha, \dots, H_r^\alpha)$$

as above is the *mapping torus of \mathcal{H} with respect to α* if it contains for each conjugacy class in G_α of any H_i^α , for $i = 1, \dots, r$, precisely one representative.

The following Combination Theorem has been proved by the first author. For a reproof using somewhat different methods compare also [33].

Theorem 2.9 ([21]). *Let G be a finitely generated group, let $\alpha \in \text{Aut}(G)$ be an automorphism, and let $G_\alpha = G \rtimes_\alpha \mathbb{Z}$ be the mapping torus group of α . Let $\mathcal{H} = (H_1, \dots, H_r)$ be a finite family of finitely generated subgroups of G , and suppose that α is hyperbolic relative to \mathcal{H} .*

- (a) *If G is weakly hyperbolic relative to \mathcal{H} , then G_α is weakly hyperbolic relative to \mathcal{H} .*
- (b) *If G is strongly hyperbolic relative to \mathcal{H} , then G_α is strongly hyperbolic relative to the mapping torus \mathcal{H}_α of \mathcal{H} with respect to α .*

3 Polynomial growth subgroups

Let $\alpha \in \text{Aut}(\mathbb{F}_n)$ be an automorphism of \mathbb{F}_n . A subgroup H of \mathbb{F}_n is of *polynomial α -growth* if every element $w \in H$ is of *polynomial α -growth*: there are constants $C > 0$, $d \geq 0$ such that the inequality

$$\|\alpha^t(w)\| \leq Ct^d$$

holds for all integers $t \geq 1$, where $\|w\|$ denotes the cyclic length of w with respect to some basis of \mathbb{F}_n . Of course, passing over to another basis (or, for the matter, to any other finite generating system of \mathbb{F}_n) only affects the constant C in the above inequality.

We verify easily that, if $H \subset \mathbb{F}_n$ is a subgroup of polynomial α -growth, then it is also of polynomial β^k -growth, for any $k \in \mathbb{Z}$ and any $\beta \in \text{Aut}(\mathbb{F}_n)$ that represents the same outer automorphisms as α . Also, any conjugate subgroup $H' = gHg^{-1}$ is also of polynomial growth.

A family of polynomially growing subgroups $\mathcal{H} = (H_1, \dots, H_r)$ is called *exhaustive* if every element $g \in \mathbb{F}_n$ of polynomial growth is conjugate to an element contained in some of the H_i . The family \mathcal{H} is called *minimal* if no H_i is a subgroup of any conjugate of some H_j with $i \neq j$.

The following proposition is well known (compare [17]). For completeness we state it in full generality, although some ingredients (for example “very small” actions) are not specifically used here. The paper [30] may serve as an introductory text for the objects concerned.

Proposition 3.1. *Let $\alpha \in \text{Aut}(\mathbb{F}_n)$ be an arbitrary automorphism of \mathbb{F}_n . Then either the whole group \mathbb{F}_n is of polynomial α -growth, or else there is a very small action of \mathbb{F}_n on some \mathbb{R} -tree T by isometries, which has the following properties:*

- (a) *The \mathbb{F}_n -action on T is α -invariant with respect to a stretching factor $\lambda > 1$: one has*

$$\|\alpha(w)\|_T = \lambda\|w\|_T$$

for all $w \in \mathbb{F}_n$, where $\|w\|_T$ denotes the translation length of w on T , i.e. the value given by $\|w\|_T := \inf\{d(wx, x) \mid x \in T\}$.

- (b) *The stabilizer in \mathbb{F}_n of any non-degenerate arc in T is trivial:*

$$\text{Stab}([x, y]) = \{1\} \quad \text{for all } x \neq y \in T$$

- (c) *There are only finitely many orbits $\mathbb{F}_n \cdot x$ of points $x \in T$ with non-trivial stabilizer $\text{Stab}(x) \subset \mathbb{F}_n$. In particular, the family of such stabilizers $H_k = \text{Stab}(x_k)$, obtained by choosing an arbitrary point x_i in each of these finitely many \mathbb{F}_n -orbits, is α -invariant up to conjugation.*

- (d) *For every $x \in T$ the rank of the point stabilizer $\text{Stab}(x)$ is strictly smaller than n .*

We now define a finite iterative procedure, in order to identify all elements in \mathbb{F}_n which have polynomial α -growth: One applies Proposition 3.1 again to the non-trivial point stabilizers H_k as exhibited in part (c) of this proposition, where α is replaced by the restriction to H_k of a suitable power of α , composed with an inner automorphism of \mathbb{F}_n . By Property (d) of Proposition 3.1, after finitely many iterations this procedure must stop, and thus one obtains a partially ordered finite collection of such invariant \mathbb{R} -trees T_j .

In every tree T_j which is minimal in this collection, we choose a point in each of the finitely many orbits with non-trivial stabilizer, to obtain a finite family \mathcal{H} of finitely generated subgroups H_i of \mathbb{F}_n . It follows directly from this definition that every H_i has polynomial α -growth, and that the family \mathcal{H} is α -invariant up to conjugation.

The family \mathcal{H} is exhaustive, as, in each of the T_j , any path of non-zero length grows exponentially, by property (a) of Proposition 3.1. From property (b) we derive the minimality of \mathcal{H} : Indeed, we obtain the stronger property, that any two conjugates of distinct H_i can intersect only in the trivial subgroup $\{1\}$ (see Proposition 3.3).

It follows that the family \mathcal{H} is uniquely determined (by exhaustiveness and minimality), contrary to the above collection of invariant trees T_j , which is non-unique, as the tree T in Proposition 3.1 is in general not uniquely determined by α . The different choices, however, are well understood: a brief survey of the underlying structural analysis of α is given in §10.5 of the Appendix.

We summarize:

Proposition 3.2. (a) *Every automorphism $\alpha \in \text{Aut}(\mathbb{F}_n)$ possesses a finite family $\mathcal{H}(\alpha) = (H_1, \dots, H_r)$ of finitely generated subgroups H_i that are of polynomial growth, and $\mathcal{H}(\alpha)$ is exhaustive and minimal.*

(b) *The family $\mathcal{H}(\alpha)$ is uniquely determined, up to permuting the H_i or replacing any H_i by a conjugate.*

(c) *The family $\mathcal{H}(\alpha)$ is α -invariant.* □

The family $\mathcal{H}(\alpha) = (H_1, \dots, H_r)$ exhibited by Proposition 3.2 is called the *characteristic family of polynomial growth* for α . This terminology is slightly exaggerated, as the H_i are really only well determined up to conjugacy in \mathbb{F}_n . But on the other hand, the whole concept of a group G relative to a finite family of subgroups H_i is in reality a concept of G relative to a conjugacy class of subgroups H_i , and it is only for notational simplicity that one prefers to name the subgroups H_i rather than their conjugacy classes.

Proposition 3.3. *For every automorphism $\alpha \in \text{Aut}(\mathbb{F}_n)$ the characteristic family of polynomially growing subgroups $\mathcal{H}(\alpha)$ is quasi-convex and malnormal.*

Proof. The quasi-convexity is a direct consequence of the fact that the subgroups in $\mathcal{F}(\alpha)$ are finitely generated: Indeed, every finitely generated subgroup of a free group is quasi-convex, as is well known and easy to prove.

To prove malnormality of the family $\mathcal{H}(\alpha)$ we first observe directly from Definition 2.4 that if $\mathcal{H}' = (H'_1, \dots, H'_s)$ is a malnormal family of subgroups of some group G , and for each

$j \in \{1, \dots, s\}$ one has within H'_j a family of subgroups $\mathcal{H}''_j = (H''_{j,1}, \dots, H''_{j,r(j)})$ which is malnormal with respect to H'_j , then the total family

$$\mathcal{H} = (H''_{j,k})_{(j,k) \in \{1, \dots, s\} \times \{1, \dots, r(j)\}}$$

is a family of subgroups that is malnormal in G .

A second observation, also elementary, shows that given any \mathbb{R} -tree T with isometric G -action that has trivial arc stabilizers, every finite system of points $x_1, \dots, x_r \in T$ which lie in pairwise distinct G -orbits gives rise to a family of subgroups $(Stab(x_1), \dots, Stab(x_r))$ which is malnormal in G .

These two observations, together with Proposition 3.1, give directly the claimed malnormality of the characteristic family of polynomial α -growth. \square

4 β -train tracks

A new kind of train track maps $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$, called *partial train track maps with Nielsen faces*, has been introduced. by the second author (see Appendix). Here \mathcal{G}^2 consists of

- (a) a disjoint union X (called *the relative part*) of finitely many *vertex spaces* X_v ,
- (b) a finite collection $\widehat{\Gamma}$ (called *the train track part*) of edges e_j with endpoints in the X_v , and
- (c) a finite collection of 2-cells Δ_k with boundary in $\mathcal{G}^1 := X \cup \widehat{\Gamma}$.

The map f maps X to X and \mathcal{G}^1 to \mathcal{G}^1 . A path γ_0 in \mathcal{G}^1 is called a *relative backtracking path* if γ_0 is in \mathcal{G}^2 homotopic rel. endpoints to a path entirely contained in X . A path γ in \mathcal{G}^1 is said to be *relatively reduced* if any relative backtracking subpath of γ is contained in X .

Convention 4.1. (1) Note that throughout this paper we will only consider paths γ that are immersed except possibly at the vertices of \mathcal{G}^2 . (Recall that by hypothesis (b) above all vertices of \mathcal{G}^2 belong to X .) In other words, γ is either a classical edge path, or else an edge path with first and/or last edge that is only partially traversed. In the latter case, however, we require that this partially traversed edge belongs to $\widehat{\Gamma}$.

(2) Furthermore, for subpaths χ of γ that are entirely contained in X , we are only interested in the homotopy class in X relative endpoints. In the context considered in this paper, X will always be a graph, so that we can (and will tacitly) assume throughout the remainder of the paper that such χ is a reduced path in the graph X .

(3) We denote by $\bar{\gamma}$ the path γ with inverted orientation.

In particular, it follows from convention (2) that every relatively reduced path γ as above is reduced in the classical sense, when viewed as path in the graph \mathcal{G}^1 . The converse is wrong,

because of the 2-cells Δ_k in \mathcal{G}^2 . (Compare also part (b) of Definition-Remark 5.3, and the subsequent discussion.)

A path γ in \mathcal{G}^1 is called *legal* if for all $t \geq 1$ the path $f^t(\gamma)$ is relatively reduced. The space \mathcal{G}^2 and the map f satisfy furthermore the following properties:

- The map f has the *partial train track property relative to X* : every edge e of the train track part $\widehat{\Gamma}$ is legal.
- Every edge e from the train track part $\widehat{\Gamma}$ is *expanding*: there is a positive iterate of f which maps e to an edge path that runs over at least two edges from the train track part.
- For every path (or loop) γ in \mathcal{G}^1 there is an integer $t = t(\gamma) \geq 0$ such that $f^t(\gamma)$ is homopopic rel. endpoints (or freely homotopic) in \mathcal{G}^2 to a legal path (or loop) in \mathcal{G}^1 .

We say that $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ represents an automorphism α of F_n if there is a *marking* isomorphism $\theta : \pi_1 \mathcal{G}^2 \rightarrow F_n$ which conjugates the induced morphism $f_* : \pi_1 \mathcal{G}^2 \rightarrow \pi_1 \mathcal{G}^2$ to the outer automorphism $\widehat{\alpha}$ given by α .

Building on deep work of Bestvina-Handel [3], the second author has shown [28] that every automorphism α of F_n has a partial train track representative with Nielsen faces $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$, and all conjugacy classes represented by loops in the relative part have polynomial α -growth. However, for the purpose of this paper an additional property is needed, which in [28], [29] only occurs for the “top stratum” of $\widehat{\Gamma}$, namely that legal paths lift to quasi-geodesics in the universal covering $\widetilde{\mathcal{G}}^2$.

This is the reason why one needs to work here with *β -train track maps*, and with *strongly legal* paths, which have this additional property. This improvement, and some other technical properties of β -train tracks needed later are presented in detail in the next section.

The following result is presented in the Appendix. Note that all properties of β -train tracks maps which are used below are explicitly listed here.

Theorem 4.2. *Every automorphism α of F_n is represented by a β -train track map. This is a partial train track map with Nielsen faces $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ relative to a subspace $X \subset \mathcal{G}^2$, which satisfies:*

- (a) *Every connected component X_v of X is a graph, and the marking isomorphism $\theta : \pi_1 \mathcal{G}^2 \rightarrow F_n$ induces a monomorphism $\pi_1 X_v \rightarrow F_n$. Every conjugacy class represented by a loop in X has polynomial growth.*
- (b) *There is a subgraph $\Gamma \subset \mathcal{G}^1$, which contains all of the train track part $\widehat{\Gamma}$, and there is a homotopy equivalence $r : \mathcal{G}^2 \rightarrow \Gamma$ which restricts to the identity on Γ , such that the composition-restriction $f_\Gamma = r \circ f|_\Gamma : \Gamma \rightarrow \Gamma$ is a classical relative train track map as defined in [3].*
- (c) *Every edge e of the train track part of \mathcal{G}^2 is strongly legal (see Definition 5.4) and thus in particular legal.*
- (d) *Every strongly legal path in \mathcal{G}^1 is mapped by f to a strongly legal path.*

- (e) Every edge e from the train track part $\widehat{\Gamma}$ is expanding: there is a positive iterate of f which maps e to an edge path that runs over at least two edges from $\widehat{\Gamma}$.
- (f) The lift of any strongly legal path γ to the universal covering $\widetilde{\mathcal{G}}^2$ is a quasi-geodesic with respect to the simplicial metric on $\widetilde{\mathcal{G}}^2$ (where every edge in either, the train track and the relative part, is given length 1), for some fixed quasi-geodesy constants independent of the choice of γ .
- (g) Every reduced path in Γ lifts also to a quasi-geodesic in $\widetilde{\mathcal{G}}^2$. Every path that is mapped by the retraction r to a reduced path in Γ lifts also to a quasi-geodesic in $\widetilde{\mathcal{G}}^2$. In particular, every path which derives from a strongly legal path by applying r to any collection of subpaths does lift to a quasi-geodesic in $\widetilde{\mathcal{G}}^2$.
- (h) For every path γ in \mathcal{G}^1 there is an integer $\widehat{t} = \widehat{t}(\gamma) \geq 0$ such that $f^{\widehat{t}}(\gamma)$ is homotopic rel. endpoints in \mathcal{G}^2 to a strongly legal path in \mathcal{G}^1 . The integer $\widehat{t}(\gamma)$ depends only on the number of illegal turns (compare Definition 5.4) in γ and not on γ itself.

For further use of β -train track maps, in particular with respect to a structural analysis of automorphisms of \mathbb{F}_n , we refer the reader to §10.5 of the Appendix.

5 Strongly legal paths and INP's in β -train tracks

Let $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ be a β -train track map as described in the previous section. Recall from the beginning of the last section that a path γ in \mathcal{G}^1 is *legal* if, for any $t \geq 1$, the image path $f^t(\gamma)$ is relatively reduced, i.e. every relative backtracking subpath of $f^t(\gamma)$ is completely contained in the relative part $X \subset \mathcal{G}^1$. For the precise definition of a “path” recall Convention 4.1.

Definition 5.1. An INP is a reduced path $\eta = \eta' \circ \eta''$ in \mathcal{G}^1 which has the following properties:

- (0) The first and the last edge (or non-trivial edge segment) of the path η belongs to the train track part $\widehat{\Gamma} \subset \mathcal{G}^1$.
- (1) The subpaths η' and η'' (called the *branches* of η) are legal.
- (2) The path $f^t(\eta)$ is not legal, for any $t \geq 0$.
- (3) For some integer $t_0 \geq 1$ the path $f^{t_0}(\eta)$ is homotopic relative to its endpoints, in \mathcal{G}^1 , to the path η .

We would like to alert the reader that in the literature one requires sometimes in property (3) above that $t_0 = 1$, and that for $t_0 \geq 2$ one speaks of a *periodic INP*. We will not make this notational distinction in this paper.

For every INP η there is an associated *auxiliary edge* e in the relative part $X \subset \mathcal{G}^1$ which has the same endpoints as η . The relative part $X \subset \mathcal{G}^1$ consists precisely of all auxiliary

edges and of all edges e' of $\Gamma \setminus \widehat{\Gamma}$. In other words: \mathcal{G}^1 is the union of Γ with the set of all auxiliary edges.

The canonical retraction $r : \mathcal{G}^2 \rightarrow \Gamma$ from Theorem 4.2 (b) is given on \mathcal{G}^1 as power \widehat{r}^n of the map $\widehat{r} : \mathcal{G}^1 \rightarrow \mathcal{G}^1$ which is the identity on Γ and maps every auxiliary edge e to the associated INP-path $\widehat{r}(e) = \eta$. Recall in this context that there are only finitely many INP's and thus only finitely many auxiliary edges, for any given β -train track map.

Aside 5.2. Technically speaking, an auxiliary edge e is in truth the union of two *auxiliary half-edges*, which meet at an *auxiliary vertex* which is placed in the center of e and belongs to the relative part. The reason for this particularity lies in the fact that otherwise 3 (or more) auxiliary edges could form a non-trivial loop γ in X which is contractible in \mathcal{G}^2 .

To avoid this phenomenon (compare the “expansion of a Nielsen face” in Definition 3.7 of [28]), in this case there is only one auxiliary vertex which is the common center of the three auxiliary edges, and only three auxiliary half-edges, arranged in the shape of a tripod with the auxiliary vertex as center: the union of any two of the auxiliary half edges defines one of the three auxiliary edge we started out with. As a consequence, the above loop γ is in fact a contractible loop in the tripod just described. For more detail and the relation with attractive fixed points at $\partial\mathbb{F}_n$ see [28], Definition 3.7.

Definition-Remark 5.3. (a) A *turn* is a path in \mathcal{G}^1 of the type $e \circ \chi \circ e'$, where e and e' are edges (or non-trivial edge segments) from the train track part $\widehat{\Gamma}$ of \mathcal{G}^1 , while χ is an edge path (possibly trivial !) entirely contained in the relative part $X \subset \mathcal{G}^1$. We recall (Convention 4.1) that one is only interested in χ up to homotopy rel. endpoints, within the subspace X , and thus one always assumes that χ has been isotoped to be a reduced path in the graph X .

(b) A path γ is not legal if and only if for some $t \geq 1$ the path $f^t(\gamma)$ contains a turn $e \circ \chi \circ e'$ as above which (i) either is not relatively reduced, i.e. χ is a contractible loop and $\bar{e} = e'$, or else (ii) the path χ is (after reduction) an auxiliary edge e_0 with associated INP $\widehat{r}(e_0) = \eta$, such that η starts in \bar{e} and ends in e' .

(c) A path γ in \mathcal{G}^1 is legal if and only if all of its turns are legal. In particular, every legal path is relatively reduced (and thus reduced in the graph \mathcal{G}^1 , see Convention 4.1). The converse implication is false.

(d) Every INP $\eta = \eta' \circ \eta''$ as in Definition 5.1 has precisely one turn that is not legal, called the *tip* of η . This is the turn from the last train track edge of η' to the first train track edge of η'' . More specifically, for all $t \geq 1$ the path $f^t(\eta)$ contains precisely one turn (= the turn from the last train track edge of $f^t(\eta')$ to the first train track edge of $f^t(\eta'')$) that is not relatively reduced, as above in alternative (i) of part (b).

Although not needed in the sequel, we would like to explain the case (ii) of part (b) above:

For some sufficiently large exponent $t' \geq 1$ there will be a terminal segment e_1 of e and an initial segment e'_1 of e' with $f^{t'}(e_1 \circ e_0 \circ e'_1) = \bar{\eta}' \circ e_0 \circ \bar{\eta}''$. Thus the subpath $f^{t'}(e_1 \circ e_0 \circ e'_1)$ of $f^{t+t'}(\gamma)$, while not contained in X , is relatively backtracking, since $\bar{\eta}' \circ e_0 \circ \bar{\eta}''$ is contractible

in \mathcal{G}^2 . This is because \mathcal{G}^2 contains for every auxiliary edge e_0 a 2-cell Δ_{e_0} (called a *Nielsen face*) with boundary path $\bar{e}_0 \circ \eta$. By definition it follows that γ is not legal.

Definition-Remark 5.4. (a) A *half turn* is a path in \mathcal{G}^1 of the type $e \circ \chi$ or $\chi \circ e'$, where e and e' are edges (or non-trivial edge segments) from the train track part $\widehat{\Gamma}$ of \mathcal{G}^1 , while χ is a non-trivial reduced edge path entirely contained in the relative part $X \subset \mathcal{G}^1$.

Every finite path γ contains only finitely many maximal (as subpaths of γ) half turns, namely precisely two at each turn, plus a further half turn at the beginning and another one at the end of γ .

(b) A path γ in \mathcal{G}^1 is called *strongly legal* if it is legal (and thus reduced in \mathcal{G}^1), and if in addition it has the following property: The path γ' , obtained from γ through replacing every auxiliary edge e_i on γ by the associated INP $\eta_i = \widehat{r}(e_i)$, contains as only illegal turns the tips of the INPs η_i .

(c) A legal (and hence reduced) path γ in \mathcal{G}^1 is strongly legal if and only if all maximal half turns in γ are strongly legal. A half turn $e \circ \chi$ (or similarly $\chi \circ e'$) in γ is not strongly legal if and only if the first edge of χ is an auxiliary edge e' with $\widehat{r}(e') = \eta$, and for some $t \geq 1$ the first edge of $f^t(\eta)$ is precisely the first edge of the legal path $f^t(\bar{e})$.

(d) A turn is called *illegal* if it is not legal, or if any of its two maximal sub-half-turns is not strongly legal.

We will now treat explicitly a technical subtlety which is relevant for the next section:

If η is an INP in \mathcal{G}^1 , decomposed as above into two legal (actually they turn out to be always strongly legal !) branches $\eta = \eta' \circ \eta''$, then it can happen that η' (or η'') contains an auxiliary edge e_1 . Replacing now e_1 by its associated INP $\widehat{r}(e_1) = \eta_1$, the same phenomenon can occur again: the legal branches of η_1 may well run over an auxiliary edge. However, this process can repeat only a finite number of times.

This is the reason why above we distinguish between an INP $\eta = \widehat{r}(e)$ with associated auxiliary edge e on one hand, and the path $r(e)$ in $\Gamma \subset \mathcal{G}^1$ obtained through finitely iteration of \widehat{r} on the other hand. For any auxiliary edge e we call the reduced path $r(e)$ in Γ a *pre- INP* , and we observe that such a pre- INP may well contain another such pre- INP as subpath (although not as boundary subpath, by property (0) of Definition 5.1).

Definition 5.5. A pre- INP $r(e)$ in a reduced path γ in Γ is called *isolated*, if any other pre- INP in γ that intersects $r(e)$ in more than a point is contained as subpath in $r(e)$.

Clearly, replacing each such isolated pre- INP $r(e_i)$ of γ by the associated auxiliary edge e_i yields a path γ' in \mathcal{G}^1 which does not depend on the order in which these replacements are performed, and is thus uniquely determined by γ . It also satisfies $r(\gamma') = \gamma$, which is a reduced path, by hypothesis. Such a path γ' is called a *normalized* path in \mathcal{G}^1 ; they will be investigated more thoroughly in the next section.

6 Normalized paths in β -train tracks

Throughout this section we assume that a β -train track map $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ is given as defined in the previous two sections, and that f represents an automorphism α of \mathbb{F}_n . We will use in this section both, the *absolute* and the *relative* length of a path γ in \mathcal{G}^1 : The absolute length $|\gamma|_{abs}$ is given by associating to every edge e of \mathcal{G}^1 , i.e. of $\widehat{\Gamma}$ and of X , the length $L(e) = 1$. The relative length $|\gamma|_{rel}$ is given by associating to every edge e in the train track part $\widehat{\Gamma} \subset \mathcal{G}^1$ the length $L(e) = 1$, while every edge e' in the relative part $X \subset \mathcal{G}^1$ is given length $L(e') = 0$.

We will now start with our study of normalized paths. The reader should keep in mind that lifts of normalized paths to the universal cover $\widetilde{\mathcal{G}}^2$ of \mathcal{G}^2 are meant (and shown below) to be strong analogues of geodesic segments in a tree. For example, one can see directly from the definition that a normalized path is reduced in \mathcal{G}^1 and relatively reduced in \mathcal{G}^2 , and that a concatenation of normalized paths, even if not normalized, is necessarily relatively reduced in \mathcal{G}^2 if it is reduced in \mathcal{G}^1 .

Definition 6.1. A path γ in \mathcal{G}^1 is *normalized*, if

- (i) the path $r(\gamma)$ in Γ is reduced, and
- (ii) the path γ is obtained from $r(\gamma)$ through replacing every isolated pre-INP $r(e)$ of $r(\gamma)$ by the associated auxiliary edge e .

Proposition 6.2. *For every path γ in \mathcal{G}^1 there is a unique normalized path γ_* in \mathcal{G}^1 which is (in \mathcal{G}^2) homotopic to γ relative to its endpoints.*

Proof. To prove existence, it suffices to apply the retraction r to γ , followed by a subsequent reduction, to get a reduced path in $\Gamma \subset \mathcal{G}^1$ that is homotopic rel. endpoints in \mathcal{G}^2 to γ . One then replaces iteratively every isolated pre-INP by the associated auxiliary edge to get γ_* .

Since reduced paths in Γ are uniquely determined with respect to homotopy rel. endpoints, to prove uniqueness of γ_* we only have to verify that for every normalized path the above explained procedure reproduces the original path. This follows directly from the definition of an “isolated” pre-INP at the end of §5. \square

For an arbitrary path γ in the graph \mathcal{G}^1 we always denote by γ_* the normalized path obtained from γ as given in Proposition 6.2.

Proposition 6.3. *Let $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ be a β -train track map.*

- (a) *Every strongly legal path γ in \mathcal{G}^1 is normalised.*
- (b) *If γ is a path in \mathcal{G}^1 that is entirely contained in the relative part $X \subset \mathcal{G}^2$, then the normalized path γ_* is also entirely contained in X .*
- (c) *Normalized paths lift in the universal covering $\widetilde{\mathcal{G}}^2$ to quasi-geodesics, with respect to the absolute metric on $\widetilde{\mathcal{G}}^2$.*

Proof. Statement (a) follows directly from the above definition of a normalized path. For statement (b) the same is true, but we also need the subtlety involved when introducing the auxiliary edges that has been indicated in Aside 5.2. Part (c) follows directly from Theorem 4.2 (g). \square

Lemma 6.4. *There exists a “composition constant” $E > 0$ which has the following property:*

(1) *Let γ_1 and γ_2 be two normalized paths in \mathcal{G}^1 , and let $\gamma = \gamma_1 \circ \gamma_2$ be the (possibly non-reduced or non-normalized) concatenation. Then there are decompositions $\gamma_1 = \gamma'_1 \circ \gamma''_1$ and $\gamma_2 = \gamma''_2 \circ \gamma'_2$ such that the normalized path γ_* can be written as concatenation*

$$\gamma_* = \gamma'_1 \circ \gamma_{1,2} \circ \gamma'_2,$$

where the path $\gamma_{1,2}$ has absolute length

$$|\gamma_{1,2}|_{abs} \leq E.$$

(2) *If one assumes that the concatenation $\gamma = \gamma_1 \circ \gamma_2$ is reduced, then one can furthermore conclude that also the paths γ''_1 and γ''_2 have absolute length $\leq E$.*

Proof. (1) We first observe that by definition of normalized paths the two paths $r(\gamma_1)$ and $r(\gamma_2)$ in Γ are reduced. Hence there is an initial subpath r_1 of $r(\gamma_1)$ as well as a terminal subpath r_2 of $r(\gamma_2)$ such that the (possibly non-reduced) concatenation $r(\gamma_1) \circ r(\gamma_2)$ of the reduced paths $r(\gamma_i)$ can be simplified to give the reduced path $r_1 \circ r_2$. The claim now is a direct consequence of the following observation: Any pre-INP $r(e)$ in the subpaths r_i is isolated in r_i if and only if it is isolated in the concatenation $r_1 \circ r_2$, unless $r(e)$ is contained in a neighborhood of the concatenation point. But the size of this neighborhood only depends on the maximal absolute length of any pre-INP in \mathcal{G}^1 and is hence independent of the particular paths considered.

(2) In order to prove the stronger claim (2) it suffices to show that, if the concatenation $\gamma_1 \circ \gamma_2$ is reduced, then the possible cancellation in $r(\gamma_1) \circ r(\gamma_2)$ is bounded.

By way of contradiction, assume that the reduced paths $r(\gamma_1)$ (= the path $r(\gamma_1)$ with orientation reversed) and $r(\gamma_2)$ have a long common initial segment γ_0 . By the argument given above in part (a), the occurrences of isolated pre-INP's in γ_0 , other than in a terminal subsegment of γ_0 of a priori bounded length, do not depend on whether we consider the segment γ_0 as part of $r(\gamma_1)$ or of $r(\gamma_2)$. But then the normalized paths $\bar{\gamma}_1$ and γ_2 will also have a long common initial segment, which contradicts the assumption that $\gamma_1 \circ \gamma_2$ is reduced. \square

The following is crucially used in the next section:

Corollary 6.5. *For any constant $D > 0$ there exists a bound $K > 0$ which has the following property: Let $\gamma = \gamma_1 \circ \gamma_2 \circ \gamma_3$ be a concatenated path in \mathcal{G}^1 , and assume that their lengths satisfy:*

- (i) $|\gamma_1|_{abs} \leq D$
- (ii) $|\gamma_2|_{rel} = 0$

(iii) $|\gamma_3|_{abs} \leq D$

Then the normalized path γ_* has relative length

$$|\gamma_*|_{rel} \leq K.$$

Proof. We consider the normalized paths γ_{i*} and observe that, by Proposition 6.3 (c), the absolute length of γ_{1*} and γ_{3*} is bounded above by a constant only dependent on D . Furthermore, the relative length is always smaller or equal to the absolute one. Hence the sum of the relative lengths of the γ_{i*} depends only on D , and Lemma 6.4 (1) implies directly that the same is true for the relative length of the normalized path γ_* . \square

Lemma 6.6. *There exists a constant $J \geq 1$ such that for any path γ in \mathcal{G}^1 the following holds, where $ILLT(\cdot)$ denotes the number of illegal turns in a path:*

(a) *If γ is non-reduced, then the path γ' obtained from γ through reduction in \mathcal{G}^1 satisfies*

$$ILLT(\gamma') \leq ILLT(\gamma)$$

(b) *If γ is reduced, and γ_* is obtained from γ through normalization, then one has:*

$$ILLT(\gamma_*) \leq J \cdot ILLT(\gamma)$$

Proof. (a) This is a direct consequence of the fact that at every turn $e \circ \chi \circ e'$ of γ , where χ is a reduced path in X , either γ is reduced (in the graph \mathcal{G}^1), or χ is trivial and $e' = \bar{e}$, in which case the turn is illegal.

(b) We first use Proposition 6.3 (b) to observe that the maximal strongly legal subpaths of γ are normalized. We then use iteratively Lemma 6.4 (1) to obtain $k = ILLT(\gamma)$ subpaths γ_i of γ_* , each of absolute length bounded above by the constant E from Lemma 6.4 (1), such that every complementary subpath of the union of the γ_i in γ_* is strongly legal. But the number of illegal turns in any γ_i cannot exceed the absolute length of γ_i , which gives directly the claim. \square

Proposition 6.7. *Let $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ be a β -train track map. Then there is a integer $K \geq 1$ such that for any normalized path γ in \mathcal{G}^1 , the number $ILLT(\cdot)$ of illegal turns satisfies:*

$$ILLT(\gamma) \geq 2 ILLT(f^K(\gamma)_*)$$

Proof. We first consider any path γ'' in \mathcal{G}^1 with at most $2J + 1$ illegal turns, for $J \geq 1$ as given in Lemma 6.6. By property (h) of Theorem 4.2 there is a constant K such that $f^K(\gamma'')_*$ is strongly legal, for all such paths γ'' .

We now subdivide γ into $k + 1 \leq \frac{ILLT(\gamma)}{2J}$ subpaths such that each subpath has $\leq 2J + 1$ illegal turns. We consider the normalized f^K -image of each subpath, which is strongly legal, and their concatenation $f^K(\gamma)$, which satisfies $ILLT(f^K(\gamma)) \leq k$, but is a priori not reduced, and after reduction a priori not normalized. We then apply Lemma 6.6 to obtain $ILLT(f^K(\gamma)_*) \leq J \cdot k$ and hence $ILLT(f^K(\gamma)_*) \leq \frac{1}{2} ILLT(\gamma)$. \square

Lemma 6.8. *There is a “cancellation bound” $C = C(f) > 0$ such that for any concatenated normalized path $\gamma = \gamma_1 \circ \gamma_2$ the normalized image path decomposes as $f(\gamma)_* = \gamma'_1 \circ \gamma_{1,2} \circ \gamma'_2$, with $f(\gamma_1)_* = \gamma'_1 \circ \gamma''_1$ and $f(\gamma_2)_* = \gamma''_2 \circ \gamma'_2$, and all three, $\gamma_{1,2}, \gamma''_1$ and γ''_2 have length $\leq C$.*

Proof. The analogous statement, with every normalized path replaced by its (reduced !) image in Γ under the retraction r , follows directly from the fact that f represents an automorphism and hence induces a quasi-isometry on the universal covering of \mathcal{G}^2 , with respect to the absolute metric.

To deduce now the desired statement for normalized paths it suffices to apply the same arguments as in the proof of Lemma 6.4 (2). \square

Let γ be a path in \mathcal{G}^1 , and let $C > 0$ be any constant. We say that a strongly legal subpath γ' of γ has *strongly legal C -neighborhood* in γ if γ' occurs as subpath of a larger strongly legal subpath γ'' of γ which is of the form $\gamma'' = \gamma_1 \circ \gamma' \circ \gamma_2$, where each of the γ_i either has relative length $|\gamma_i|_{rel} = C$, or else γ_i is a boundary subpath (possibly of length 0) of γ .

In other words, there is no illegal turn in γ that has relative distance $< C$ within γ from the subpath γ' .

Let $b(f) \geq 1$ denote the *expansion exponent* of f , defined to be the smallest positive exponent such that for any edge $e \in \widehat{\Gamma}$ the image $f^{b(f)}(e)$ is an edge path of relative length ≥ 2 . The existence of $b(f)$ is a direct consequence of statement (e) of Theorem 4.2.

Proposition 6.9. *Let $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ be a β -train track map, let $b = b(f)$ be the expansion exponent of f , and let $C = C(f^b)$ be the cancellation bound for f^b as given in Lemma 6.8. Then for any normalized edge path γ in \mathcal{G}^1 the following holds:*

Every strongly legal subpath γ_0 with strongly legal C -neighborhood in γ is mapped by f^b to a strongly legal path γ'_0 which is contained as subpath with strongly legal C -neighborhood in $f^b(\gamma)_$. Furthermore, their relative lengths satisfy:*

$$|\gamma'_0|_{rel} \geq 2|\gamma_0|_{rel}$$

Proof. By definition of the exponent b every strongly legal path γ_0 is mapped to a path $f^b(\gamma_0)$ of relative length $|f^b(\gamma_0)|_{rel} \geq 2|\gamma_0|_{rel}$.

Now, every strongly legal path is normalized (by Proposition 6.3 (a)), and the image of a strongly legal path is again strongly legal (by Theorem 4.2 (d)). Since γ_0 has strongly legal C -neighborhood, and b is the expansion constant of f , the path $f^b(\gamma_0) = f^b(\gamma_0)_*$ has strongly legal $2C$ -neighborhood in the (possibly unreduced and after reduction not normalized) path $f^b(\gamma)$. But then Lemma 6.8 proves directly that in the normalized path $f^b(\gamma)_*$ the path $f^b(\gamma_0)$ has still strongly legal C -neighborhood. \square

Corollary 6.10. *For every $\lambda > 1$ there exist an integer $N \geq 1$ such that, if γ is a normalized path in \mathcal{G}^1 then*

(a) either the normalized path $f^N(\gamma)_$ has relative length*

$$|f^N(\gamma)_*|_{rel} \geq \lambda|\gamma|_{rel}$$

(b) or else any normalized path γ' in \mathcal{G}^1 with $f^N(\gamma')_* = \gamma$ satisfies

$$|\gamma'|_{rel} \geq \lambda |\gamma|_{rel}.$$

Proof. Let $k \geq 0$ be the number of illegal turns in γ and set $C = C(f^b(f))$ as in Proposition 6.9. There are finitely many (at most $k + 1$ ones) maximal strongly legal subpaths γ_i with strongly legal C -neighborhood in γ . If $|\gamma|_{rel} \geq 3Ck$, then the total relative length of the γ_i exceeds $\frac{1}{3}|\gamma|_{rel}$. Applying Proposition 6.9 iteratively to each of them gives directly the claim (a).

If $|\gamma|_{rel} < 3Ck$ we apply iteratively Proposition 6.7. Since the relative length of any of the strongly legal subpath of γ' between two adjacent illegal turns is bounded below by 1 (= the length of any edge in the train track part), we derive directly the existence of a constant $N \geq 1$ that has the property claimed in statement (b). \square

Remark 6.11. Corollary 6.5 above is used crucially in the next section. The proof given in this section relies on the particular properties of normalized paths as introduced in this paper. In this remark we would like to propose an alternative proof, which dates back to the original plan for this paper. It is conceptually simpler in that it doesn't directly appeal to train track technology and to the very intimate knowledge of normalized paths which we have used earlier in this section. It uses though some of the later material of this section, such as Corollary 6.10. However, the latter is anyway used crucially in §8 below.

We do not present this alternative proof in full detail, but we believe that the interested reader can recover the latter from the sketch given here. We first collect the following observations; only the last one is non-trivial, and none of them uses Corollary 6.5:

1. The relative length of a path is smaller or equal to the absolute length of the same path.
2. The map f induces a quasi-isometry for both, the relative and absolute metric.
3. The relative part is quasi-convex in the given 2-complex $\tilde{\mathcal{G}}^2$.
4. The absolute length of any geodesic in the tree $\tilde{\Gamma}$ which connect two points in a same connected component of the relative part grows at most polynomially, under iteration into the future and also into the past.
5. Any subpath of a normalized path is normalized, up to adding subpaths with absolute length uniformly bounded above at its extremities.
6. The relative length of each normalized path is expanded by a factor $\lambda > 1$ after N iterations, either into the future or into the past. (This is the content of Corollary 6.10.)

We conclude:

Let c be a normalized path. If $c' \subset c$ has its endpoints at absolute distance less than C from two points in a same connected component of the relative part, then the relative length of c' is bounded above by some constant depending only on C . (This is essentially the content of Corollary 6.5.)

The proof of this conclusion is not really hard and only requires to manipulate properly some of the inequalities that are given by the above observations: One compares the exponential expansion of the relative length to the polynomial expansion of the absolute length, which eventually leads to a contradiction, unless the above conclusion holds.

7 Normalized paths are relative quasi-geodesics

In this section we consider the universal covering $\tilde{\mathcal{G}}^2$ of \mathcal{G}^2 with respect to both, the *absolute metric* d_{abs} and the *relative metric* d_{rel} , which are defined by lifting the absolute and the relative edge lengths respectively from \mathcal{G}^2 to $\tilde{\mathcal{G}}^2$. We also “lift” the terminology: for example, the *relative part* of $\tilde{\mathcal{G}}^2$ is the lift of the relative part $X \subset \mathcal{G}^2$.

We first note that every connected component of the relative part of $\tilde{\mathcal{G}}^2$ is quasi-convexly embedded, with respect to the absolute metric, since the fundamental group of any connected component of X is a finitely generated subgroup of the free group $\pi_1 \mathcal{G}^2 = \mathbb{F}_n$.

Next, we recall that $\tilde{\mathcal{G}}^2$, with respect to the absolute metric, is quasi-isometric to a metric tree: Such a quasi-isometry is given for example by any lift of the retraction $r : \mathcal{G}^2 \rightarrow \Gamma$ to $\tilde{r} : \tilde{\mathcal{G}}^2 \rightarrow \tilde{\Gamma}$, which is again a retraction, and $\tilde{\Gamma}$ is a metric simplicial tree with free \mathbb{F}_n -action.

Finally, let us recall that a path γ is a (λ, μ) -quasi-geodesic, for given constants $\lambda > 0$, $\mu \geq 0$, if and only if for every subpath γ' of γ , with endpoints x' and y' , one has:

$$|\gamma'| \leq \lambda d(x', y') + \mu$$

Proposition 7.1. *For all constants $\lambda, \mu > 0$ there are constants $\lambda', \mu', C > 0$, such that the following holds in $\tilde{\mathcal{G}}^2$:*

For every absolute (λ, μ) -quasi-geodesic γ there exists a relative (λ', μ') -quasi-geodesic $\hat{\gamma}$ which is of absolute Hausdorff distance $\leq C$ from γ .

Proof. We consider a relative geodesic γ' with same endpoints as γ , as well as their images $\tilde{r}(\gamma)$ and $\tilde{r}(\gamma')$. The path $\tilde{r}(\gamma)$ is contained in an absolute neighborhood of the geodesic segment $[x, y]$ in the tree $\tilde{\Gamma}$, where x and y are the endpoints of $\tilde{r}(\gamma)$.

Since $\tilde{\Gamma}$ is a tree, the path $\tilde{r}(\gamma')$ must run over all of $[x, y]$, so that we can consider a minimal collection of subpaths γ'_i of γ' such that the union of all $\tilde{r}(\gamma'_i)$ contains the segment $[x, y]$. (Here “minimal” means that no collection of proper subpaths of the γ'_i has the same property). We note that the number of such subpaths is bounded above by the absolute length of $[x, y]$.

We now enlarge these subpaths by a bounded amount, to ensure that they are edge paths: This ensures that the preimage γ'_i of any such $\tilde{r}(\gamma'_i)$ is either

- (i) completely contained in the relative part, or else

(ii) it is of relative length ≥ 1 .

Now, the adjacent endpoints of any two subsequent γ'_i can be connected by paths γ'_j of bounded absolute length in $\tilde{\mathcal{G}}^2$, and, if the two endpoints belong to the same connected component of the relative part, then by the absolute quasi-convexity of the latter we can assume that γ'_j as well belongs to this component. In particular, we observe that the number of paths γ'_j that are not contained in the relative part is bounded above by the relative length of γ' .

Hence the path $\hat{\gamma}$, defined as alternate concatenation of the γ'_i and γ'_j , has relative length given as sum of the relative length of the pairwise disjoint subpaths γ'_i of the relative geodesic γ' , plus the relative length of the γ'_j , which is uniformly bounded. Since the number of γ'_j is also bounded by the relative length of γ' , it follows that there are constants as in the proposition which bound the relative length of $\hat{\gamma}$.

Since the very same arguments extend to all subpaths of $\hat{\gamma}$, it follows directly that $\hat{\gamma}$ is a relative quasi-geodesic as claimed. \square

Below we need the following lemma; its proof follows directly from the definition of a quasi-geodesic and the inequality $d_{rel}(\cdot, \cdot) \leq d_{abs}(\cdot, \cdot)$.

Lemma 7.2. *For any constants $\lambda, \mu > 0$, every relative (λ, μ) -quasi-geodesic γ in $\tilde{\mathcal{G}}^2$, which does not traverse any edge from the relative part, is also an absolute (λ, μ) -quasi-geodesic. \square*

Proposition 7.3. *There exist constants $\lambda, \mu > 0$ such that in $\tilde{\mathcal{G}}^2$ any lift γ of a normalized path γ_0 in \mathcal{G}^1 is a relative (λ, μ) -quasi-geodesic.*

Proof. We note that it suffices to prove:

- (*) There exist constants $C_1, C_2, C_3 \geq 0$ as well as $\lambda' \geq 1, \mu' \geq 0$, such that for any subpath γ' of γ , with endpoints x', y' (of γ'), there exist a relative (λ', μ') -quasi-geodesic $\hat{\gamma}'$ with endpoints \hat{x}', \hat{y}' , such that $d_{rel}(x', \hat{x}') \leq C_1$ and $d_{rel}(y', \hat{y}') \leq C_1$, and

$$|\gamma'|_{rel} \leq C_2 |\hat{\gamma}'|_{rel} + C_3.$$

By Proposition 6.3 (c), the lift γ of the normalized path γ_0 is an absolute quasi-geodesic, for quasi-geodesy constants independent of the choice of γ . Now, Proposition 7.1 gives a relative quasi-geodesic $\hat{\gamma}$ in an absolute Hausdorff neighborhood of γ , where the seize of this neighborhood as well as the quasi-geodesy constants are again independent of γ . As a consequence, for any subpath γ' of γ we find a corresponding subpath $\hat{\gamma}'$ of $\hat{\gamma}$ which satisfies the endpoint conditions in (*) for some constant $C_1 > 0$ independent of our choices.

Without loss of generality we can assume that the path $\hat{\gamma}$ is contained in the 1-skeleton of $\tilde{\mathcal{G}}^2$, and that furthermore $\hat{\gamma}'$ is an edge path, i.e. starts and ends at a vertex of $\tilde{\mathcal{G}}^2$.

We now consider the set $\hat{\mathcal{L}}$ of maximal subpaths $\hat{\gamma}_i$ of $\hat{\gamma}'$ which are contained in the relative part. The collection of closed subpaths $\hat{\gamma}_j$ of $\hat{\gamma}'$ complementary to those in $\hat{\mathcal{L}}$ is denoted by $\hat{\mathcal{L}}^c$. We observe that, by Lemma 7.2, every such $\hat{\gamma}_j$ is an absolute quasi-geodesic, with quasi-geodesy constants depending only on C_1 and not on our choice of $\hat{\gamma}'$. Furthermore,

every such $\widehat{\gamma}_j$ has absolute length ≥ 1 (= the relative length of any edge outside the relative part), and we have:

$$|\widehat{\gamma}_j|_{abs} = |\widehat{\gamma}_j|_{rel}$$

The path γ' inherits a natural “decomposition” $\mathcal{L} \sqcup \mathcal{L}^c$ from the decomposition of $\widehat{\gamma}'$ into $\widehat{\mathcal{L}} \sqcup \widehat{\mathcal{L}}^c$: In order to define the set \mathcal{L} , we associate to each element $\widehat{\gamma}_i$ of $\widehat{\mathcal{L}}$ the maximal subpath γ_i of γ' with the endpoints that are C_1 -close to the endpoints of $\widehat{\gamma}_i$. We now apply Corollary 6.5, to obtain that the relative length of each such path γ_i in \mathcal{L} is smaller than some constant $K > 0$ which is dependent on the seize of C_1 but independent of all our choices.

We now define the collection \mathcal{L}^c of subpaths of γ' simply as those subpaths γ_j which connect the endpoints of the corresponding subsequent subpaths γ_i from \mathcal{L} as defined above. Of course, the γ_j may have length 0, or if two γ_i overlap, they may run in the opposite direction than γ' . But all this does not matter, as the concatenation of all subsequent paths from \mathcal{L} and \mathcal{L}^c clearly runs through all of γ' , and hence has bigger or equal relative length than γ' .

Now, by definition, for every path γ_j in \mathcal{L}^c there is a corresponding path $\widehat{\gamma}_j$ in $\widehat{\mathcal{L}}^c$ that has endpoints C_1 -close to the endpoints of γ_j . Since both, γ_j and $\widehat{\gamma}_j$ are absolute quasi-geodesics, since the relative length is always bounded above by the absolute length, i.e. $|\gamma_j|_{rel} \leq |\gamma_j|_{abs}$, and since we derived above $|\widehat{\gamma}_j|_{rel} = |\widehat{\gamma}_j|_{abs}$, there are constants $D_1, D_2 > 0$ such that

$$|\gamma_j|_{rel} \leq D_1 \cdot |\widehat{\gamma}_j|_{rel} + D_2$$

But the number of alternating subpaths from \mathcal{L} and \mathcal{L}^c is equal to that of $\widehat{\mathcal{L}}$ and $\widehat{\mathcal{L}}^c$ and thus bounded above by the relative length of $\widehat{\gamma}'$. Since the relative length of each γ_i in \mathcal{L} is bounded by the constant K , we obtain directly the existence of constants C_1, C_2 and C_3 as claimed above in (*). \square

8 Proof of the Main theorem

We first prove Proposition 1.4 as stated in the Introduction. The notion of a relative hyperbolic automorphism is recalled in Definition 2.6:

Proof of Proposition 1.4.

We consider the universal covering $\widetilde{\mathcal{G}}^2$ of the β -train track \mathcal{G}^2 from the β -train track representative $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ of α . We lift the relative length on edges to $\widetilde{\mathcal{G}}^2$ to make $\widetilde{\mathcal{G}}^2$ into a pseudo-metric space, and we pass over to the associated metric space $\widehat{\mathcal{G}}^2$ by contracting every edge of length 0. This amounts precisely to contracting every connected component \widetilde{X}_i of the full preimage \widetilde{X} of the relative part $X \subset \mathcal{G}^2$ to a single point \widehat{X}_i .

We now lift the train track map f to a map $\widetilde{f} : \widetilde{\mathcal{G}}^2 \rightarrow \widetilde{\mathcal{G}}^2$ which represents α in the following sense: For any $w \in \mathbb{F}_n$ and any point $P \in \widetilde{\mathcal{G}}^2$ one has:

$$\alpha(w)\widetilde{f}P = \widetilde{f}w(P)$$

Since f maps X to itself, the map \tilde{f} induces canonically a map $\hat{f} : \hat{\mathcal{G}}^2 \rightarrow \hat{\mathcal{G}}^2$ that satisfies similarly, for any $w \in \mathbb{F}_n$ and any point $P \in \hat{\mathcal{G}}^2$:

$$\alpha(w)\hat{f}P = \hat{f}w(P)$$

For our purposes below we also want, in addition to this “twisted commutativity property”, that \hat{f} fixes a vertex of $\hat{\mathcal{G}}^2$ outside of the union \hat{X} of all \hat{X}_i . To ensure this we apply property (e) of Theorem 4.2 and raise f to a sufficiently high power f^k in order to find a fixed point in the interior of an edge e of $\hat{\Gamma}$ (i.e. outside of X): We then subdivide edges finitely many times in order to make this fixed point into a f^k -fixed vertex of \mathcal{G}^2 . We then lift f^k to the map \hat{f}^k constructed above and compose it with the deck transformation action of a suitable element $v \in \mathbb{F}_n$ so that some lift of this f^k -fixed vertex is fixed by $v\hat{f}^k$. It follows that $v\hat{f}^k$ “twistedly commutes” with $\iota_v \alpha^k$ in the above meaning, where ι_v denotes the inner automorphisms $\iota_v : \mathbb{F}_n \rightarrow \mathbb{F}_n, w \mapsto v w v^{-1}$.

By virtue of Remark 2.7 we can continue to work with $v\hat{f}^k$ and $\iota_v \alpha^k$ rather than with \hat{f} and α as above, without loss of generality in our proof. However, for simplicity of notation we stick for the rest of the proof to \hat{f} and α , but we assume that \hat{f} has a fixed vertex $Q = \hat{f}(Q) \in \hat{\mathcal{G}}^2 \setminus \hat{X}$.

We now consider any generating system S of \mathbb{F}_n , and the associated coned Cayley graph $\Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$ as given in Definition 2.2. We define an \mathbb{F}_n -equivariant map

$$\psi : \Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n) \rightarrow \hat{\mathcal{G}}^2$$

by sending the base point $V(1)$ to the above \hat{f} -fixed vertex $Q \in \hat{\mathcal{G}}^2 \setminus \hat{X}$. Every cone vertex of $\Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$ is mapped to the corresponding contracted connected component \hat{X}_i of \hat{X} . The correspondence here is given through the subgroup of \mathbb{F}_n which stabilizes a cone vertex of $\Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$, since the same subgroup stabilizes also the “corresponding” contracted connected component \hat{X}_i of \hat{X} . Every edge e of $\Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$ is sent to an edge path $\psi(e)$ in $\hat{\mathcal{G}}^2$ of length $L(\psi(e)) > 0$: By construction no two distinct vertices of $\Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$ are mapped by ψ to the same vertex in $\hat{\mathcal{G}}^2$.

It follows that those edges of $\Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$ that are adjacent to the same cone vertex are mapped by ψ to edge paths that all have the same length. It is easy to see directly that the map ψ is a quasi-isometry (alternatively one can use Proposition 6.1 of [20]). Since we are only interested in estimating the distance of vertices (which are mapped by ψ again to vertices), and any distinct two vertices in either space have distance $\geq \frac{1}{2}$, we can suppress the additive constant in the quasi-isometry inequalities to obtain a constant $C > 0$ such that for all vertices $P, R \in \tilde{\Gamma}_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$ one has:

$$\frac{1}{C} d(P, R) \leq d(\psi(P), \psi(R)) \leq C d(P, R)$$

Similarly, the canonical inequalities obtained from Proposition 7.3, which describe that every normalized path in \mathcal{G}^1 lifts to a quasi-geodesic in $\hat{\mathcal{G}}^2$, will only be applied to edge paths

which are either of relative length 0 or are bounded away from 0 by 1 (= the length of the shortest edge in $\widehat{\mathcal{G}}^1$). Hence we obtain directly, for a suitable constant $A > 0$ and any two vertices $P, R \in \widehat{\mathcal{G}}^2$ that are connected by a normalized edge path $\gamma(P, R)$, the inequalities:

$$d(P, R) \leq |\gamma(P, R)|_{rel} \leq A d(P, R)$$

Thus we can calculate, for any $w \in \mathbb{F}_n$ and for $\lambda > 0$ as given in Corollary 6.10, for which we first assume that alternative (a) holds:

$$\begin{aligned} |w|_{S, \mathcal{H}} &= d(V(1), V(w)) \\ &\leq C d(\psi(V(1)), \psi(V(w))) \\ &\leq C |\gamma(\psi(V(1)), \psi(V(w)))|_{rel} \\ &\leq \frac{C}{\lambda} |\widetilde{f}^N(\gamma(\psi(V(1)), \psi(V(w))))_*|_{rel} \\ &\leq \frac{C}{\lambda} A d(\widetilde{f}^N(\psi(V(1))), \widetilde{f}^N(\psi(V(w)))) \\ &\leq \frac{AC}{\lambda} d(\widetilde{f}^N(Q), \widetilde{f}^N(wQ)) \\ &\leq \frac{AC}{\lambda} d(\widetilde{f}^N(Q), \alpha^N(w)\widetilde{f}^N(Q)) \\ &\leq \frac{AC}{\lambda} d(Q, \alpha^N(w)Q) \\ &\leq \frac{AC}{\lambda} d(\psi(V(1)), \psi(V(\alpha^N(w)))) \\ &\leq \frac{AC}{\lambda} C d(V(1), V(\alpha^N(w))) \\ &= \frac{AC^2}{\lambda} |\alpha^N(w)|_{S, \mathcal{H}} \end{aligned}$$

Since the constants A and C are independent of N , a sufficiently large choice of λ in Corollary 6.10 gives the desired conclusion (compare Definition 2.6).

The calculation for case (b) in Corollary 6.10 is completely analogous and not carried through here. The only additional argument to be mentioned here is to ensure the existence of a path γ' as in Corollary 6.10 (b). But this follows directly from the fact that the β -train track map $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ represents an automorphisms of \mathbb{F}_n , so that we can assume that f (and thus \widetilde{f}) is surjective: Otherwise one could replace \mathcal{G}^2 by a proper f -invariant subcomplex, and the corresponding restriction of f would be again a β -train track map which has otherwise the same properties as f . □

We can now give the proof of the main theorem of this paper as stated in the Introduction:

Proof of Theorem 1.1. From Proposition 3.3 we know that $\mathcal{H}(\alpha)$ is quasi-convex and malnormal. Thus Lemma 2.5 implies that \mathbb{F}_n is strongly hyperbolic relative to $\mathcal{H}(\alpha)$. Furthermore, from Proposition 1.4 we know that α is hyperbolic relative to $\mathcal{H}(\alpha)$. Hence Theorem 2.9 implies directly the claim. □

9 The Rapid Decay property

The *Rapid Decay property* (or *property (RD)*) was originally established by U. Haagerup for finitely generated free groups [23]. Indeed, it has also been called ‘‘Haagerup inequality’’

(compare [36]). The first to formalize and study systematically property (RD) was Jolissaint in [25]. Subsequent to his pioneering work, property (RD) was shown to hold for various classes of groups, in particular for hyperbolic groups [12], certain classes of groups acting on CAT(0)-complexes [10], relatively hyperbolic groups [14], etc.

The main importance of property (RD) comes from its applications to the Novikov conjecture and the Baum-Connes conjecture, see [11] and [26]. In particular, property (RD) is useful in constructing explicit isomorphisms in this branch of K-theory.

We review the basic definitions and results that we need to prove Corollary 1.2. Property (RD) may be stated in many equivalent ways. We borrow from [25] the definition which seems to be the simplest one for somebody with expertise in geometric group theory.

Consider any metric d on a group G which is equivariant with respect to the action of G on itself by left-multiplication. The function $L : G \rightarrow \mathbb{R}, g \mapsto d(1, g)$ is called a *length function on G* .

Such a length function, together with the choice of an exponent $s > 0$, is used to define a norm $\|\cdot\|_{2,s,L}$ on the group algebra $\mathbb{C}[G]$, which is given for any $\phi = \sum \phi(g)g \in \mathbb{C}[G]$ by:

$$\|\phi\|_{2,s,L} = \sqrt{\sum_{g \in G} \phi(g)\bar{\phi}(g)(1 + L(g))^{2s}}$$

We now consider the Hilbert space $l^2(G)$ and interpret any $\phi \in \mathbb{C}[G]$ as linear operator on $l^2(G)$, where the image of any $\psi \in l^2(G)$ is given by the convolution $\phi * \psi$, defined as usual by $(\phi * \psi)(g) = \sum_{h \in G} \phi(h)\psi(h^{-1}g)$. We can now consider also for any $\phi \in \mathbb{C}[G]$ the operator norm

$$\|\phi\| = \sup_{\psi \in l^2(G)} \frac{\|\phi * \psi\|_2}{\|\psi\|_2},$$

where $\|\psi\|_2 = \sqrt{\sum_{g \in G} \psi(g)\bar{\psi}(g)}$ denotes the classical l_2 -norm of $\psi \in l^2(G)$.

Definition 9.1 ([25]). A group G has *property (RD) with respect to a length function L* if there exist constants $c > 0$ and $s > 0$ such that, for any $\phi \in \mathbb{C}[G]$, one has:

$$\|\phi\| \leq c \|\phi\|_{2,s,L}$$

It is proved in [25] (Lemma 1.1.4 and Remark 1.1.7) that, if a group G has property (RD) with respect to the word-length function given by a finite generating set, then it has property (RD) with respect to any length function.

Definition 9.2. A finitely generated group G has *property (RD)* if it has property (RD) with respect to the word-length function defined by any finite generating set of G .

It is shown in [23] that finitely generated free groups have property (RD). Hence in the context of this section we don't need to refer to any other length function.

For any homomorphism β of a group G with finite generating set S one defines

$$a(\beta) = \max_{s \in S} |\beta(s)|_S,$$

where $|g|_S$ denotes the word length of $g \in G$ with respect to S . Following [25], for a second group Γ with finite generating system Σ , a homomorphism $\theta: \Gamma \rightarrow \text{Aut}(G)$ is said to have *polynomial amplitude*, if there exist positive numbers c and r such that, for any $\gamma \in \Gamma$ one has:

$$a(\theta(\gamma)) \leq c(1 + |\gamma|_\Sigma)^r$$

This notion is easily seen to be independent of the particular choice of the generating systems S and Σ . The following result has been shown by Jolissaint:

Proposition 9.3 ([25]). *Let G and Γ be two finitely generated groups, and let $\theta: \Gamma \rightarrow \text{Aut}(G)$ be a homomorphism with polynomial amplitude. If Γ and G have property (RD), then so does the semi-direct product $G \rtimes_\theta \Gamma$ defined by θ .*

We will also need the following theorem due to Drutu-Sapir:

Theorem 9.4 ([14]). *Let G be a group which is strongly hyperbolic relative to a finite family \mathcal{H} of finitely generated subgroups \mathcal{H}_j . If all the subgroups in \mathcal{H}_j have property (RD), then so does G .*

Proof of Corollary 1.2. We first consider the special case where $\alpha \in \text{Aut}(\mathbb{F}_n)$ is an automorphism of polynomial growth: in this case one deduces directly that the map $\theta: \mathbb{Z} \rightarrow \text{Aut}(\mathbb{F}_n)$, defined by $\theta(t) = \alpha^t$, has polynomial amplitude. Thus it follows from Proposition 9.3 that $\mathbb{F}_n \rtimes_\alpha \mathbb{Z}$ has property (RD).

We now consider an arbitrary automorphism $\alpha \in \text{Aut}(\mathbb{F}_n)$. From part (2) of Theorem 1.1 we know that $\mathbb{F}_n \rtimes_\alpha \mathbb{Z}$ is strongly hyperbolic, relative to the canonical family \mathcal{H}_α of mapping torus subgroups \mathcal{H}_j over subgroups $H_j \subset \mathbb{F}_n$ where the restriction of α has polynomial growth (see §3). By the above argument, each one of \mathcal{H}_j has property (RD). Thus we can conclude from Theorem 9.4 that $\mathbb{F}_n \rtimes_\alpha \mathbb{Z}$ as well must have property (RD). \square

References

- [1] M. Bestvina, M. Feighn and M. Handel. The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms. *Ann. of Math.* **151**, 517–623 (2000).
- [2] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.* **35**, 85–101 (1992). With an Addendum and correction in *J. Differential Geom.* **43**, 783–788 (1996).
- [3] M. Bestvina and M. Handel. Train tracks for automorphisms of the free group. *Annals of Math.* **135**, 1–51 (1992).

- [4] M. Bridson and D. Groves. The quadratic isoperimetric inequality for mapping tori of free group automorphisms I: Positive automorphisms. ArXiv math.GR/0211459.
- [5] M. Bridson and D. Groves. Free-group automorphisms, train tracks and the beaded decomposition. ArXiv math.GR/0507589.
- [6] M. Bridson and D. Groves. The quadratic isoperimetric inequality for mapping tori of free group automorphisms II: The general case. ArXiv math.GR/0610332.
- [7] P. Brinkmann. Hyperbolic automorphisms of free groups. *Geom. Funct. Anal.* **10**, 1071–1089 (2000).
- [8] B. H. Bowditch. Relatively hyperbolic groups, Preprint 1999. (<http://www.maths.soton.ac.uk/staff/Bowditch/preprints.html>)
- [9] I. Bumagin. On definitions of relatively hyperbolic groups, in “Geometric methods in group theory”. *Contemporary Mathematics series* **372**, 189–196 (2005).
- [10] I. Chatterji and K. Ruane. Some geometric groups with Rapid Decay. *Geom. and Funct. Anal.* **15**, 311–339 (2005).
- [11] A. Connes and H. Moscovici. Cyclic cohomology, the Novikov conjecture and hyperbolic groups. *Topology* **29**, 345–388 (1990).
- [12] P. de la Harpe. Groupes hyperboliques, algèbres d’opérateurs et un théorème de Jolissaint. *Comptes Rendus de l’Académie des Sciences* **307**, 771–774 (1988).
- [13] C. Drutu. Relatively hyperbolic groups: geometry and quasi-isometric invariance. ArXiv math.GR/0605211
- [14] C. Drutu and M. Sapir. Relatively Hyperbolic Groups with Rapid Decay Property. *Int. Math. Res. Notices* **19**, 1181–1194 (2005).
- [15] B. Farb. Relatively hyperbolic groups. *Geom. Funct. Anal.* **8**, 1–31 (1998).
- [16] M. Feighn and M. Handel. The Recognition Theorem for $\text{Out}(\mathbb{F}_n)$. ArXiv math.GR/0612702
- [17] D. Gaboriau, A. Jäger, G. Levitt, and M. Lustig. An index for counting fixed points of automorphisms of free groups. *Duke Math. J.* **93**, 425–45 (1998).
- [18] D. Gaboriau and G. Levitt. The rank of actions on \mathbf{R} -trees. *Ann. Sci. École Norm. Sup.* **28**, 549–570, 1995.
- [19] F. Gautero. Hyperbolicity of mapping-torus groups and spaces. *L’Enseignement Mathématique* **49**, 263–305 (2003).

- [20] F. Gautero and M. Lustig. Relative hyperbolization of (one-ended hyperbolic)-by-cyclic groups. *Math. Proc. Camb. Phil. Soc.* **137**, 595–611 (2004).
- [21] F. Gautero. Geodesics in trees of hyperbolic and relatively hyperbolic groups. Preprint 2005 (revised 2007). ArXiv math.GR/0710.4079v2
- [22] M. Gromov. Hyperbolic groups. In “Essays in Group Theory”. *Math. Sci. Res. Inst. Publ.* **8**, 75–263 (1987).
- [23] U. Haagerup. An example of non-nuclear C^* -algebra which has the metric approximation property. *Inventiones Mathematicae* **50**, 279–293 (1979).
- [24] A. Jäger and M. Lustig. Free groups automorphisms with many fixed points at infinity. Preprint 1998 (revised 2007). To appear in: *Geometry & Topology Monographs* **14**, 101–113 (2008).
- [25] P. Jolissaint. Rapidly Decreasing Functions in Reduced C^* -Algebras of Groups. *Transactions of the A.M.S.* **317**, 167–196 (1990).
- [26] V. Lafforgue. KK-théorie bivariante pour les algèbres de Banach et conjecture de Baum-Connes. *Inventiones Mathematicae* **149**, 1–95 (2002).
- [27] G. Levitt. Counting growth types of automorphisms of free groups. ArXiv math.GR/0801.4844
- [28] M. Lustig. Structure and conjugacy for automorphisms of free groups I. *MPI-Bonn preprint series 2000*, No. 241. (<http://www.mpim-bonn.mpg.de/preprints>)
- [29] M. Lustig. Structure and conjugacy for automorphisms of free groups II. *MPI-Bonn preprint series 2001*, No. 4. (<http://www.mpim-bonn.mpg.de/preprints>)
- [30] M. Lustig, *Conjugacy and centralizers for iwip automorphisms of free groups*, Geometric Group Theory, Trends in Mathematics, pp. 197–224, Birkhäuser, Basel 2007
- [31] M. Lustig. Better train tracks. Preliminary preprint, 2006.
- [32] N. Macura. Quadratic isoperimetric inequality for mapping tori of polynomially growing automorphisms of free groups. *Geom. Funct. Anal.* **10**, 874–901 (2000).
- [33] Mj. Mahan and L. Reeves. A combination theorem for strong relative hyperbolicity. ArXiv math.GR/0611601.
- [34] D. V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties and algorithmic problems. *Memoirs of the American Mathematical Society* **179**, 2006.

- [35] A. Szczepański. Relatively hyperbolic groups. *Michigan Math. Journal* **45**, 611–618 (1998).
- [36] M. Talbi. Inégalité de Haagerup et Géométrie des Groupes. Ph.D.-Thesis, Université de Lyon I (2001).

10 Appendix to “The mapping torus group of a free group automorphism is hyperbolic relative to the canonical subgroups of polynomial growth”

by Martin Lustig

Bestvina-Handel have proved in [3] that every automorphism of \mathbb{F}_n can be represented by a relative train track map $f : \Gamma \rightarrow \Gamma$, recalled below in subsection 10.4. The goal of this appendix is to explain how one can derive from such a relative train track map a β -train track map as defined in section 4. We give here in a detailed and careful manner all ingredients needed in this construction, and we sketch the proofs. A fully expanded version of this appendix will be given in the forthcoming paper [31].

10.1 Partial train track maps and Nielsen faces

A *graph-of-spaces* \mathcal{G} relative to X consists of a finite collection X of pathwise connected *vertex spaces* X_v , and a finite collection $\widehat{\Gamma}$ of edges with endpoints in X . We call X the *relative part* of \mathcal{G} , and the edges in $\widehat{\Gamma}$ are referred to as *edges of* \mathcal{G} (by which we exclude possible edges in X !). A subpath γ_0 of a path γ in \mathcal{G} is called *backtracking* if the endpoints of γ_0 coincide, and if the resulting loop is contractible in \mathcal{G} . A path γ in \mathcal{G} is called *reduced* (rel. X) if every backtracking subpath γ_0 of γ is contained in X .

Definition 10.1. Let \mathcal{G} be a graph-of-spaces with vertex space collection $X \subset \mathcal{G}$, and let $f : \mathcal{G} \rightarrow \mathcal{G}$ be a continuous map with $f(X) \subset X$.

- (1) A path γ in \mathcal{G} is called *legal* (with respect to f) if for every $t \geq 1$ the path $f^t(\gamma)$ is reduced.
- (2) The map $f : \mathcal{G} \rightarrow \mathcal{G}$ is a *partial train track map relative to X* if every edge of \mathcal{G} is legal. In this case the collection $\widehat{\Gamma}$ of edges in \mathcal{G} is called the *train track part* of \mathcal{G} .
- (3) The map f is called *expanding* if for every edge $e \in \widehat{\Gamma}$ some iterate $f^t(e)$ runs over 2 or more edges from $\widehat{\Gamma}$.

A path η in \mathcal{G} is called an *indivisible Nielsen path (INP)* if $f(\eta)$ is homotopic rel. endpoints to η , and if η is a concatenation $\eta = \gamma \circ \gamma'$ of two legal paths γ and γ' which are not contained in X . Note that η can not be legal, if f is expanding. Note also that the endpoints of η may be situated in the interior of an edge of $\widehat{\Gamma}$. A path η is called *periodic indivisible Nielsen path (periodic INP)* if η is an INP for some positive iterate f^t of f . We do not distinguish between periodic INP's that are homotopic via a homotopy (not necessarily fixing endpoints) that takes entirely place in the relative part of \mathcal{G} .

Definition 10.2. Let $f : \mathcal{G} \rightarrow \mathcal{G}$ be a partial train track map relative to the collection X of vertex spaces of a graph-of-spaces \mathcal{G} .

- (a) A concatenation $\gamma \circ \gamma'$ of two paths γ and γ' in \mathcal{G} is called *legal concatenation* if there exists a terminal subpath γ_0 of γ and an initial subpath γ'_0 of γ' , both not entirely contained in X , such that the concatenation $\gamma_0 \circ \gamma'_0$ is legal.
- (b) A path γ in \mathcal{G} is called *pseudo-legal* if γ is a legal concatenation of legal paths and periodic INP's.

It is known that every INP of $f : \mathcal{G} \rightarrow \mathcal{G}$ defines a branch point orbit in an \mathbb{R} -tree with isometric $\pi_1\mathcal{G}$ -action, which can be obtained from the partial train track via a row-eigenvector of the geometric transition matrix of f (compare §4 of [30] and the references given there). If $\pi_1\mathcal{G}$ is a free group \mathbb{F}_n of finite rank n , then the number of such branchpoints and their “multiplicity” is bounded in terms of n , see [18]. As a consequence, one obtains:

Proposition 10.3. *For any expanding partial train track map $f : \mathcal{G} \rightarrow \mathcal{G}$, with finitely generated free group $\pi_1\mathcal{G}$, there are only finitely many periodic INP's in \mathcal{G} .*

The following proposition has been shown in §3 of [28].

Proposition 10.4. (a) *Let \mathcal{G} be a graph-of-spaces with vertex space collection X , and let $f : \mathcal{G} \rightarrow \mathcal{G}$ be an expanding partial train track map relative to X . Then for every path γ in \mathcal{G} there is an exponent $t(\gamma) \geq 1$ such that $f^{t(\gamma)}(\gamma)$ is homotopic rel. endpoints to a pseudo-legal path.*

(b) *There is an upper bound to the exponent $t(\gamma)$, which depends only on the number q of factors in any decomposition of $\gamma = \gamma_1 \circ \dots \circ \gamma_q$ as concatenation of legal paths γ_i , and not on the particular choice of γ itself.*

Definition 10.5. Let $f : \mathcal{G} \rightarrow \mathcal{G}$ be an expanding partial train track map relative to the vertex space collection X of a graph-of-spaces \mathcal{G} .

- (a) Let η be an INP of f . Then gluing an *auxiliary edge* e along $\partial e = \partial\eta$ to \mathcal{G} , and simultaneously a *Nielsen face*, i.e. a 2-cell Δ^2 , along the boundary path $\partial\Delta^2 = \eta^{-1} \circ e$, is called *expanding a Nielsen face at the INP η* . One extends the map f by the identity on e and by mapping Δ^2 to $\Delta^2 \cup f(\eta)$, to obtain a *partial train track map with Nielsen faces* $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ relative to \widehat{X} , for the resulting space $\mathcal{G}^2 = \mathcal{G} \cup e \cup \Delta^2$ and $\widehat{X} = X \cup e$.
- (b) Similarly one defines the expansion of an f -orbit of Nielsen faces at the f -orbit of a periodic INP. It is also possible to expand Nielsen faces at several periodic INP's simultaneously.
- (c) Every Nielsen face Δ^2 , expanded together with an auxiliary edge e at some (periodic) INP η , defines a homotopy which deforms a path that runs over e to a path which runs over η instead. Thus the collection of Nielsen faces that have been expanded at (periodic) INP's of \mathcal{G} defines a strong deformation retraction $\widehat{r} : \mathcal{G}^2 \rightarrow \mathcal{G}$, which maps every auxiliary edge e_i to the corresponding (periodic) INP η_i and leaves every point of \mathcal{G} fixed.

- (d) A path γ in $\mathcal{G}^1 = \widehat{\Gamma} \cup \widehat{X} \subset \mathcal{G}^2$ is called *strongly reduced* if $\widehat{r}(\gamma)$ is reduced in \mathcal{G} (relative to X). It follows that in this case γ is also reduced, as path in \mathcal{G}^1 relative to \widehat{X} , but also as path in \mathcal{G}^2 relative to \widehat{X} . (The subtle difference here is caused by the above definition of a “backtracking subpath”: a subpath of γ may well be backtracking in \mathcal{G}^2 but not in \mathcal{G}^1 !)
- (e) A path γ in $\mathcal{G}^1 \subset \mathcal{G}^2$ is called *strongly legal* if for any $t \geq 1$ the path $f^t(\gamma)$ is strongly reduced. Note that every strongly legal path is legal.
- (f) The train track map $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ is said to be *strong* if every edge of the train track part $\widehat{\Gamma} \subset \mathcal{G}^2$ is strongly legal.

Remark 10.6. Let $f : \mathcal{G} \rightarrow \mathcal{G}$ be a partial train track map relative to a collection $X \subset \mathcal{G}$ of vertex spaces, and let $\widehat{f} : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ be obtained from f and \mathcal{G} by expanding Nielsen faces at finitely many (periodic) INP’s.

- (a) Then the restriction $f_1 : \mathcal{G}^1 \rightarrow \mathcal{G}^1$ of \widehat{f} to the union \mathcal{G}^1 of \mathcal{G} with all added auxiliary edges is a partial train track map (without Nielsen faces !) relative to \widehat{X} , where the fundamental group $\pi_1 \mathcal{G}^1$ has been increased with respect to $\pi_1 \mathcal{G}$ through adding the auxiliary edges to the relative part X to get \widehat{X} .
- (b) If at every periodic INP in \mathcal{G} a Nielsen face has been expanded to obtain \mathcal{G}^2 , then every pseudo-legal path in \mathcal{G} is homotopic in \mathcal{G}^2 to a strongly legal path.

For notational purposes we now extend the above introduced notation slightly. Note however that all 2-cells attached to partial train tracks in our context will be Nielsen faces, but in a iterated fashion which for notational convenience we prefer to suppress.

Definition 10.7. Let $\mathcal{G}^1 = \widehat{\Gamma} \cup \widehat{X}$ be a graph-of-spaces relative to \widehat{X} , and let \mathcal{G}^2 be obtained from \mathcal{G}^1 by attaching finitely many 2-cells along their boundary to \mathcal{G}^1 . We call \mathcal{G}^2 a *graph-of-spaces with 2-cells*.

For some subspace $X \subset \widehat{X}$, which contains all endpoints of edges of $\widehat{\Gamma}$, let $\mathcal{G} = \widehat{\Gamma} \cup X$ be a graph-of-spaces (without 2-cells !) relative to X , and assume that $r : \mathcal{G}^2 \rightarrow \mathcal{G}$ is a strong deformation retraction.

A map $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ is a *strong train track map with 2-cells, relative to \widehat{X}* , if

- (a) $f(\mathcal{G}^1) \subset \mathcal{G}^1$ and $f(\widehat{X}) \subset \widehat{X}$,
- (b) the induced map $f_1 : \mathcal{G}^1 \rightarrow \mathcal{G}^1$ is an expanding train track map relative to \widehat{X} , and
- (c) every edge $e \in \widehat{\Gamma}$ is *strongly legal* with respect to r : For any $t \geq 1$ the path $r f_1^t(e)$ is reduced (rel. X) in \mathcal{G} .

10.2 Building up strong train tracks

The following lemma is an important tool in the proof of our main result presented in subsection 10.4. A proof appeared already in [29] (at the end of §6), modulo a minor switch in the terminology employed.

Lemma 10.8. *Let $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ be a strong partial train track map of a graph-of-spaces \mathcal{G}^2 with 2-cells, relative to a subspace $X \subset \mathcal{G}^2$.*

- (a) *Let \mathcal{G}_1^2 and $f_1 : \mathcal{G}_1^2 \rightarrow \mathcal{G}_1^2$ be obtained from \mathcal{G}^2 by attaching a further edge e at its endpoints to points $P_0, P_1 \in X$, and by extending f via $f_1(e) = \gamma_0 \circ e \circ \gamma_1$, where the γ_i are strongly legal paths in \mathcal{G} . Then one can homotope in \mathcal{G}^2 the attaching points P_0 and P_1 to points P'_0 and P'_1 , and replace γ_0 and γ_1 correspondingly by homotopic paths γ'_0 and γ'_1 , so that e becomes strongly legal.*

In particular, the resulting map $f'_1 : \mathcal{G}_1'^2 \rightarrow \mathcal{G}_1'^2$ is a strong partial train track map with 2-cells, relative to X , provided that at least one of the paths γ'_i is not entirely contained in X (in order to ensure that f'_1 is expanding).

- (b) *The analogous statement is true if \mathcal{G}_1^2 is constructed from \mathcal{G}^2 by attaching finitely many edges e^j to X , with $f_1(e^j) = \gamma_0^j \circ e^{\pi(j)} \circ \gamma_1^j$ for some permutation π .*
- (c) *The above given homotopies lead canonically to a homotopy equivalence $h : \mathcal{G}_1^2 \rightarrow \mathcal{G}_1'^2$, which restricts on \mathcal{G}^2 to a selfmap that is homotopic to the identity and satisfies $h(X) \subset X$, such that $f'_1 h$ and $h f_1$ are homotopic.*

Definition 10.9. A partial train track map $f : \mathcal{G} \rightarrow \mathcal{G}$ relative X satisfies the *initial-segments condition* if for every edge e of the train track part of \mathcal{G} some initial and some terminal segment of e are mapped by f onto an edge of the train track part of \mathcal{G} .

The following lemma can be derived in a direct manner from the above definitions.

Lemma 10.10. *Let $f : \mathcal{G} \rightarrow \mathcal{G}$ be a partial train track map of a graph-of-spaces \mathcal{G} relative to the vertex space collection X of \mathcal{G} , which satisfies the initial-segments condition, and with the property that all edges of \mathcal{G} are attached to some subspace $X' \subset X$.*

Assume that X is itself a graph-of-spaces $X = \mathcal{G}'^2$ with 2-cells, relative to X' , and assume that the restriction f' of f to $X = \mathcal{G}'^2$ is a strong partial train track map with 2-cells, relative to X' .

Assume furthermore that for every edge e of \mathcal{G} any subpath γ of $f(e)$ that is entirely contained in $X = \mathcal{G}'$ is strongly legal.

Then \mathcal{G} is a graph-of-spaces with 2-cells, relative to X' , and f is a strong partial train track map with 2-cells, relative to X' .

10.3 The attaching-iteration method

In this subsection we consider pairs of (not necessarily connected) spaces $X \subset Y$ and maps $f : Y \rightarrow Y$ which satisfy $f(X) \subset X$. Assume that Y is obtained from the disjoint union of X and a space Z by gluing a subspace $Z_0 \subset Z$ to X via an attaching map $\phi : Z_0 \rightarrow X$:

$$Y = Z \cup X / \langle z = \phi(z) \mid z \in Z_0 \rangle$$

Let X' be an f -invariant union of some of the connected components of X , and let $Z'_0 \subset Z_0$ the subset consisting of those points that are glued via ϕ to X' . We can now construct a new space in the following way: We fix an integer $t \geq 1$. Then we unglue every point $z' \in Z'_0$ from X' and reglue it to $f^t(z')$, to obtain a new space

$$Y_1 = Z \cup X / \langle z = \phi(z), z' = f^t \phi(z') \mid z \in Z_0 \setminus Z'_0, z' \in Z'_0 \rangle .$$

We define a map $h : Y \rightarrow Y_1$ which restricts to the identity on the subspace Z as well as on $X \setminus X'$, and maps every point $x \in X'$ to the point $f^t(x)$. It is easy to verify that these definitions are compatible with the gluing maps. We observe:

Remark 10.11. If the restriction of f to a self-map of X' is a homotopy equivalence, then also the map h is a homotopy equivalence.

We now define a map $f_1 : Y_1 \rightarrow Y_1$ as follows, where we distinguish three cases according to the position of the point $y_1 \in Y_1$ and to that of $f(y) \in Y$, where y denotes the point “corresponding” to y_1 in the identical copy in Y of the subspace $X \subset Y_1$ or $(Z \setminus Z_0) \subset Y_1$. The case $y_1 \in Z_0$ can be discarded, as any such y_1 is identified via the map ϕ with some point of X . We define:

1. If $y_1 \in X' \subset Y_1$, then we set $f_1(y_1) = f(y)$.
2. If $y_1 \in (Z \setminus Z_0) \cup (X \setminus X') \subset Y_1$ and $f(y) \in X' \subset Y$, then we set $f_1(y_1) = f^{t+1}(y)$.
3. If $y_1 \in (Z \setminus Z_0) \cup (X \setminus X') \subset Y_1$ and $f(y) \in Y \setminus X'$, then we set $f_1(y_1) = f(y)$.

These definitions extend continuously to define a map on Z_0 which is compatible with the gluing maps, and hence one obtains directly a well defined map $f_1 : Y_1 \rightarrow Y_1$ with $f_1(X) \subset X$. The proof of the following proposition is now an exercise:

Proposition 10.12. *Let $f : Y \rightarrow Y, f_1 : Y_1 \rightarrow Y_1$ and $h : Y \rightarrow Y_1$ be as above.*

(1) *The maps f and f_1 commute via h :*

$$hf = f_1h$$

(2) *If f and h are homotopy equivalences, then so is f_1 .*

- (3) If Y is a graph-of-spaces with vertex space collection X , and if f is a partial train track map relative X , then so are Y_1 and f_1 .

Remark 10.13. Notice that the assumption in Remark 10.11, that f restricts to a homotopy equivalence of X' , is necessary in order to get a homotopy equivalence h as above, with $hf = f'h$. If one is content with a more general map h which satisfies this equation, but is only an isomorphism on π_1 , then the weaker assumption suffices that f induces a π_1 -isomorphism on each connected component of X' . It is, however, unavoidable that X' contains no *inessential* component X_v of X , i.e. X_v satisfies $\pi_1 X_v = \{1\}$. Otherwise $\pi_1 Y_1$ would be different from $\pi_1 Y$.

In the context considered below it turns out that case 3. in the definition of the map f_1 before Proposition 10.12 does never occur. In order to simplify the notation, we define:

Definition 10.14. Let $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ be a partial train track map with 2-cells relative $X \subset \mathcal{G}^2$. A connected component of \mathcal{G}^2 is called *essential*, if it is mapped by f via a homotopy equivalence to another connected component. A connected component is called *pre-essential* if it is mapped by f to an essential component.

We now want to further specify the particular application of the attaching-iteration method that will be used in the next subsection, to construct strong partial train track maps via an iterative procedure. To be specific, other than Proposition 10.12 we also use Proposition 10.4 and Remark 10.6 (b) to obtain part (4) of the following:

Corollary 10.15. Let $f' : \mathcal{G}'^2 \rightarrow \mathcal{G}'^2$ be a strong partial train track map with 2-cells, relative to a subspace $X' \subset \mathcal{G}'^2$, and assume that in \mathcal{G}'^2 Nielsen faces have been expanded at every periodic INP. Assume furthermore that every connected component of \mathcal{G}'^2 is either essential or pre-essential.

Let \mathcal{G}_0 be a graph-of-spaces relative to \mathcal{G}'^2 , where \mathcal{G}_0 is obtained from \mathcal{G}'^2 by attaching a finite collection $\widehat{\Gamma}_0$ of edges to X' . Let $f_0 : \mathcal{G}_0 \rightarrow \mathcal{G}_0$ be an extension of the map f' , which is a partial train track map relative to \mathcal{G}'^2 .

Then there exists a graph-of-spaces \mathcal{G} , given by attaching a collection $\widehat{\Gamma}$ of edges to a collection X of vertex spaces, as well as a partial train track map $f : \mathcal{G} \rightarrow \mathcal{G}$ relative to X , which have the following properties:

- (1) There is a homotopy equivalence $h : \mathcal{G}_0 \rightarrow \mathcal{G}$ with $h(\mathcal{G}'^2) \subset X$ such that $f_0 h$ is homotopic to $h f$. The map h restricts to a homeomorphism $h_\Gamma : \widehat{\Gamma}_0 \rightarrow \widehat{\Gamma}$, with $f h(x) = h f_0(x)$ for all points $x \in \widehat{\Gamma}_0$ with $f_0(x) \in \widehat{\Gamma}_0$.
- (2) In particular, if the partial train track map (rel. \mathcal{G}'^2) f_0 satisfies the initial-segments condition, then so does f .
- (3) There is a homeomorphism $\psi : \mathcal{G}'^2 \rightarrow X$ and an integer $t \geq 0$, such that on every essential component of \mathcal{G}'^2 the map h is equal to $\psi f'^t$, while on every pre-essential

component of \mathcal{G}'^2 the map h is equal to ψ . Moreover, the restriction $f_X : X \rightarrow X$ of f is a strong partial train track map with 2-cells relative to $\psi(X')$, and f_X is homotopic to $\psi f' \psi^{-1}$ on every essential component and to $\psi f'^{t+1} \psi^{-1}$ on every pre-essential component of X .

- (4) The image $f(e)$ of any edge e of $\widehat{\Gamma}$ is an alternate concatenation of subpaths that are either contained in $\widehat{\Gamma}$, or else they are strongly legal paths in X .

The above proposition is a crucial step in our iterative procedure given in subsection 10.4 to build β -train track maps from relative train track maps. The reader should be warned, however, that the resulting map $f : \mathcal{G} \rightarrow \mathcal{G}$ is not necessarily yet a partial train track map relative to $\psi(X')$. This conclusion would in general be wrong, despite of the fact that f is a partial train track map relative to X , that the restriction of f to X is a partial train track map relative to $\psi(X')$, and that property (4) of the above proposition holds: A priori, it may still happen that a path $f^t(e)$ is not reduced relative to $\psi(X')$, for some edge e of $\widehat{\Gamma}$ and some $t \geq 1$.

10.4 Construction of a β -train track map from a relative train track map

Bestvina-Handel have proved in [3] that for every automorphism α of \mathbb{F}_n there exists a finite connected graph Γ with identification $\pi_1 \Gamma = \mathbb{F}_n$, such that α can be represented by a relative train track map $g : \Gamma \rightarrow \Gamma$. This means, using the terminology introduced in the previous subsections, that there is an g -invariant filtration $\Gamma_0 \subset \Gamma_1 \subset \dots \subset \Gamma_s = \Gamma$ of (not necessarily connected) subgraphs, where Γ_0 is the vertex set of Γ , such that the following conditions are satisfied:

Properties 10.16. For any $k \in \{1, \dots, s\}$ we denote by g_k the restriction of g to the g -invariant subgraph $\Gamma_k \subset \Gamma$.

- (1) The map $g_k : \Gamma_k \rightarrow \Gamma_k$ is a partial train track map relative to Γ_{k-1} .
- (2) If the map g_k is expanding (relative to Γ_{k-1}), then it satisfies the initial-segments condition (see Definition 10.9).
- (3) If the map g_k is non-expanding (relative to Γ_{k-1}), then either all of Γ_k is mapped by f_k to Γ_{k-1} , or else g_k is, modulo Γ_{k-1} , a transitive permutation of the edges of $\Gamma_k \setminus \Gamma_{k-1}$.

We now describe the construction that derives from such a relative train track map $g : \Gamma \rightarrow \Gamma$ a β -train track map $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$. In a first attempt we concentrate on the weaker property that f is a strong partial train track map with 2-cells. Subsequently we show that the construction defined below yields indeed a map that satisfies the additional properties claimed in Theorem 4.2.

Our construction proceeds iteratively, moving at each step one level up, i.e. from g_{k-1} to g_k . To start this iterative process, note that for $k = 1$ the map $g_1 : \Gamma_1 \rightarrow \Gamma_1$ is an

absolute train track map (in the sense of [3]), since the relative part Γ_0 consists precisely of the vertices of Γ . Thus g_1 is in particular a strong partial train track map with 2-cells (where the retraction r is simply the identity map).

Let us now assume, by induction, that there is a strong partial train track map with 2-cells $f_{k-1} : \mathcal{G}_{k-1}^2 \rightarrow \mathcal{G}_{k-1}^2$, relative to a subspace X_{k-1} , as well as a homotopy equivalence $h_{k-1} : \Gamma_{k-1} \rightarrow \mathcal{G}_{k-1}^2$ such that $h_{k-1}g_{k-1}$ is homotopic to $f_{k-1}h_{k-1}$. We then consider a copy $\widehat{\Gamma}$ of the edges of $\Gamma_k \setminus \Gamma_{k-1}$, and we attach each edge \widehat{e} of $\widehat{\Gamma}$, with copy e in $\Gamma_k \setminus \Gamma_{k-1}$, at the points $h_{k-1}(\partial e)$ to \mathcal{G}_{k-1}^2 . We define the map h_k on Γ_k to agree with h_{k-1} on Γ_{k-1} , and every edge e of $\Gamma_k \setminus \Gamma_{k-1}$ is mapped by h_k to its copy \widehat{e} in $\widehat{\Gamma}$.

We define the map f_k on $\widehat{\Gamma} \cup \mathcal{G}_{k-1}^2$ to agree with f_{k-1} on \mathcal{G}_{k-1}^2 , and for each edge $\widehat{e} \in \widehat{\Gamma}$ we define $f_k(\widehat{e})$ to be the concatenation $\gamma_0 \circ h_k g_k(e) \circ \gamma_1$. Here the γ_i are the paths traced out by the points $h_{k-1}(\partial e)$ during the homotopy between $h_{k-1}g_{k-1}$ and $f_{k-1}h_{k-1}$. Thus we obtain a graph-of-spaces $\mathcal{G}_k = \widehat{\Gamma} \cup \mathcal{G}_{k-1}^2$ with vertex space collection \mathcal{G}_{k-1}^2 , a map $f_k : \mathcal{G}_k \rightarrow \mathcal{G}_k$ with $f_k(\mathcal{G}_{k-1}^2) \subset \mathcal{G}_{k-1}^2$, and a homotopy equivalence $h_k : \Gamma_k \rightarrow \mathcal{G}_k$ with $h_k(\Gamma_{k-1}) \subset \mathcal{G}_{k-1}^2$, such that $h_k g_k$ is homotopic to $f_k h_k$.

As first step in our iterative construction we expand Nielsen faces in \mathcal{G}_{k-1}^2 , until at all periodic INP's of the partial train track map f_{k-1} rel. X_{k-1} there is a Nielsen face attached. All auxiliary edges introduced in this procedure are added to the relative subspace X_{k-1} . We extend the strong deformation retraction r_{k-1} on \mathcal{G}_{k-1}^2 , which exists by induction, by precomposing it with the strong deformation retraction \widehat{r}_k which is defined as is the map \widehat{r} in Definition 10.5 (c): every auxiliary edge is pushed over the corresponding Nielsen face, and any of the original points of \mathcal{G}_k^2 is left fixed.

Next we apply Corollary 10.15, to obtain that the image $f_k(\widehat{e})$ of any edge \widehat{e} in $\widehat{\Gamma}$ is an alternating concatenation of subpaths in $\widehat{\Gamma}$ and of strongly legal paths in \mathcal{G}_{k-1}^2 (relative to X_{k-1}). For simplicity we keep the same names, thus suppressing notationally the homotopy equivalence h as well as the homeomorphisms h_Γ and ψ from Corollary 10.15.

Proposition 10.17. *The resulting map f_k is (after a homotopically irrelevant modification in Case 4 of the proof) a strong partial train track map $f_k : \mathcal{G}_k^2 \rightarrow \mathcal{G}_k^2$ with 2-cells, relative to the subspace $X_k \subset \mathcal{G}_k^2$.*

The relative part X_k is equal to X_{k-1} , except for Case 3 of the proof, where X_k is specified in the proof.

Proof. We will distinguish four cases as follows:

CASE 1: Assume that f_k is expanding (relative to \mathcal{G}_{k-1}^2). In this case we have the initial-segments condition given as hypothesis by Property 10.16 (2), so that we can apply Lemma 10.10 to obtain directly the statement of the above proposition, for $X_k := X_{k-1}$.

CASE 2: If f_k is not expanding (relative to \mathcal{G}_{k-1}^2), and if $f_k(\widehat{\Gamma})$ is entirely contained in \mathcal{G}_{k-1}^2 but not in X_{k-1} , then after the above application of Corollary 10.15 each of the paths $f_k(\widehat{e})$ is a strongly legal path in \mathcal{G}_{k-1}^2 , so that again the claim follows directly, for $X_k := X_{k-1}$.

CASE 3: If f_k is not expanding (relative to \mathcal{G}_{k-1}^2), and if $f_k(\widehat{\Gamma})$ is entirely contained in

$\widehat{\Gamma} \cup X_{k-1}$, we define $X_k = \widehat{\Gamma} \cup X_{k-1}$ and obtain again directly the above claim. Note that in this case the train track part of Γ_k grows polynomially under iteration of f .

CASE 4: Assume that f_k is not expanding (relative to \mathcal{G}_{k-1}^2), but that $f_k(\widehat{\Gamma})$ is not entirely contained in \mathcal{G}_{k-1}^2 , and also not in $\widehat{\Gamma} \cup X_{k-1}$. In this case for every edge e_i of $\widehat{\Gamma}$ one has $f_k(e_i) = \gamma_0^i \circ e^{\pi(i)} \circ \gamma_1^i$, for some permutation π of the edges of $\widehat{\Gamma}$ and legal paths γ_0^i, γ_1^i in \mathcal{G}_{k-1}^2 relative to X_{k-1} . Furthermore, not all of the γ_0^i, γ_1^i are contained in X_{k-1} (or else we would be in Case 3). Thus we can apply Lemma 10.8 to define a modification of \mathcal{G}_k and f_k (which is homotopically trivial rel. X_{k-1}), and with this modification our claim is now proved by Lemma 10.8, for $X_k := X_{k-1}$. \square

We now verify inductively that the additional properties from Theorem 4.2 are satisfied, i.e. that the map f_k is indeed a β -train track map. To be precise, for this purpose one has first to expand further Nielsen faces in \mathcal{G}_k^2 until at every periodic INP of f_k a Nielsen face is expanded. (Note that this is anyway the first modification in the above described iterative construction to produce strong partial train tracks with 2-cells, when passing to the next level, with index $k + 1$).

In particular, this expansion of Nielsen faces, at every periodic INP, has to be done at the very last level of our iterative procedure, to obtain from $f_s : \mathcal{G}_s^2 \rightarrow \mathcal{G}_s^2$ a β -train track representative $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ of the same automorphism of \mathbb{F}_n that was originally represented by the relative train track map $g : \Gamma \rightarrow \Gamma$.

We state the following theorem for the map f , but we prove it via induction by passing from f_{k-1} to f_k .

Theorem 10.18. *The strong partial train track map with 2-cells $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$, relative to X , is a β -train track map, in that it has the following additional properties:*

- (1) *The map f is expanding relative to X .*
- (2) *The map f has polynomial growth on the relative part X .*
- (3) *Every 2-cell of \mathcal{G}^2 is a Nielsen face that has been expanded at a periodic INP of some of the strong partial train track maps $f_k : \mathcal{G}_k^2 \rightarrow \mathcal{G}_k^2$ rel. X_k , in the iterative construction for any of the steps $k = 1, \dots, s$.*
- (4) *The above defined retraction $r = r_s \circ \widehat{r}$ maps \mathcal{G}^2 to a subgraph Γ' of \mathcal{G}_k^2 , such that the restriction $r \circ f|_{\Gamma'} : \Gamma' \rightarrow \Gamma'$ is a relative train track map with respect to the iteratively defined filtration $\Gamma'_0 \subset \Gamma'_1 \subset \dots \subset \Gamma'_s$. The graph Γ' is obtained from the 1-skeleton of \mathcal{G}^2 by omitting all auxiliary edges introduced when expanding a Nielsen face for any of the intermediate partial train track maps $f_k : \mathcal{G}_k \rightarrow \mathcal{G}_k$ rel. \mathcal{G}_{k-1}^2 , for $k = 1, \dots, s$.*
- (5) *Every path γ in \mathcal{G}^2 has an iterate $f^t(\gamma)$ which is homotopic rel. endpoints to a strongly legal path. There is an upper bound to the exponent t , which depends only on the number q of factors in any decomposition of $\gamma = \gamma_1 \circ \dots \circ \gamma_q$ as concatenation of strongly legal paths γ_i , and not on the particular choice of γ itself.*

The analogous statement is true for free homotopy classes of loops in \mathcal{G}^2 .

(6) The relative part X is a graph, and the marking map (given for example by the retraction r together with the identification $\pi_1\Gamma' = \mathbb{F}_n$) restricts on each connected component to a monomorphism.

(7) Strongly legal paths in \mathcal{G}^2 , or strongly legal paths with r applied to some of its subpaths, lift to quasi-geodesics in the universal covering $\tilde{\mathcal{G}}^2$.

Proof. (1) This property is actually part of the fact that f_k is strong, see Definition 10.7. However, to be explicit, we note that in Cases 1 and 3 this property follows directly from the induction hypotheses, as the train track part of \mathcal{G}_k^2 is equal to that of \mathcal{G}_{k-1}^2 . In Cases 2 and 4, the expansiveness of f_k is a direct consequences of the particular properties stated in each of those cases. Please note that in Case 4 we need to use the hypothesis from Property 10.16 (3) that the permutation π is transitive.

(2) To show this, we first observe that all auxiliary edges added to X_{k-1} , in the first step of our iterative construction, are permuted by f_k among themselves. Thus, in Cases 1, 2 and 4 we can use directly the inductive hypothesis that f_{k-1} has polynomial growth on its relative part, since in these three cases, up to adding the auxiliary edges, the relative part has not been changed when passing from f_{k-1} to f_k . In Case 3 we observe that the edges added to X_{k-1} are also permuted among themselves, up to adding initial and terminal subpaths to them which are entirely contained in X_{k-1} , so that again the inductive hypothesis about polynomial growth of f_{k-1} suffices to derive the claim.

(3) This follows directly from the definition of our iterative construction of f_k and \mathcal{G}_k^2 .

(4) To see this, we first observe that before stating Proposition 10.17 and considering the 4 cases, we applied the attaching-iteration method through Corollary 10.15 to $f_k : \mathcal{G}_k \rightarrow \mathcal{G}_k$. It is easy to see that any time one applies the attaching-iteration method to f_k , one can simultaneously apply the same operations to the given relative train track map $f : \Gamma \rightarrow \Gamma$ (or, more precisely, to $g_k : \Gamma_k \rightarrow \Gamma_k$), and the resulting map is again a relative train track map. Thus we can use by induction that the claim is true for f_{k-1} after having applied Corollary 10.15, and thus obtain the claim for f_k directly from the definitions in each of the 4 above cases.

(5) By Remark 10.6 (a) and Proposition 10.4 (a) there is an iterate $f_k^t(\gamma)$ that is homotopic to a pseudo-legal path. Since in \mathcal{G}_k^2 Nielsen faces have been expanded at every periodic INP, the pseudo-legal path can be homotoped, at each of its periodic INP's, over the corresponding Nielsen face, to give after finitely many of such alterations a strongly legal path that is homotopic to $f_k^t(\gamma)$.

It follows directly from Proposition 10.4 (b) that the number $t \geq 0$ of iterates of f_k , needed above to make $f_k^t(\gamma)$ pseudo-legal, can be bounded above as function of the number of illegal turns in the originally given path γ .

(6) The fact that X_k is a graph follows directly from the induction hypothesis that X_{k-1} is a graph, since in the above inductive procedure only edges have been added. The (not really essential) fact that the marking map is injective on each component, however, would in general be wrong, unless we actually introduce, instead of auxiliary edges, an auxiliary

vertex and auxiliary half edges, as explained in the Aside 5.2. For more detail see Definition 3.7 of [28].

(7) The retraction $r_k : \mathcal{G}_k^2 \rightarrow \Gamma'_k \subset \mathcal{G}_k^2$ is a deformation retraction and as such homotopic (in \mathcal{G}_k^2) to the identity map of \mathcal{G}_k^2 , and it maps strongly legal paths to reduced paths in $\Gamma'_k \subset \mathcal{G}_k^2$. Since reduced paths in Γ'_k lift to geodesics in the universal covering $\tilde{\Gamma}'_k$ and thus to quasi-geodesics in $\tilde{\mathcal{G}}_k^2$, it follows that strongly legal paths, or strongly legal paths with r_k applied to some of its subpaths, also lift to quasi-geodesics in $\tilde{\mathcal{G}}_k^2$. \square

10.5 The structure of automorphisms of \mathbb{F}_n

Partial train track maps with Nielsen faces as introduced in [28], and hence in particular the β -train track maps considered here, have a crucial advantage over all other train tracks, classical [3] or improved [1] or improved-improved [6], [16], etc: The structure of the train track transition matrix $M(f) = (m_{e,e'})_{e,e' \in \hat{\Gamma}}$ is an invariant of the conjugacy class of the outer automorphism $\hat{\alpha} \in \text{Out}(\mathbb{F}_n)$ defined by α . Here the coefficient $m_{e,e'}$ is given by the number of times that the (legal) path $f(e')$ crosses over e or its inverse \bar{e} . The following result has been shown in [28], §4. For a reader friendly exposition of train tracks, invariant \mathbb{R} -trees, and the precise relationship to the transition matrix and its eigen vectors, see [30].

Theorem 10.19. (a) *For any β -train track representative $f : \mathcal{G}^2 \rightarrow \mathcal{G}^2$ of $\alpha \in \text{Aut}(\mathbb{F}_n)$ there is a canonical bijection between the set of α -invariant \mathbb{R} -trees T as given in Proposition 3.1 (a) and the set of row eigen vectors \vec{v}_* of $M(f)$ with real eigen value $\lambda > 1$.*

(b) *If T is given by the eigenvector \vec{v}_* as above, then every conjugacy class of non-trivial point stabilizers in T , unless it is of polynomial α -growth, is given by a non-trivial $M(f)$ -invariant subspace of $\mathbb{R}^{\hat{\Gamma}}$ on which \vec{v}_* has coefficients of value 0. These invariant subspaces are in 1-1 relationship with those complementary components of the support of \vec{v}_* in \mathcal{G}^1 that are not contained in X . In particular, the induced automorphism on these point stabilizers is represented by a sub-train-track of \mathcal{G}^2 , given by those complementary components, provided with the corresponding restriction of the train track map f .*

The use of this structure theorem is highlighted by the fact that, after replacing f by a suitable power, there are (up to scalar multiples) finitely many eigen vectors of $M(f)$ which have as support a subspace of $\mathbb{R}^{\hat{\Gamma}}$ on which $M(f)$ has an irreducible matrix with irreducible powers. Here “irreducible” refers to the standard use of this terminology in the context of non-negative matrices. The resulting invariant \mathbb{R} -trees, called *partial pseudo-Anosov* trees in [28], are the smallest building blocks out of which the exponentially growing part of α is iteratively built. [However, a word of caution seems to be appropriate here: Even if $M(f)$ consists of a single irreducible block with irreducible powers, and if the relative part of \mathcal{G}^1 is empty, one can not conclude that α is an iwip automorphism. This conclusion is only possible after a further local analysis at the vertices of \mathcal{G}^1 , see [30], §7 and [24], §IV.]

In §3 iteratively constructed invariant trees T_j have been considered, in order to find the characteristic family $\mathcal{H}(\alpha)$ of (conjugacy classes of) subgroups H_i where α has polynomial

growth. All of these T_j are given as in Theorem 10.19 by eigenvectors \vec{v}_* of the transition matrix $M(f)$ or of “submatrices” of $M(f)$ describing the induced β -train track map on an f^t -invariant subgraph of \mathcal{G}^2 . In particular, we obtain:

Proposition 10.20. *The connected components X_i of the relative part X of \mathcal{G}^2 are in 1-1 correspondence with the subgroups H_i from the canonical family $\mathcal{H}(\alpha)$ of polynomial α -growth: one has*

$$H_i = \pi_1 X_i,$$

up to conjugation and permutation of the H_i .