

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

F. Gautero, M. Lustig

July 2, 2007

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 $\mathcal{H}(\alpha)$.

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 [3] that elements $w \in \mathbb{F}_n$ grow either exponentially or 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

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

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 [16]. It has been fundamentally shaped by Farb [11] 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 [26], [9] and [25]. The relevant facts about relative hyperbolicity are recalled in §2 below.

A 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 1.2 (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 [23]). It is shown by Farb [11] 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 1.2.

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 [13]) and Brinkmann [7]. In [14] 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 [13] have been further extended in [15] to give a general combination theorem for relatively hyperbolic groups. 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 in [22], 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 1.3. *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 [18]. 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 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 [10]), 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 [16] in the absolute case (i.e. all H_i are trivial) and generalized subsequently in [15].

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

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

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 are (up to conjugation) 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 [15]. For a reproof using somewhat different methods compare also [24].

Theorem 2.9. *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 [12]). For completeness we state in full generality, although some ingredients (for example “very small” actions) are not specifically used here. The paper [21] 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 $\lambda > 1$, i.e.*

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

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 $H_k = \text{Stab}(x_i)$, 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 find the canonical subgroups of 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 unique, 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 (up to linear combinations finite) set of all such finite partially ordered collections of iteratively defined invariant \mathbb{R} -trees, which is an important invariant of the conjugacy class of α in $Out(\mathbb{F}_n)$, is obtained through the structural analysis of α derived in §4 of [19], summarized below as Theorem 4.3. The latter gives also an alternative proof for the uniqueness of the family \mathcal{H} .

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 in [19]. 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 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. Since X is a graph, we can (and will tacitly from now on) always assume that such χ is a reduced path in 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], in [19] (Theorem 4.1 and Proposition 3.12) it has been shown 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 a further property is needed, which in [19], [20] only occurs for the “top stratum” of $\widehat{\Gamma}$, namely that legal paths lift to quasi-geodesics in the universal covering $\widetilde{\mathcal{G}}^2$.

In a recent paper [22] the original construction of [19], [20] has been improved to define *β -train track maps*, which have this additional property, by introducing the notion of *strongly legal* paths, which are in particular legal in the sense defined above. This improvement, and some other technical properties of β -train tracks are presented in detail in the next section.

The following is proved in detail in [22] (compare also [19]). Note that all properties of these “better” β -train tracks¹ 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:*

- 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.*
- 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].*
- Every edge e of the train track part of \mathcal{G}^2 is strongly legal (see Definition 5.4) and thus in particular legal.*
- Every strongly legal path in \mathcal{G}^1 is mapped by f to a strongly legal path.*
- 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 $\widehat{\Gamma}$.*

¹The second author would like to thank the first author for having coined this term as translation of the word “mieux” into français, at the occasion of his habilitation defense.

(f) The lift of any strongly legal path γ to the universal covering $\tilde{\mathcal{G}}^2$ is a quasi-geodesic with respect to the simplicial metric on $\tilde{\mathcal{G}}^2$ (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 $\tilde{\mathcal{G}}^2$. Every path that is mapped by the retraction r to a reduced path in Γ lifts also to a quasi-geodesic in $\tilde{\mathcal{G}}^2$. 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 $\tilde{\mathcal{G}}^2$.

(h) For every path γ in \mathcal{G}^1 there is an integer $\hat{t} = \hat{t}(\gamma) \geq 0$ such that $f^{\hat{t}}(\gamma)$ is homotopic rel. endpoints in \mathcal{G}^2 to a strongly legal path in \mathcal{G}^1 . The integer $\hat{t}(\gamma)$ depends only on the number of illegal turns (compare Definition 5.4) in γ and not on γ itself.

We use the remainder of this section to draw the connection to Theorem 3.1 and the subsequent discussion in §3.

Partial train track maps with Nielsen faces, and hence in particular β -train tracks, have another crucial advantage over all other train tracks, classical [3] or improved [1] or improved-improved [6], 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 [19], §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 [21].

Theorem 4.3. (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 [19], are the smallest building blocks out of which the exponentially growing part of α is iteratively built. Compare the discussion in the previous section, before Proposition 3.2.

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 [21], §7 and [17], §IV.

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 [19]), 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 [19] or [22].

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 \bar{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 $\tilde{\mathcal{G}}^2$ to quasi-geodesics, with respect to the absolute metric on $\tilde{\mathcal{G}}^2$.*

Proof. Statement (a) follows directly from the above definition of a normalized path. For statement (b) one uses the technical finesse employed in the introduction of the auxiliary edges in [19], compare 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 seize 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 $ILT(\cdot)$ 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*

$$ILT(\gamma') \leq ILT(\gamma)$$

(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 $|\gamma_0|_{rel} \geq 2 |\gamma_0|_{rel}$.

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

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 $\widehat{\gamma}$.

Since the very same arguments extend to all subpaths of $\widehat{\gamma}$, it follows directly that $\widehat{\gamma}$ 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 $\widetilde{\mathcal{G}}^2$, 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 $\widetilde{\mathcal{G}}^2$ 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 $\widehat{\gamma}'$ with endpoints $\widehat{x}', \widehat{y}'$, such that $d_{rel}(x', \widehat{x}') \leq C_1$ and $d_{rel}(y', \widehat{y}') \leq C_1$, and

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

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 $\widehat{\gamma}$ 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 $\widehat{\gamma}'$ of $\widehat{\gamma}$ 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 $\widehat{\gamma}$ is contained in the 1-skeleton of $\widetilde{\mathcal{G}}^2$, and that furthermore $\widehat{\gamma}'$ is an edge path, i.e. starts and ends at a vertex of $\widetilde{\mathcal{G}}^2$.

We now consider the set $\widehat{\mathcal{L}}$ of maximal subpaths $\widehat{\gamma}_i$ of $\widehat{\gamma}'$ which are contained in the relative part. The collection of closed subpaths $\widehat{\gamma}_j$ of $\widehat{\gamma}'$ complementary to those in $\widehat{\mathcal{L}}$ is denoted by $\widehat{\mathcal{L}}^c$. We observe that, by Lemma 7.2, every such $\widehat{\gamma}_j$ is an absolute quasi-geodesic, with quasi-geodesy constants depending only on C_1 and not on our choice of $\widehat{\gamma}'$. 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 1.3 as stated in the Introduction. The notion of a relative hyperbolic automorphism is recalled in Definition 2.6:

Proof of Proposition 1.3. 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 \widehat{X}_i of the full preimage \widehat{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 \widetilde{f} induces canonically a map $\widehat{f} : \widehat{\mathcal{G}}^2 \rightarrow \widehat{\mathcal{G}}^2$ that satisfies similarly, for any $w \in \mathbb{F}_n$ and any point $P \in \widehat{\mathcal{G}}^2$:

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

For our purposes below we also want, in addition to this “twisted commutativity property”, that \widehat{f} fixes a vertex of $\widehat{\mathcal{G}}^2$ outside of the union \widehat{X} of all \widehat{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 $\widehat{\Gamma}$ (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 \widehat{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 $v\widehat{f}^k$. It follows that $v\widehat{f}^k$ “twistedly commutes” with $\iota_v \alpha^k$ in the above meaning, where ι_v denotes the inner automorphisms $\iota_v : \mathbb{F}_n \rightarrow \mathbb{F}_n, w \mapsto vwv^{-1}$.

By virtue of Remark 2.7 we can continue to work with $v\widehat{f}^k$ and $\iota_v \alpha^k$ rather than with \widehat{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 \widehat{f} and α , but we assume that \widehat{f} has a fixed vertex $Q = \widehat{f}(Q) \in \widehat{\mathcal{G}}^2 \setminus \widehat{X}$.

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

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

by sending the base point $V(1)$ to the above \widehat{f} -fixed vertex $Q \in \widehat{\mathcal{G}}^2 \setminus \widehat{X}$. Every cone vertex of $\Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$ is mapped to the corresponding contracted connected component \widehat{X}_i of \widehat{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 \widehat{X}_i of \widehat{X} . Every edge e of $\Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$ is sent to an edge path $\psi(e)$ in $\widehat{\mathcal{G}}^2$ 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 $\widehat{\mathcal{G}}^2$.

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 [14]). 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 $P, R \in \Gamma_S^{\mathcal{H}(\alpha)}(\mathbb{F}_n)$ one has:

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

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 $\widehat{\mathcal{G}}^2$, 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:

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

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 \tilde{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 1.3 we know that α is hyperbolic relative to $\mathcal{H}(\alpha)$. Hence Theorem 2.9 implies directly the claim. □

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] C. Drutu. Relatively hyperbolic groups: geometry and quasi-isometric invariance. ArXiv math.GR/0605211
- [11] B. Farb. Relatively hyperbolic groups. *Geom. Funct. Anal.* **8**, 1–31 (1998).
- [12] 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).
- [13] F. Gautero. Hyperbolicity of mapping-torus groups and spaces. *L’Enseignement Mathématique* **49**, 263–305 (2003).
- [14] F. Gautero and M. Lustig. Relative hyperbolization of (one-ended hyperbolic)-by-cyclic groups. *Math. Proc. Camb. Phil. Soc.* **137**, 595–611 (2004).
- [15] F. Gautero. Geodesics in trees of hyperbolic and relatively hyperbolic groups. Preprint 2005 (revised 2007).
- [16] M. Gromov. Hyperbolic groups. In “Essays in Group Theory”. *Math. Sci. Res. Inst. Publ.* **8**, 75–263 (1987).
- [17] A. Jäger and M. Lustig. Free groups automorphisms with many fixed points at infinity. Preprint 1998 (revised 2007).
- [18] G. Levitt and M. Lustig. manuscript.
- [19] 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>)

- [20] 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>)
- [21] M. Lustig. Conjugacy and centralizers for iwip automorphisms of free groups. Preprint, 2006. (to appear in the proceedings of the conferences on “*Asymptotic and geometric methods in group theory*”, Geneva and Barcelona 2005).
- [22] M. Lustig. Better train tracks. Preliminary preprint, 2006.
- [23] N. Macura. Quadratic isoperimetric inequality for mapping tori of polynomially growing automorphisms of free groups. *Geom. Funct. Anal.* **10**, 874–901 (2000).
- [24] Mj. Mahan and L. Reeves. A combination theorem for strong relative hyperbolicity. ArXiv math.GR/0611601.
- [25] D. V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties and algorithmic problems. *Memoirs of the American Mathematical Society* **179**, 2006.
- [26] A. Szczepański. Relatively hyperbolic groups. *Michigan Math. Journal* **45**, 611–618 (1998).