

STRETCHING FACTORS, METRICS AND TRAIN TRACKS FOR FREE PRODUCTS

STEFANO FRANCAVIGLIA AND ARMANDO MARTINO

ABSTRACT. In this paper, we develop the metric theory for the outer space of a free product of groups. This generalizes the theory of the outer space of a free group, and includes its relative versions. The outer space of a free product is made of G -trees with possibly non-trivial vertex stabilisers. The strategies are the same as in the classical case, with some technicalities arising from the presence of infinite-valence vertices.

We describe the Lipschitz metric and show how to compute it; we prove the existence of optimal maps; we describe geodesics represented by folding paths.

We show that train tracks representative of irreducible (hence hyperbolic) automorphisms exist and that their are metrically characterized as minimal displaced points, showing in particular that the set of train tracks is closed (in particular, answering to some questions raised in *Axis in outer space* (2011) concerning the axis bundle of irreducible automorphisms).

Finally, we include a proof of the existence of simplicial train tracks map without using Perron-Frobenius theory.

A direct corollary of this general viewpoint is an easy proof that relative train track maps exist in both the free group and free product case.

1. Introduction

In this paper, we are interested in studying the outer space of a free product of groups. Namely, given a group of the form $G = G_1 * \cdots * G_p * F_k$, we study the set of trees where G acts with vertex stabilizers the G_i 's. In the case $G = G_1 * \cdots * G_p * F_k$ is the free product decomposition of a finitely generated

Received November 11, 2015; received in final form August 25, 2016.
2010 *Mathematics Subject Classification*. 20E06, 20E36, 20E08.

group G , then this was introduced by Guirardel and Levitt in [36] in the case where this is the Grushko decomposition of G . That is, when each G_i is freely indecomposable and not isomorphic to \mathbb{Z} . However, we shall consider such spaces with respect to an arbitrary free product decomposition, and not necessarily the natural Grushko one. Similar spaces are studied by Sykiotis in [34].

The theory is similar to that of the case of free groups, with the advantage that this unified viewpoint covers at once both the general case of a free product as well as many “relative” cases of the classical Outer space. The group of isomorphisms that acts on G will be that of automorphisms that preserve the set of conjugacy classes of the G_i (which coincides with $\text{Aut}(G)$ in the case of the Grushko decomposition, by the Kurosh subgroup theorem).

In particular, one can define the Lipschitz metric (see [1], [20], [21] for the classical case). The presence of vertex stabilizers involves some technical complications (for instance, the Ascoli–Arzelà theorem does not hold for spaces that are not locally compact) but the main results of the classical case hold *mutatis mutandis*. For instance, optimal maps exist and Lipschitz factors can be computed on a list of simple candidates. Also, geodesics are constructed via folding paths.

For the study of automorphisms a very useful tool is the theory of train track maps, developed by Bestvina and Handel [8] (see also [4], [5], [6], [7]) and extensively used in literature. This tool is available also in the present setting.

For studying train tracks, we chose to follow the metric viewpoint as in [3]. In particular, we show that for an irreducible automorphism the set of train tracks coincides with the set of minimally displaced elements. We remark that there is no uniform definition of train track maps in the literature, even if the difference from one definition to another is minimal. As the set of minimally displaced elements is closed, this gives in particular a proof that the set of train tracks is closed, hence answering to a question raised in [23], where the authors give a characterizations of the axis bundle of an irreducible automorphism (see Remark 8.22). We would also like to mention the very recent preprint [33] about axis bundles.

Many of the results about train tracks that we are going to describe are well known (at least to the experts) in the case of free groups, and the proofs in our general setting do not require substantial changes. We give here explicit and fully detailed proofs of all these facts for completeness. We refer the reader also to the very recent and nice preprint [32] that deals with local finite trees with possibly non trivial edge-stabilizers, from the same viewpoint of us. As Meinert pointed out, the fact the we work with trivial edge-stabilizers is crucial, as Lemma 4.2 may fail in general. In this work, we do not develop the theories of geodesic currents and laminations ([10], [11], [12], [13], [14], [15],

[16], [19], [22], [24], [25], [26], [27], [28], [29], [30], [31]) that would certainly be of interest in this general case.

2. G -Trees and lengths

For any simplicial tree T (not necessarily locally compact), we denote by VT and ET the set of vertices and edges of T , respectively. A *simplicial metric tree*, is a simplicial tree equipped with a complete path metric such that edges are isometric to closed intervals of \mathbb{R} . Note that the simplicial structure on a metric tree is an additional structure that is not necessarily determined by the metric structure. However, we do require that all branch points be vertices, and generally we will simply take the set of vertices to be the set of branch points (which is determined by the metric structure).

For $x, y \in T$, we denote by $[x, y]_T$ (or simply by $[x, y]$ if there is no ambiguity concerning T) the unique path from x to y , and for a path γ in T we denote by $l_T(\gamma)$ the length of γ in T .

Let G be a group. In this work, by a G -tree we mean a simplicial metric tree $T = (T, d_T)$, where G acts simplicially on T and for all $g \in G$ and $e \in ET$, e and ge are isometric. In other words, G is acting on T by isometries and preserving the simplicial structure.

If T is a G -tree, then the quotient space $G \backslash T$ is a graph. We denote by $\pi_T : T \rightarrow G \backslash T$ the projection map.

In general, a path γ in $G \backslash T$ may have many lifts to T , even if we fix the initial point of the lift. This is because each time γ passes through an edge whose initial vertex has a lift with non-trivial stabilizer, we have many choices for the lift of the edge.

Let T be a G -tree. The following definitions depend on the action of G on T . An element $g \in G$ is called *hyperbolic* if it fixes no points. Any hyperbolic element g of G acts by translation on a subtree of T homeomorphic to the real line, called the *axis* of g and denoted by $axis_T(g)$. The *translation length* of g is the distance that g translates the axis. The action of G on T defines a *length function* denoted by l_T

$$l_T : G \rightarrow \mathbb{R}, \quad l_T(g) := \inf_{x \in T} d_T(x, gx).$$

REMARK 2.1. We note that, under our hypothesis, this inf is always achieved (see for example [17, 1.3]). In particular, $g \in G$ is hyperbolic if and only if $l_T(g) > 0$.

3. The outer space of a free product

We follow [36]. We will consider groups G of the form

$$G = G_1 * \cdots * G_p * F_k,$$

where $(G_i)_{i=1}^p$ is a family of groups, and F_k denotes the free group of rank $0 \leq k < \infty$.

We will be mainly concerned with the case where G admits a co-compact action on a tree with trivial edge stabilisers and indecomposable vertex stabilisers, or equivalently a group of finite Kurosh rank. However, in general we will not assume that that G_i are indecomposable. That is, while G may admit a decomposition as a free product of finitely many freely indecomposable groups, we are interested in developing the subsequent theory in the situation where our given free product decomposition is not necessarily of that kind. For instance, we will apply the theory in the case that G is free, and the G_i are certain free factors of G .

Let $\mathcal{T}(G)$ denote the set of simplicial metric G -trees. We say that two elements $T, T' \in \mathcal{T}(G)$ are *equivalent*, and we write $T \sim T'$, if there exists a G -equivariant isometry $f: T \rightarrow T'$.

Let $T \in \mathcal{T}(G)$. A vertex $v \in VT$ is *redundant*, if it has degree two, and any g that fixes v also fixes the edges adjacent to v . It is *terminal* if $T - \{v\}$ is connected. A terminal vertex is often referred to as a *leaf*. We will consider G -trees with no *redundant* vertices.

Let $\mathcal{O} = \mathcal{O}(G, (G_i)_{i=1}^p, F_k)$ be the subset of $\mathcal{T}(G)/\sim$ of simplicial, metric G -trees T , up to equivariant isometry, satisfying that:

- (C0) T has no redundant vertices;
- (C1) the G -action of T is *minimal* (i.e., there exist no proper invariant subtree), with trivial edge stabilizers;
- (C2) for each $i = 1, \dots, p$, there is exactly one orbit of vertices with stabilizer conjugate to G_i and all edge stabilizers are trivial;
- (C3) all other vertices have trivial stabilizer. We will often refer to such vertices as *free* vertices.

It may be worth to mention that under such assumptions, for any $T \in \mathcal{O}$ the quotient $G \backslash T$ is a finite graph.

The space \mathcal{O} admit a natural action of $(0, \infty)$ defined by rescaling the metric, that is to say, multiplying all lengths of the edges by the same number. The quotient space of \mathcal{O} by that action is denoted by $\mathcal{PO} = \mathcal{PO}(G, (G_i)_{i=1}^p, F_k)$ and is called the *outer space* of G . Sometimes \mathcal{O} will be referred to as the *unprojectivized outer space* of G .

There is a natural map from $\mathcal{T}(G)$ to \mathbb{R}^G , mapping T to $(l_T(g))_{g \in G}$. This map clearly factors through $\mathcal{T}(G)/\sim$. The following fact is proved in [17, Theorem 3.7] (see also [2]).

LEMMA 3.1. *The restriction of the translation length function to $\mathcal{O} \rightarrow \mathbb{R}^G$ is injective.*

The *axes topology* on \mathcal{O} is the topology induced as a subspace of \mathbb{R}^G .

As in [36], there are in fact two topologies on \mathcal{O} . There is, as in Outer Space, the simplicial topology which is different from the Gromov topology (which coincides with the axes topology). The metric we study in the following discussion induces the same topology as the axes topology.

Definition 3.2. The group $\text{Aut}(G, \mathcal{O})$ is the group of automorphisms that preserve the set of conjugacy classes of the G_i 's. Namely $\phi \in \text{Aut}(G)$ belongs to $\text{Aut}(G, \mathcal{O})$ if $\phi(G_i)$ is conjugate to one of the G_i 's.

In the case of the Grushko decomposition, $\text{Aut}(G) = \text{Aut}(G, \mathcal{O})$. The group $\text{Aut}(G, \mathcal{O})$ acts on $\mathcal{T}(G)$ by changing the action. That is, for $\phi \in \text{Aut}(G)$ and T in $\mathcal{T}(G)$, the image of T under ϕ is the G -tree with the same underlying tree as T , endowed with the action given by $(g, x) \in G \times T \mapsto \phi(g)x \in T$. If ϕ_h is the automorphism of G given by conjugation by $h \in G$ ($g \mapsto h^{-1}gh$), then for every $T \in \mathcal{T}(G)$, $T \sim \phi_h(T)$ via the map $T \rightarrow \phi_h(T)$, $x \mapsto h^{-1}x$. Thus, $\text{Out}(G, \mathcal{O}) = \text{Aut}(G, \mathcal{O})/\text{Inn}(G)$ acts on $\mathcal{T}(G)/\sim$.

4. The metric

4.1. \mathcal{O} -Maps. Let T be a simplicial metric G -tree. Denote by $\text{Hyp}(T)$ set of elements $g \in G$ whose the action on T is hyperbolic (see [17] for details). If $[T] \in \mathcal{O}$ and $g \notin \text{Hyp}(T)$, then g fixes a vertex of T , and by (C2) there exists $i \in \{1, \dots, p\}$ such that g lies in a G -conjugate of G_i . Conversely, if g lies in a G -conjugate of some G_i , $i \in \{1, \dots, p\}$, by (C2) g fixes a vertex, and then it is not hyperbolic. Therefore, $g \in \text{Hyp}(T)$ if and only if it is hyperbolic for any other element of \mathcal{O} . The set of hyperbolic elements of G for some (and hence for all) T in \mathcal{O} is denoted by $\text{Hyp}(\mathcal{O})$.

Definition 4.1 (\mathcal{O} -maps). Let $A, B \in \mathcal{O}$. An \mathcal{O} -map $f: A \rightarrow B$ is a G -equivariant, Lipschitz continuous, surjective function. (Note that we don't require to f to be a graph morphism.) We denote by $\text{Lip}(f)$ the Lipschitz constant of f , that is the smallest constant $K \geq 0$ such that, for all $x_1, x_2 \in A$

$$d_B(f(x_1), f(x_2)) \leq K d_A(x_1, x_2).$$

LEMMA 4.2. For every pair $A, B \in \mathcal{O}$, there exists a \mathcal{O} -map $f: A \rightarrow B$. Moreover, any two \mathcal{O} -maps from A to B coincide on the non-free vertices.

See [18] or [35, Theorem 3.8] for a proof.

LEMMA 4.3. Let $A, B \in \mathcal{O}$. For any \mathcal{O} -map $f: A \rightarrow B$, we have

$$\sup_{g \in \text{Hyp}(\mathcal{O})} \frac{l_B(g)}{l_A(g)} \leq \text{Lip}(f).$$

Proof. Let $f: A \rightarrow B$ be an \mathcal{O} -map, $g \in G$, and $x \in A$. Since f is continuous $[f(x), f(gx)]_B \subseteq f([x, gx]_A)$, then

$$\begin{aligned} l_B(g) &\leq d_B(f(x), gf(x)) \\ &= d_B(f(x), f(gx)) \\ &\leq l_B(f([x, gx]_A)) \\ &\leq \text{Lip}(f)l_A([x, gx]_A). \end{aligned}$$

By Remark 2.1, there exists $x_g \in A$ realizing $l_A(g)$, that is $l_A([x_g, gx_g]_A) = l_A(g)$. Using x_g in the previous inequality, we conclude that

$$(1) \quad l_B(g) \leq \text{Lip}(f)l_A(g). \quad \square$$

4.2. The metrics.

Definitions 4.4. For any pair $A, B \in \mathcal{O}$ we define the right and left maximal stretching factors

$$\Lambda_R(A, B) := \sup_{g \in \text{Hyp}(\mathcal{O})} \frac{l_B(g)}{l_A(g)}, \quad \Lambda_L(A, B) := \sup_{g \in \text{Hyp}(\mathcal{O})} \frac{l_A(g)}{l_B(g)} = \Lambda_R(B, A)$$

and *asymmetric pseudo-distances*

$$d_R(A, B) := \log \Lambda_R(A, B), \quad d_L(A, B) := \log \Lambda_L(A, B) = d_R(B, A).$$

We define $\Lambda(A, B) := \Lambda_R(A, B)\Lambda_L(A, B)$ and the *distance* between A and B as

$$d(A, B) := \log \Lambda(A, B).$$

Directed triangular inequalities are readily checked for d_R and d_L , thus triangular inequality holds for d . Moreover, d is a genuine distance on \mathcal{O} as $d(A, B) = 0$ gives $l_B(g) = l_A(g)$ for any element of G , and this implies that $A = B$ by Lemma 3.1. The functions d_R and d_L become asymmetric distances once restricted to the subset of \mathcal{O} of G -trees with co-volume one, which can be identified with \mathcal{PO} . (See, for example, [20], [21], [1] for the study of such functions in the case of outer space of free groups.) We notice that without volume-restriction d_R and d_L can be negative. Also, the fact that $d_R(X, Y) = 0$ implies $X = Y$ is true only with the volume-restriction (see [20, Lemma 4.16] for a proof in the classical case that extends with no changes to the actual case).

LEMMA 4.5. *The action of $\text{Aut}(G, \mathcal{O})$ on \mathcal{O} is by isometries.*

Proof. If $\phi \in \text{Aut}(G, \mathcal{O})$, then it preserves the conjugacy classes of the G_i 's. Therefore, g is hyperbolic if and only if $\phi(g)$ is. Thus,

$$\sup_{g \in \text{Hyp}(\mathcal{O})} \frac{l_B(g)}{l_A(g)} = \sup_{g \in \text{Hyp}(\mathcal{O})} \frac{l_B(\phi(g))}{l_A(\phi(g))} = \sup_{g \in \text{Hyp}(\mathcal{O})} \frac{l_\phi(B)(g)}{l_\phi(A)(g)}. \quad \square$$

5. Equivariant Ascoli–Arzelá

In this section, we provide a tool for computing stretching factors. We follow the approach of [20]. The main issue is that given $A, B \in \mathcal{O}$, one needs to find a map between them which optimizes the Lipschitz constant. Since elements of \mathcal{O} are not locally compact, Ascoli–Arzelá does not apply directly, and we need to control local pathologies by hand.

The lazy reader may skip this section by paying the small price of missing out on some definitions and the beautiful proof of the equivariant version of Ascoli–Arzelá theorem.

Definition 5.1. A map $f : A \rightarrow B$ between metric graphs is called *piecewise linear* if it is continuous and for all edges e of A , there exists a positive number $S_{f,e}$, called the stretching factor of f at e , such that the restriction of f to e has constant speed $S_{f,e}$. More precisely, f is piecewise linear if for any $e \in EA$, the following diagram commutes and the vertical functions are local isometries:

$$\begin{array}{ccc}
 e \subset A & \xrightarrow{f|_e} & f(e) \subset B \\
 \downarrow & & \uparrow \\
 \mathbb{R} \supset [0, l_A(e)] & \xrightarrow{t \mapsto S_{f,e}t} & [0, S_{f,e}l_A(e)] \subset \mathbb{R}
 \end{array}$$

We remark that piecewise linear maps are locally injective on edges.

Definition 5.2 (PL-map for trees). Let $A, B \in \mathcal{O}$. We say that a function $f : A \rightarrow B$ is a *PL-map* if it is a piecewise linear \mathcal{O} -map. For any \mathcal{O} -map $f : A \rightarrow B$ we define the map $PL(f)$ as the unique *PL-map* that coincides with f on vertices.

REMARK 5.3. Let $f : A \rightarrow B$ be an \mathcal{O} -map and $e \in EA$. If $l_B(f(e))$ denotes the distance between the images of the vertices of e , then by construction we have $S_{PL(f),e} = l_B(f(e))/l_A(e) \leq \text{Lip}(f|_e) \leq \text{Lip}(f)$, for all $e \in EA$. Therefore, $\text{Lip}(PL(f)) \leq \text{Lip}(f)$.

Before proving the equivariant Ascoli–Arzelá, we discuss an example.

Example 5.4. Consider a segment $[0, 3]$ with free vertices and a segment $[0, 1]$ with a vertex non free, say 0, with associated group \mathbb{Z} . Consider the associated trees $A = [0, 3]$ and B . B is a star-shaped tree with an infinite valence vertex, say 0, from which emanate infinite copies of $[0, 1]_n$ labeled by $n \in \mathbb{Z}$. Now consider the map $f : [0, 3] \rightarrow [0, 1]$

$$f(t) = \begin{cases} t & t \in [0, 1], \\ 2 - t & t \in [1, 2], \\ t - 2 & t \in [2, 3]. \end{cases}$$

For any $n \in \mathbb{Z}$ there exist a lift $f_n : A \rightarrow B$ of f such that the segment $[0, 2]$ is mapped to $[0, 1]_n$ and $[2, 3]$ is mapped to $[0, 1]_0$. The sequence f_n has no sub-sequence that converges, but clearly if one “straightens” f_n by collapsing $[0, 2]$ to 0, this sequence becomes constant. Of course, this is safe because there is no G -action on A .

This is more or less everything that can go wrong. We now introduce the precise notion of collapsible and non-collapsible map.

Definition 5.5. Let $A \in \mathcal{O}$. A subset $X \subset A$ is *collapsible* if $gX \cap X = \emptyset$ for any $\text{Id} \neq g \in G$.

Definition 5.6. Let $A, B \in \mathcal{O}$ and $f : A \rightarrow B$ be an \mathcal{O} -map. A *collapsible component* of f is a connected component of $A \setminus f^{-1}(v)$, for a $v \in VB$ non-free, which is collapsible.

Definition 5.7. Let $A, B \in \mathcal{O}$ and $f : A \rightarrow B$ be an \mathcal{O} -map. f is said *collapsible* if it has a collapsible component. f is said *non-collapsible* if it is not collapsible.

Note that f is non-collapsible if any component C of $A \setminus f^{-1}(v)$ either contains a non-free vertex or there is point $w \in A$ and $\text{id} \neq g \in G$ so that both w and gw belong to C .

Definition 5.8. Let $A, B \in \mathcal{O}$ and $f : A \rightarrow B$ be an \mathcal{O} -map. f is σ -PL if:

- σ is a simplicial structure $(V\sigma, E\sigma)$ on A obtained by adding 2-valent vertices to A .
- The number of G -orbits of edges of σ is finite.
- For any $v \in VB$ non-free, $f^{-1}(v)$ is a forest (union of trees) with leaves in $V\sigma$.
- f is PL w.r.t. σ .

Any PL-map $f : A \rightarrow B$ is σ -PL for the f -pull-back structure induced on A by B .

LEMMA 5.9. *Let $A, B \in \mathcal{O}$ and $f : A \rightarrow B$ be a σ -PL map. Then the number of orbits of collapsible components of f is finite.*

Proof. First note that $A \setminus f^{-1}(gv) = g(A \setminus f^{-1}(v))$. Hence, the orbits of components corresponding to the orbit of v have representatives in $A \setminus f^{-1}(v)$. Since there are finitely many orbits of vertices, it is enough to show that the collapsible components in $A \setminus f^{-1}(v)$ are contained in finitely many orbits.

We argue by contradiction and assume that we have infinitely many collapsible components C_i of $A \setminus f^{-1}(v)$ in distinct orbits.

Since there are finitely many orbits of (open) edges of σ , we may assume that the orbit of some edge e meets every C_i ; hence there are $g_i \in G$ such that $e \subset g_i C_i$. Moreover, for the same reason and from the definition of collapsible component, we deduce that there is a uniform bound on the number of edges in any collapsible component.

Without loss of generality, we may assume that the number of edges of C_0 is maximal amongst the C_i and that $g_0 = \text{Id}$. Since C_i and C_0 are not in the same orbit, $g_i C_i \neq C_0$. On the other hand $g_i C_i \cap C_0 \neq \emptyset$, thus one of them contains a leaf of the other. Since C_0 is maximal there is a leaf x_i of $g_i C_i$

in C_0 . Leaves of C_i are σ -vertices, and since C_0 has finitely many vertices and edges, we may assume that $x_i = x$ is independent of i , and that there is an edge ξ of C_0 contained in $g_i C_i$ for all i (we note that $g_i C_i \cap C_0$ contains at least one edge because it is the intersection of open sets, guaranteeing the existence of ξ).

As x is a leaf of $g_i C_i$, $f(x) = g_i v$ for all i . In particular,

$$g_i^{-1} g_j(v) = v.$$

Since $\xi \subset g_i C_i \cap g_j C_j$ we have $C_i \cap g_i^{-1} g_j C_j \neq \emptyset$. However, C_i is a component of $A \setminus f^{-1}(v)$ and $g_i^{-1} g_j C_j$ is a component of $A \setminus f^{-1}(g_i^{-1} g_j v) = A \setminus f^{-1}(v)$. Hence, they are equal contradicting the fact that the C_i 's are in distinct orbits. \square

LEMMA 5.10. *Let $A, B \in \mathcal{O}$. Let $f : A \rightarrow B$ be a collapsible σ -PL map. Then there is an \mathcal{O} -map $f_\bullet : A \rightarrow B$ such that:*

- f_\bullet is σ -PL (same σ).
- $\text{Lip}(f_\bullet) \leq \text{Lip}(f)$.
- The number of orbits of collapsible components of f_\bullet is strictly smaller than that of f .

Proof. Let $v \in VB$ non-free and let C be a collapsible component of $A \setminus f^{-1}(v)$. Collapse C by defining $f_\bullet|_C = v$. Extend f_\bullet by equivariance on the orbit of C . This is possible since $gC \cap C = \emptyset$ for $g \neq id$. On the remaining part of A , let $f_\bullet = f$. Clearly, f_\bullet is an \mathcal{O} -map which is σ -PL and satisfies $\text{Lip}(f_\bullet) \leq \text{Lip}(f)$.

Since $g(A \setminus f^{-1}(v)) = A \setminus f^{-1}(gv)$, it follows that $A \setminus f^{-1}(v)$ contains a representative for every orbit of components. In passing from f to f_\bullet , the components of $A \setminus f^{-1}(v)$ which are not of the form gC are unchanged, while the orbit of C is removed. More precisely, $A \setminus f_\bullet^{-1}(v) = \{A \setminus f^{-1}(v)\} \setminus GC$. Thus, the number of orbits of collapsible components in $A \setminus f^{-1}(v)$ is decreased by 1.

Now, consider the non-free vertices of B that are not in the orbit of v and chose orbit-representatives w_1, \dots, w_k . Define G -sets

$$\begin{aligned} U_i &= \{D : D \text{ is a component of } A \setminus f^{-1}(gw_i) \text{ for some } g \in G\}, \\ \widehat{U}_i &= \{D : D \text{ is a component of } A \setminus f_\bullet^{-1}(gw_i) \text{ for some } g \in G\}, \\ V_i &= \{D : D \text{ is a collapsible component of } A \setminus f^{-1}(gw_i) \text{ for some } g \in G\}, \\ \widehat{V}_i &= \{D : D \text{ is a collapsible component of } A \setminus f_\bullet^{-1}(gw_i) \text{ for some } g \in G\}. \end{aligned}$$

Since $w_i \notin Gv$, then $f_\bullet^{-1}(w_i) \subseteq f^{-1}(w_i)$. Therefore, any component K of $A \setminus f^{-1}(w_i)$ is contained in a unique component K_\bullet of $A \setminus f_\bullet^{-1}(w_i)$. Moreover, if K_\bullet is collapsible, so is K .

This inclusion defines a G -equivariant surjection $\iota : U_i \rightarrow \widehat{U}_i$ such that $\widehat{V}_i \subset \iota(V_i)$. Therefore, the number of G -orbits in V_i is greater than or equal to the number of G -orbits in \widehat{V}_i . □

COROLLARY 5.11 (Existence of Coll). *Let $A, B \in \mathcal{O}$. Let $f : A \rightarrow B$ be a σ -PL map. Then there is an \mathcal{O} -map $\text{Coll}(f) : A \rightarrow B$ such that:*

- $\text{Coll}(f)$ is σ -PL (same σ).
- $\text{Lip}(\text{Coll}(f)) \leq \text{Lip}(f)$.
- $\text{Coll}(f)$ is non-collapsible.

Proof. This follows by induction from Lemmas 5.9 and 5.10. □

In the sequel we use the following conventions:

- When we write $\text{Coll}(f)$ we mean any map given by Corollary 5.11.
- We say that P is true eventually on n if $\exists n_0$ so that P is true for all $n > n_0$, and we write P is true $\forall n \gg 0$.
- P is true frequently if $\forall n \exists m > n$ so that P is true for m .
- A sequence *sub-converges* if it converges up to passing to sub-sequences.

Now we are in position to prove the existence of a map that minimizes the Lipschitz factor.

THEOREM 5.12 (Equivariant Ascoli–Arzelá). *Let $A, B \in \mathcal{O}$, then there exists a PL-map $F : A \rightarrow B$ with*

$$\text{Lip}(F) = \inf \{ \text{Lip}(\varphi) : \varphi \text{ is an } \mathcal{O}\text{-map from } A \text{ to } B \}.$$

Proof. For the entire proof—which requires several lemmas—we fix a minimizing sequence $f'_n : A \rightarrow B$ of PL-maps so that

$$\lim_{n \rightarrow \infty} \text{Lip}(f'_n) = \inf \{ \text{Lip}(\varphi) : \varphi \text{ is an } \mathcal{O}\text{-map from } A \text{ to } B \}$$

and we define

$$f_n = \text{Coll}(f'_n).$$

By definition of Coll we have that the f_n are non-collapsible, uniformly L -Lipschitz and

$$\lim_{n \rightarrow \infty} \text{Lip}(f_n) = \inf \{ \text{Lip}(\varphi) : \varphi \text{ is an } \mathcal{O}\text{-map from } A \text{ to } B \}.$$

By Ascoli–Arzelá the maps $\pi_B \circ f_n \circ \pi_A^{-1} : G \backslash A \rightarrow G \backslash B$ sub-converge to a map \bar{f}_∞ . We will show that \bar{f}_∞ is in fact the projection of a map $A \rightarrow B$ which is the limit of f_n . From now on, we restrict to a sub-sequence and we suppose that $\pi_B \circ f_n \circ \pi_A^{-1}$ uniformly converges to \bar{f}_∞ .

Let \mathcal{T} be the set of pairs (T, f) such that:

- $T \subset A$ is a G -invariant subset of A (not necessarily simplicial).
- $f : T \rightarrow B$ is G -equivariant and L -Lipschitz.
- $\pi_B(f(t)) = \bar{f}_\infty(\pi_A(t))$ for any $t \in T$.
- $f_n|_T$ sub-converges to f .

The set \mathcal{T} is ordered by inclusion/consistency: $(T, f) < (Q, u)$ if $T \subset Q$ and $f = u|_T$. (Note that $\mathcal{T} \neq \emptyset$, because f_n is constant on non-free vertices.)

We need a couple of standard facts on Lipschitz functions, that we collect in the following lemma whose proof is left to the reader. \square

LEMMA 5.13. *Let $X \subset Y$ be metric spaces and let Z be a complete metric space. Denote by \bar{X} the closure of X in Y . Then:*

(1) *If $f : X \rightarrow Z$ is a L -Lipschitz map, then there is a L -Lipschitz map $\bar{f} : \bar{X} \rightarrow Z$ so that $\bar{f}|_X = f$.*

(2) *If $u_n : X \rightarrow Z$ is a sequence of L -Lipschitz maps and $u_\infty : \bar{X} \rightarrow Z$ is such that $u_n \rightarrow u_\infty$ on X , then the extensions $\bar{u}_n : \bar{X} \rightarrow Z$ converge to u_∞ .*

(3) *Suppose in addition that \bar{X} is compact, then the point-wise convergence of u_n is uniform.*

LEMMA 5.14. *If $(T, f) \in \mathcal{T}$ then $f_n|_T$ sub-converges uniformly to f .*

Proof. We restrict to the sub-sequence where $f_n|_T$ sub-converges. G acts by isometries on A, B . \bar{T} is G -invariant and admits a compact fundamental domain K . Since f_n are uniformly Lipschitz, we can apply Lemma 5.13, point 3 to K and get uniform convergence on K . The uniform convergence on T follows from G -equivariance of f_n and f . \square

If $\{(T_i, \varphi_i)\}$ is a chain in \mathcal{T} then, by Lemma 5.14 and a standard argument on sub-sequences, $(\bigcup_i T_i, \bigcup_i \varphi_i)$ is an upper bound. Therefore, \mathcal{T} has a maximal element.

Let (T, f_∞) be a maximal element of \mathcal{T} . If we show that $T = A$ we are done because $f_\infty = \lim f_n$ realizes the minimum Lipschitz constant and $F = PL(f_\infty)$ will be PL and with the same Lipschitz constant.

LEMMA 5.15. *T contains all non-free vertices and it is closed.*

Proof. Both claims follow from maximality of T . The first is because G -equivariance implies that $f_n(v) = \text{Fix}_B(\text{Stab}_A(v))$ is a constant sequence. The second is an immediate consequence of Lemma 5.13. \square

Assuming that $T \neq A$ and maximal we shall derive a contradiction. Let $x \in \partial T \cap A$ be fixed for the remainder of the proof. As T is closed $x \in T$. Define

$$\lambda_A = \min_{w \in VA: w \neq x} d_A(x, w), \quad \lambda_B = \min_{w \in VB: w \neq f_\infty(x)} d_B(f_\infty(x), w),$$

$$\lambda = \min(\lambda_A, \lambda_B/2L).$$

Choose $y \notin T$ such that $d_A(x, y) < \lambda$.

Since $f_n(x) \rightarrow f_\infty(x)$ eventually on n we have

$$f_n(y) \in B(f_\infty(x), \lambda_B).$$

Note that $B(f_\infty(x), \lambda_B)$ is star-shaped, namely it contains at most one vertex and contains exactly one vertex if and only if $f_\infty(x)$ is a vertex of B .

If $f_n(y)$ sub-converges, we can extend f_∞ to y and then extend equivariantly contradicting the maximality of T . Therefore $f_n(y)$ does not sub-converge. In particular, this implies that $B(f_\infty(x), \lambda_B)$ does not have compact closure, hence

$$f_\infty(x) = v$$

is non-free vertex of B . Also, for the same reason $f_n(y) \neq v$ eventually on n , and so after passing to a sub-sequence we may assume that $f_n(y) \neq v \forall n$.

Define

$$C_n(y) = \text{the connected component of } A \setminus f_n^{-1}(v) \text{ containing } y.$$

The rest of the argument is devoted to proving that $C_n(y)$ is collapsible eventually on n . This contradicts the fact that f_n are not collapsible and completes the proof.

LEMMA 5.16. *Let $C \subset A \setminus f_n^{-1}(v)$ be a connected subset such that there is a $w \in C$ and $\text{Id} \neq g \in G$ with $gw \in C$. Then, there exist two connected component $D_1(g)$ and $D_2(g)$ of $B \setminus v$, depending only on g , such that*

$$f_n(C) \subset D_1(g) \cup D_2(g).$$

Proof. Since C is connected $f_n(C)$ is contained in a single component of $B \setminus v$, the point is the independence from n .

Set $\text{Inv}_B(g)$ to be $\text{axis}_B(g)$ if g is hyperbolic and $\text{Inv}_B(g) = \text{Fix}_B(g)$ if g is elliptic. $\text{Inv}_B(g)$ is either a line or a single point. Therefore, it intersects at most two components of $B \setminus v$ that we denote $D_1(g)$ and $D_2(g)$ (possibly $D_1(g) = D_2(g)$).

The segment $[w, gw]_A$ is contained in C and $f_n([w, gw]_A) \cap \text{Inv}_B(g) \neq \emptyset$. Therefore $f_n(C)$ is contained in $D_1(g) \cup D_2(g)$. □

LEMMA 5.17. *For any $\text{Id} \neq g \in G$, we have $gC_n(y) \cap C_n(y) = \emptyset$ eventually on n .*

Proof. We argue by contradiction. So there exist g such that $gC_n(y) \cap C_n(y) \neq \emptyset$ frequently on n . By Lemma 5.16, $f_n(C_n(y)) \subset D_1(g) \cup D_2(g)$ frequently on n . Therefore, $f_n(y) \in B(v, \lambda_B) \cap (D_1(g) \cup D_2(g))$ which is the union of at most two open segments and hence has compact closure. In particular, $f_n(y)$ sub-converges contradicting the maximality of T . □

Lemma 5.17 is not enough to conclude that $C_n(y)$ is collapsible because a priori for any given n there may be infinitely many g such that $gC_n(y) \cap C_n(y) \neq \emptyset$.

For any half-line starting from y let w the first vertex such that there is $\text{Id} \neq g_w \in G$ with $g_w w \in [y, w]$. Let K be the union of all such segments. By construction, K is a simplicial tree containing y . Also, the diameter of K is

finite because there are finitely many orbit of vertices. Moreover, the interior of K does not contain any non-free vertex. Therefore, K is a finite simplicial tree. We denote by $\text{int}(K)$ the interior of K as a tree, which is not necessarily an open subset of A .

LEMMA 5.18. *Eventually on n we have $C_n(y) \subset \text{int}(K)$.*

Proof. Since K is finite the collection $G_K = \{g_w : w \text{ a leaf of } K\}$ is finite. By Lemma 5.17 for every $g_w \in G_K$, $g_w C_n(y) \cap C_n(y) = \emptyset$ eventually on n . Up to passing to a sub-sequence, we may suppose that this happens for any n . Therefore, $C_n(y)$ cannot contain the segment $[g_w w, w]$. Since $C_n(y)$ is connected and contains y , it follows that it does not contain any of the leaves of K . The claim follows. \square

LEMMA 5.19. *There are only finitely many $g \in G$ such that $g(\text{int}(K)) \cap \text{int}(K) \neq \emptyset$.*

Proof. Suppose $g(\text{int}(K)) \cap \text{int}(K) \neq \emptyset$. Then it contains at least either a vertex or an open segment. Let σ denote either a vertex or an open segment. Since K is a finite tree, it contains finitely many open edges and vertices. For any $\sigma \in \text{int}(K)$ there are only finitely many $g \in G$ such that $g^{-1}\sigma \in \text{int}(K)$, again because K is a finite tree, the action of G on edges is free and $\text{int}(K)$ contains only free vertices. \square

As a direct corollary of Lemmas 5.18 and 5.19, we get that the family of elements $g \in G$ such that $gC_n(y) \cap C_n(y)$ could possibly be non-empty is finite and independent of n . Therefore, by Lemma 5.17 eventually on n , for all $\text{Id} \neq g \in G$, $gC_n(y) \cap C_n(y) = \emptyset$. That is to say, $C_n(y)$ is eventually collapsible. A contradiction.

6. Optimal maps

In this section, we describe a class of maps, called optimal maps, which provide a useful tool for computing stretching factors and studying train track maps.

Definition 6.1 (Train track, from [3]). A *pre-train track structure* on a G -tree T is a G -invariant equivalence relation on the set of germs of edges at each vertex of T . Equivalence classes of germs are called *gates*. A *train track structure* on a G -tree T is a pre-train track structure with at least two gates at every vertex. A *turn* is a pair of germs of edges emanating from the same vertex. A turn is *legal* if the two germs belong to different equivalent classes. An immersed path is legal if it has only legal turns.

Definition 6.2. Given $A, B \in \mathcal{O}$ and a PL-map $f : A \rightarrow B$, we denote by $A_{\max}(f)$ (or simply A_{\max}) the subgraph of A consisting on those edges e of A for which $S_{f,e} = \text{Lip}(f)$. That is to say, the set of edges maximally stretched by f .

Note that A_{\max} is G -invariant. We notice that in literature the set A_{\max} is often referred to as *tension graph*.

Definition 6.3. Let $A, B \in \mathcal{O}$ and $f : A \rightarrow B$ be a PL-map. The pre-train track structure induced by f on A is defined by declaring germs of edges to be equivalent if they have the same *non-degenerate* f -image.

Definition 6.4 (Optimal map). Let $A, B \in \mathcal{O}$. A PL-map $f : A \rightarrow B$ is *not optimal* at v if A_{\max} has only one gate at v for the pre-train track structure induced by f . Otherwise f is *optimal* at v . The map f is *optimal* if it is optimal at all vertices.

We notice that with the above definition, f is optimal at any vertex outside A_{\max} .

REMARK 6.5. An PL-map $f : A \rightarrow B$ is optimal if and only if the pre-train track structure induced by f is a train track structure on A_{\max} . In particular if $f : A \rightarrow B$ is an optimal map, then at every vertex v of A_{\max} there is a legal turn in A_{\max} .

LEMMA 6.6. *Let $A, B \in \mathcal{O}$ and let $f : A \rightarrow B$ be a PL-map. Then every non-free vertex has at least two gates. In particular, f is optimal at non-free vertices.*

Proof. Let v be a non-free vertex of A and let x be an edge emanating from v . If the germ of x is collapsed by f to $f(v)$, then for any $1 \neq \gamma \in \text{Stab}(v)$ also γx is collapsed to $f(v)$. By definition such two germs x and γx are not equivalent in the pre-train track structure induced by f , so v has at least two gates. If the germ of x is not collapsed to $f(v)$ then by equivariance we have $f(\gamma x) = \gamma f(x)$, and since B has trivial edge-stabilizers $\gamma f(x)$ is different from $f(x)$. Therefore, x and γx have different non-degenerate images and thus are not equivalent. Hence, v has again at least two gates. □

LEMMA 6.7. *Let $A, B \in \mathcal{O}$ and let $f : A \rightarrow B$ be a PL-map. If f is not optimal, then there is a PL-map $h : A \rightarrow B$ such that either $\text{Lip}(h) \lesssim \text{Lip}(f)$ or $A_{\max}(h) \subsetneq A_{\max}(f)$ (or both).*

Proof. Let v be a (free) vertex of A_{\max} where f is not optimal, and let e be an edge of A_{\max} incident to v .

For $t \in [0, l_A(e)]$ let p_t be the point in e at distance t from v . Let f_t be the unique PL-map $A \rightarrow B$ such that for $w \in VA$

$$f_t(w) = \begin{cases} f(w) & \text{if } w \neq gv, g \in G, \\ f_t(gv) = gf(p_t) & g \in G. \end{cases}$$

For small enough t , if all the edges of A_{\max} are incident to a point in the orbit of v , then we obtain that $\text{Lip}(f_t) \lesssim \text{Lip}(f)$; otherwise, we get that $A_{\max}(f_t)$ is obtained from $A_{\max}(f)$ by removing the edges of A_{\max} incident to v and its orbit. We set $h = f_t$. □

COROLLARY 6.8. *For any $A, B \in \mathcal{O}$ there exists an optimal map $h : A \rightarrow B$. Moreover, if a PL-map $f : A \rightarrow B$ is not optimal but minimizes the Lipschitz constant, then there is an optimal map h such that $A_{\max}(h) \subsetneq A_{\max}(f)$.*

Proof. By Lemma 4.3, we know

$$\Lambda_R(A, B) \leq \inf \{ \text{Lip}(h) \text{ s.t. } h : A \rightarrow B \text{ is an } \mathcal{O}\text{-map} \}.$$

By Theorem 5.12, there exists a PL-map f with $\text{Lip}(f)$ minimal. Among such maps, we choose f so that A_{\max} is the smallest possible. By Lemma 6.7, f is optimal.

As for the second claim, recall that there are finitely many orbits of edges. So if $A_{\max}(h) \subsetneq A_{\max}(f)$, then the number of orbits of edges in $A_{\max}(h)$ is strictly less than that in $A_{\max}(f)$. Therefore, given f that minimizes the Lipschitz constant, repeated use of Lemma 6.7 gives the desired conclusion. \square

Definition 6.9. Let f be as in Corollary 6.8. By $\text{Opt}(f)$ we mean any optimal map h as given in that corollary.

Definition 6.10. Given $A, B \in \mathcal{O}$ and a PL-map $f : A \rightarrow B$, a sub-tree $L \subset A$ is *tight* if $L \subset A_{\max}$ and $f|_L$ is injective.

We notice that if $L = \text{axis}_A(g)$ for some $g \in G$, and L is tight, then $f(L) = \text{axis}_B(g)$.

THEOREM 6.11. *For any optimal map $f : A \rightarrow B$ there is an element $g \in G$ so that its axis in A is tight. In particular $l_B(g)/l_A(g) = \text{Lip}(f) = \Lambda_R(A, B)$.*

Proof. Let $f : A \rightarrow B$ be any optimal map. By Remark 6.5, every vertex of A_{\max} has a legal turn. Since pre-train track structures are G -equivariant, we can G -equivariantly associate to any edge e of A_{\max} incident to a vertex v , a legal turn $\tau(e)$ at v in A_{\max} , containing e . This defines a successor of e in A_{\max} . Starting from an edge e_0 in A_{\max} , the path obtained by concatenating successors, defines an embedded legal half-line, which eventually becomes periodic because A has finitely many orbits of edges. The period $g \in G$ has the requested properties. \square


7. Folding paths and geodesics

7.1. Local folds. For this sub-section, we fix $A, B \in \mathcal{O}$, and a PL-map $f : A \rightarrow B$.

Definition 7.1 (Isometric folding relations). For any $v \in VA$, $t \in \mathbb{R}$ and pair τ of edges $\tau = (\alpha, \beta)$ with $S_\alpha(f) = S_\beta(f)$, and emanating from v such that $f(\alpha)$ and $f(\beta)$ agree on some non-trivial segment. We define an equivalence relation $\sim_{\tau, t}$ on A as follows. First, we declare $x \in \alpha$ and $y \in \beta$ to be equivalent if $d(x, v) = d(y, v) \leq t$ and, after the isometric identification of $[v, x]$ and $[v, y]$,

we have $f|_{[v,x]} = f|_{[v,y]}$. Then we extend this relation to the orbit of τ by equivariance.

Definition 7.2. Given v, τ as above, and $t \in \mathbb{R}$ we define $A_{\tau,t} = A \sim_{\tau,t}$ equipped with the metric making the quotient map $q_{\tau,t} : A \rightarrow A_{\tau,t}$ a local isometry. The map f splits as



We say that A_t is obtained by (equivariantly) folding τ . If τ is understood, we shall abuse notation and suppress the subscript τ .

By definition, a fold depends on how f overlaps edges. When necessary, we will say that a fold is *directed by f* to emphasize this fact.

LEMMA 7.3. *In the present setting, for any t we have that either $\text{Lip}(PL(f_t)) < \text{Lip}(f)$ or $A_{t \max}(PL(f_t)) \subseteq q_t(A_{\max}(f))$.*

Proof. Fix t . Let σ be the f_t -pullback simplicial structure on A_t (that is to say, we declare new vertices the pre-images of vertices). We write A_t^σ and A_t to distinguish between σ and the original simplicial structure of A_t . Note that f_t is then a σ -PL map.¹ As q_t is a local isometry, for any edge of σ we have $S_e(f_t) = S_{q_t^{-1}(e)}(f)$. In particular, $\text{Lip}(f_t) = \text{Lip}(f)$ and $A_{t \max}^\sigma(f_t) = q_t(A_{\max}(f))$. Now the edge-stretching factors of $PL(f_t)$ are less than or equal to those of f_t . Hence, $\text{Lip}(PL(f_t)) < \text{Lip}(f_t)$ or $A_{t \max}(PL(f_t)) \subseteq A_{t \max}^\sigma(f_t)$. \square

The following lemma will be useful in the study of train track maps. For $\Phi \in \text{Aut}(G, \mathcal{O})$, $\Phi(A)$ and A are the same metric tree with different G -action. So $\Phi(A)_t = \Phi(A_t)$ and we use the same symbol q_t to denote the quotient map from $\Phi(A) \rightarrow \Phi(A_t)$. We have the following commutative diagram which defines the map h_t (Figure 1).

LEMMA 7.4. *Let $\Phi \in \text{Aut}(G, \mathcal{O})$ and suppose that $f : A \rightarrow \Phi(A)$ is a PL-map such that*

$$\text{Lip}(f) = \min\{\text{Lip}(h), h : X \rightarrow \Phi(X), X \in \mathcal{O}\}.$$

Let A_t be the tree obtained by perform a local fold directed by f . Then

$$\text{Lip}(PL(h_t)) = \text{Lip}(PL(f_t)) = \text{Lip}(f)$$

and

$$A_{t \max}(PL(h_t)) \subseteq q_t(A_{\max}(f)).$$

¹ Note that f_t may be not PL for the original structure (see Definition 5.1). This happens for example if there is an edge of A which is mapped by f through the turn τ .




FIGURE 1. The quotient map h_t .

Moreover, if $PL(h_t)$ is not optimal, then

$$A_{t_{\max}}(\text{Opt}(PL(h_t))) \subsetneq A_{t_{\max}}(PL(h_t)).$$

Proof. Since q_t is a local isometry, then $\text{Lip}(h_t) = \text{Lip}(f_t) = \text{Lip}(f)$. Passing to PL does not increase the Lipschitz constants, and by hypothesis $\text{Lip}(f)$ is minimal, thus

$$\text{Lip}(PL(h_t)) = \text{Lip}(PL(f_t)) = \text{Lip}(f).$$

Hence, by Lemma 7.3 we have that $A_{t_{\max}}(PL(f_t)) \subseteq q_t(A_{\max}(f))$. Now, note that $PL(q_t \circ PL(f_t)) = PL(h_t)$. Therefore, since q_t is a local isometry, we have $A_{t_{\max}}(PL(h_t)) \subseteq A_{t_{\max}}(PL(f_t))$. Whence, the second claim. The last claim is a direct consequence of Corollary 6.8. \square

7.2. Folding paths. Now we describe paths joining any two points of \mathcal{O} which are geodesics w.r.t. the metric d_R . The procedure is exactly that used in [20] in the case of free groups.

First, we restrict attention to the special situation where $f : A \rightarrow B$ is a map so that $A_{\max}(f) = A$ and $\text{Lip}(f) = 1$.

For a complete simple fold, we mean the path obtained by equivariantly folding two edges α and β as much as possible. That is to say, the path $[0, m] \rightarrow \mathcal{O}$


$$t \mapsto A_{\tau,t},$$

where $\tau = (\alpha, \beta)$, $M = \min\{L_A(\alpha), L_A(\beta)\}$, $\sim_{\tau,M}$ is not trivial, and m is the minimum t so that $\sim_{\tau,t} = \sim_{\tau,M}$.

PROPOSITION 7.5. *Let $A, B \in \mathcal{O}$ and $f : A \rightarrow B$ a PL -map such that $A_{\max}(f) = A$ and $\text{Lip}(f) = 1$. Then there exists a path from A to B which is a concatenation of complete simple folds directed by f (and successive maps f_t).*

Proof. Let σ be the simplicial structure induced on A by f , so that f maps σ -edges to edges. Note that by definition of σ if the initial segments of two edges of σ have the same f -image then the two edges have the same f -image. On the other hand, if f admits no simple folds, then f is a G -equivariant isometry from A to B , and hence $A = B$. Otherwise, let A_t be a tree obtained by a complete simple fold.

Then f splits as



and σ induces a simplicial structure on A_t with fewer orbits of edges. Induction completes the proof. \square

Now we come back to the general case.

Definition 7.6 (Isometric folding paths). Let $A, B \in \mathcal{O}$ and $f : A \rightarrow B$ be a PL-map such that $A_{\max}(f) = A$. A *isometric folding path* from A and B , and directed by f , is a path obtained as follows.

- First, rescale the metric on A so that $\text{Lip}(f) = 1$ and call that point A_0 . Note that $A_0 = \text{Lip}(f)A$.
- Then, consider a path $\hat{\gamma}(t) = A_t$ from A_0 to B given by Proposition 7.5, parametrized by $t \in [0, 1]$.
- Finally, rescale $\hat{\gamma}$ by $\gamma(t) = \hat{\gamma}(t) / \text{Lip}(f)^{(1-t)}$.

THEOREM 7.7. *For any $A, B \in \mathcal{PO}$ there is a d_R -geodesic in \mathcal{PO} from A to B .*

Proof. We use the following characterization of (unparameterized) geodesics.

LEMMA 7.8. *An oriented path γ from A to B in (the volume-one slice of) \mathcal{PO} is a d_R -geodesic if and only if there is a $g \in G$ hyperbolic so that*

$$\Lambda_R(C, D) = \frac{l_D(g)}{l_C(g)}$$

for any $C < D \in \gamma$.

Proof. A path is geodesic if and only if intermediate triangular inequalities are equalities (see, for instance, [20]), that is to say, for any $C \in \gamma$ we have

$$d_R(A, B) = d_R(A, C) + d_R(C, B).$$

Let $g \in G$ hyperbolic so that $\Lambda_R(A, B) = \frac{l_B(g)}{l_A(g)}$. We have

$$\Lambda_R(A, B) = \frac{l_B(g)}{l_A(g)} = \frac{l_C(g)}{l_A(g)} \frac{l_B(g)}{l_C(g)} \leq \sup_{\xi} \frac{l_C(\xi)}{l_A(\xi)} \sup_{\xi} \frac{l_B(\xi)}{l_C(\xi)} = \Lambda_R(A, C) \Lambda_R(C, B)$$

and equality holds true if and only if

$$\frac{l_C(g)}{l_A(g)} = \sup_{\xi} \frac{l_C(\xi)}{l_A(\xi)}, \quad \frac{l_B(g)}{l_C(g)} = \sup_{\xi} \frac{l_B(\xi)}{l_C(\xi)}.$$


Applying the same reasoning to the geodesic $\gamma|_{CB}$ we get that γ is geodesic if and only if g is maximally stretched form C to D for any $A \leq C < D \leq B$ in γ . □

We now describe a path in \mathcal{O} which projects to a d_R -geodesics in \mathcal{PO} .

Let $f : A \rightarrow B$ be an optimal map and let $g \in G$ be an element with tight axis in A (Theorem 6.11). First, we equivariantly rescale all the edges of A as follows

$$e \mapsto S_e(f)e$$

obtaining a tree A_0 . Call p the projection map $p : A \rightarrow A_0$ clearly f splits as



Note that if $S_e(f) = 0$ for some edge e , than f collapses e , so the tree A_0 is actually in \mathcal{O} . Moreover, the map $f_0 : A_0 \rightarrow B$ is PL and has the property that $A_{0\max}(f_0) = A_0$ and $\text{Lip}(f_0) = 1$. Finally, note that if $S_e(f) = 0$ then $e \notin A_{\max}(f)$, therefore p restricts to a homeomorphism from $A_{\max}(f)$ to its image. In particular, this implies that g has tight axis in A_0 .

We do this operation continuously so that we have an oriented path γ_1 from A to a A_0 . It is clear that for any two points C, D in γ_1

$$\Lambda_R(C, D) = \frac{l_D(g)}{l_C(g)}.$$

Let γ_2 be an isometric folding path from A_0 to B directed by f_0 . Define γ to be the concatenation of γ_1 and γ_2 .

Since g has a tight axis in A_0 , by definition f_0 is injective on $\text{axis}_{A_0}(g)$. Therefore, $\text{axis}_{A_0}(g)$ is never folded in γ_2 . In particular, $l_{\gamma_2(t)}(g)$ is constant. Since γ_2 is an isometric folding path, nothing is stretched by a factor ≥ 1 , hence g realizes $\Lambda_R(C, D) = 1$ for any two ordered points in γ_2 .

This implies that for any $C < D$ in γ

$$\Lambda_R(C, D) = \frac{l_D(g)}{l_C(g)}.$$

Finally, note that the above condition is scale invariant, in the sense that if

$$\Lambda_R(C, D) = \frac{l_D(g)}{l_C(g)}$$

then

$$\Lambda_R(\lambda C, \mu D) = \frac{l_{\mu D}(g)}{l_{\lambda C}(g)}$$

for any $\lambda, \mu > 0$.

Therefore, the projection of γ to the set of co-volume one elements of \mathcal{O} is an (unparameterized) d_R -geodesic. \square

8. Train tracks

In this section, we prove that any irreducible automorphism in $\text{Aut}(G, \mathcal{O})$ is represented by a train track map (see below the definitions). We follow the approach *à la Bers* of M. Bestvina ([3]). The following arguments are restatement of those for the case of free groups. In fact, the proofs do not require adjustment due to the fact that we are allowing non trivial stabilizer for vertices, and one could just say that the theory of train tracks for free groups passes to the case of free products without any substantial change. We refer to [4], [5], [6], [7], [8] for the train track theory for free groups. However, there are many facts that are well-known for free groups, at least to the experts, but for which there is no reference in literature. We take the occasion of the present discussion on train tracks for free product to give explicit statements and metric proofs of some of these facts, as for instance the relations between the minimal displaced set and the set of train tracks. (See also [23], [33].)

8.1. Irreducibility and minimal displacement trichotomy. Let $\Phi \in \text{Aut}(G, \mathcal{O})$. It acts on \mathcal{O} by changing the marking. Define

$$\lambda_\Phi = \inf_{X \in \mathcal{O}} \Lambda_R(X, \Phi(X)).$$

Note that since both X and $\Phi(X)$ have the same volume, the number λ_Φ cannot be smaller than one. Therefore, there are three cases: Φ is *elliptic*, if $\log \lambda_\Phi$ is zero and the infimum is attained; *parabolic*, if the infimum is not attained; *hyperbolic* if $\log \lambda_\Phi$ is positive and attained.

For $T \in \mathcal{O}$, we say that a Lipschitz surjective map $f : T \rightarrow T$ represents Φ if for any $g \in G$ and $t \in T$ we have $f(gt) = \Phi(g)(f(t))$. (In other words, if it is an \mathcal{O} -map from T to $\Phi(T)$.)

Definition 8.1. We say $\Phi \in \text{Aut}(G, \mathcal{O})$ is \mathcal{O} -irreducible (or simply irreducible for short) if for any $T \in \mathcal{O}$ and for any $f : T \rightarrow T$ representing Φ , if $W \subset T$ is a proper f -invariant G -subgraph then $G \setminus W$ is a union of trees each of which contains at most one non-free vertex.

This is related to an algebraic definition of irreducibility as follows. Suppose that G can be written as a free product, $G = G_1 * G_2 * \dots * G_p * G_\infty$, where we allow the possibility that G_∞ is trivial. Then we say that the set $\mathcal{G} = \{[G_i] : 1 \leq i \leq p\}$ is a *free factor system* for G , where $[G_i] = \{gG_i g^{-1} : g \in G\}$ is the set of conjugates of G_i . Given two free factor systems $\mathcal{G} = \{[G_i] : 1 \leq i \leq p\}$ and $\mathcal{H} = \{[H_j] : 1 \leq j \leq m\}$, we write $\mathcal{G} \sqsubseteq \mathcal{H}$ if for each i there exists a j such that $G_i \leq gH_j g^{-1}$ for some $g \in G$. We write $\mathcal{G} \sqsubset \mathcal{H}$ if either one of the previous inclusions is strict or $p < m$. We also say that $\mathcal{G} = \{[G_i] : 1 \leq i \leq p\}$ is proper if $\mathcal{G} \sqsubset \{[G]\}$.

We say that $\mathcal{G} = \{[G_i] : 1 \leq i \leq p\}$ is Φ invariant for some $\Phi \in \text{Out}(G)$ if for each i , $\Phi([G_i]) = [G_j]$ for some j . We shall restrict our attention to those free factor systems \mathcal{G} such that $\{[H]\} \sqsubseteq \mathcal{G}$ whenever H is an indecomposable free factor which is *not* a free group. In particular, this means that $G = G_1 * G_2 * \dots * G_p * G_\infty$, and $G_\infty \cong F_k$ for some free group F_k . Associated to such a free factor system $\mathcal{G} = \{[G_i] : 1 \leq i \leq p\}$ we have the space of trees $\mathcal{O} = \mathcal{O}(G, (G_i)_{i=1}^p, F_k)$ and any (outer) automorphism of G leaving \mathcal{G} invariant will act on \mathcal{O} in the usual way.

Definition 8.2. Let \mathcal{G} be a free factor system of G as above and suppose it is Φ invariant for some $\Phi \in \text{Out}(G)$. Then Φ is called irreducible relative to \mathcal{G} if \mathcal{G} is a maximal (under \sqsubseteq) proper, Φ -invariant free factor system.

LEMMA 8.3. *Suppose \mathcal{G} is a free factor system of G with associated space of trees \mathcal{O} , and further suppose that \mathcal{G} is Φ -invariant. Then Φ is irreducible relative to \mathcal{G} if and only if Φ is \mathcal{O} -irreducible.*

Proof. Suppose ϕ is \mathcal{O} -reducible. Then there is $T \in \mathcal{O}$ and $W \subset T$ a proper ϕ -invariant G -subgraph such that $W \setminus G$ contains at least a component whose fundamental group is non-trivial (as graph of groups) and either contains only free vertices or strictly contains a conjugate of the G'_i 's. In both cases, \mathcal{G} is not maximal.

On the other hand, suppose \mathcal{G} is not maximal. Then there is a ϕ -invariant free factor system $\mathcal{H} \{[H_j] : 1 \leq j \leq m\}$ (hence $G = H_1 * \dots * H_m * F_r$) such that $\mathcal{G} \sqsubset \mathcal{H}$. Let $T \in \mathcal{O}(G, (H_j)_{j=1}^m, F_r)$. Since any G_i 's is contained in some conjugate of some H_j 's, any H_j has an induced splitting whose factors are conjugated to some G_i 's. Then we can replace any non-free vertex of T stabilized by a conjugate of some H_j , with a tree in the outer space $\mathcal{O}(H_j)$ corresponding to that splitting (that may be a point if $[G_i] = [H_j]$). Let W be the collection of such trees. Since \mathcal{H} is ϕ -invariant, so is W . The fact that $\mathcal{G} \sqsubset \mathcal{H}$ assures that there $W \setminus G$ is *not* a union of trees each of which contains at most a non-free vertex. Hence, ϕ is \mathcal{O} -reducible. □

We prove now that irreducible automorphisms are hyperbolic.

THEOREM 8.4. *For any irreducible $\Phi \in \text{Aut}(G, \mathcal{O})$, $\lambda_\Phi = \inf_{X \in \mathcal{O}} \Lambda_R(X, \Phi(X))$ is a minimum and obtained for some $X \in \mathcal{O}$.*

Proof. We essentially follow the proof in [3], noting that the technical difficulties arise due to the fact that our space is not locally compact and that our action is not proper. We shall utilise the Sausages lemma result, Theorem 9.10, which is proved in the subsequent section, but whose proof is independent on the results in this section.

In order to proceed, we demonstrate the contrapositive, that an automorphism, $\Phi \in \text{Aut}(G, \mathcal{O})$, for which $\lambda_\Phi = \inf_{X \in \mathcal{O}} \Lambda_R(X, \Phi(X))$ is not a minimum

is reducible. So we suppose that $\Gamma_k \in \mathcal{O}$ is a minimising sequence for Φ . That is, $\lim_{k \rightarrow \infty} d_R(\Gamma_k, \Phi\Gamma_k) \rightarrow \log \lambda_\Phi$.

We notice that $d_R(\Gamma, \Phi\Gamma)$ is scale-invariant as a function of Γ , and hence descends to a function on \mathcal{PO} . For the remaining part of this proof, we work with the co-volume one slice of \mathcal{O} , which we still denote by \mathcal{O} for simplicity of notation.

Our first step is to show that the trees Γ_k cannot stay in the “thick” part of \mathcal{O} . The ϵ -thick part of (the co-volume one slice of) \mathcal{O} consists of all trees $X \in \mathcal{O}$ such that $l_X(g) \geq \epsilon$ for all hyperbolic $g \in G$.² Note that the ϵ -thick part of \mathcal{O} is co-compact for any $\epsilon > 0$.

More precisely, we wish to show that only finitely many of the Γ_k lie in the ϵ -thick part of \mathcal{O} for any $\epsilon > 0$. For suppose not, then passing to a subsequence we may assume that all Γ_k belong to the ϵ -thick part and then, again by taking subsequences and invoking co-compactness, we may find $\Psi_k \in \text{Aut}(G, \mathcal{O})$ such that $\Psi_k(\Gamma_k)$ converges to some Γ_∞ which is again in the ϵ -thick part of \mathcal{O} . Hence, $d_R(\Gamma_\infty, \Psi_k\Phi\Psi_k^{-1}\Gamma_\infty) \rightarrow \log \lambda_\Phi$ because of the continuity of d_R .

However, note that as we are dealing with simplicial trees, the translation lengths of the elements in a given tree form a discrete set. In fact, the set $\{l_{\Psi(\Gamma_\infty)}(g) : g \in G\}$ is the *same* discrete set for any $\Psi \in \text{Aut}(G, \mathcal{O})$. Moreover, by Theorem 9.10, $d_R(\Gamma_\infty, \Psi_k\Phi\Psi_k^{-1}\Gamma_\infty)$ are given by the quotient of the translation lengths of candidates, and there are only finitely many possible lengths of candidates in Γ_∞ (even though there will, in general, be infinitely many candidates) and therefore the distances $d_R(\Gamma_\infty, \Psi_k\Phi\Psi_k^{-1}\Gamma_\infty)$ also form a discrete set. Hence, there must exist some k (in fact infinitely many) such that $d_R(\Gamma_\infty, \Psi_k\Phi\Psi_k^{-1}\Gamma_\infty) = \log \lambda_\Phi$, whence we obtain that $\lambda_\Phi = \Lambda_R(X, \Phi(X))$, where $X = \Psi_k^{-1}\Gamma_\infty$. The first step is proved.

Now for any $\Gamma \in \mathcal{O}$, we let Γ^ϵ be the sub-forest obtained as the union of all the hyperbolic axes of elements of G whose translation length is less than ϵ , along with all the vertices. Since there are only finitely many graphs of groups arising from \mathcal{O} , each of which is finite, there exists an ϵ such that for all Γ , Γ^ϵ is a proper sub-forest of Γ (we remind that we are now working with co-volume one trees). Call this ϵ_0 .

Also notice that each such sub-forest is a G -invariant subgraph, and hence there is a bound on the length of any proper chain of such sub-forests. Call this number B .

Now let $\epsilon_i = \epsilon_0/(\lambda_\Phi + 1)^i$. Choose $\Gamma = \Gamma_k$ as above such that $d_R(\Gamma, \Phi\Gamma) < \log(\lambda_\Phi + 1)$ and Γ not in the ϵ_B -thick part (so Γ^{ϵ_B} is non-trivial). Now,

$$\Gamma \neq \Gamma^{\epsilon_0} \supseteq \Gamma^{\epsilon_1} \supseteq \dots \supseteq \Gamma^{\epsilon_B}$$

² The ϵ -thick part of the co-volume one slice of \mathcal{O} can also be defined as the set of trees $X \in \mathcal{O}$ such that the systole of $X \setminus G$ is at least ϵ . By systole here we mean the shortest length of a loop in $X \setminus G$ in the sense of graph of groups (for example, a segment whose vertices are not free is a loop whose length is the double of the metric length of the segment).

is a chain of $(B + 1)$ non-trivial sub-forests of Γ . Therefore, they cannot all be distinct. However, if $f : \Gamma \rightarrow \Phi\Gamma$ is any optimal map, then—as it is $(\lambda_\Phi + 1)$ -Lipschitz— f must send Γ^{δ_i} into $\Gamma^{\delta_{i-1}}$. Hence, we must have an f -invariant subgraph of Γ which is non-trivial (cf Definition 8.1), and hence, Φ is reducible. \square

Hence, from now on we will use the fact that all our \mathcal{O} -irreducible elements of $\text{Aut}(G, \mathcal{O})$ are hyperbolic.

8.2. Minimally displaced points. For an irreducible automorphism Φ , we introduce the set of minimally displaced points, which plays the role of a “metric” axis for Φ . We will show later that this coincides with the set of train tracks.

Definition 8.5 (Minimal displaced set). Let Φ be an \mathcal{O} -irreducible element of $\text{Aut}(G, \mathcal{O})$. We define the minimal displaced set of Φ by

$$\mathcal{M}(\Phi) = \{T \in \mathcal{O} : \Lambda_R(T, \Phi(T)) = \lambda_\Phi\}.$$

THEOREM 8.6. *Let Φ be an \mathcal{O} -irreducible element of $\text{Aut}(G, \mathcal{O})$. Then, if $T \in \mathcal{M}(\Phi)$ and $f : T \rightarrow \Phi(T)$ is an optimal map, we have*

$$T_{\max}(f) = T.$$

Proof. We consider f as either an \mathcal{O} -map from $T \rightarrow \Phi(T)$ or a map $f : T \rightarrow T$ representing Φ , without distinction. If T_{\max} is f -invariant we are done because, since by Theorem 6.11 T_{\max} contains the axis of some hyperbolic element, irreducibility implies $T = T_{\max}$.

In the subsequent argument we shall perform small perturbations on T by changing edge-lengths. The map f will induce maps on these new trees which are the same as f set-wise. Formally, we have many different pairs of trees and associated maps, but that we still call (T, f) .

Suppose that T_{\max} is not f -invariant. Then there is an edge e in T_{\max} whose image contains a sub-segment of an edge a which is not in T_{\max} . We shrink a by a small amount. If the perturbation is small enough, $\text{Lip}(f)$ is not increased, and since $\text{Lip}(f) = \lambda_\Phi$ this remains true after the perturbation. Therefore, e is no longer in T_{\max} , a is still not in T_{\max} , and T_{\max} must contain some other edge b with $S_b(f) = \lambda_\Phi$.

Note that after perturbation f might no longer be optimal. However, by Corollary 6.8 there is an optimal map $h : T \rightarrow \Phi(T)$ with

$$T_{\max}(h) \subsetneq T_{\max}(f).$$

Since $\text{Lip}(f) = \lambda_\Phi$, then also $\text{Lip}(h) = \lambda_\Phi$. If $T_{\max}(h)$ is not f -invariant, we repeat this argument recursively. After finitely many steps we must end obtaining a finite sequence of maps h_i with the properties that $\text{Lip}(h_i) = \lambda_\Phi$ and

$$\emptyset \neq T_{\max}(h_0) \subsetneq T_{\max}(h_1) \subsetneq \cdots \subsetneq T_{\max}(f) \subsetneq T.$$

Since we stopped, $T_{\max}(h_0)$ is f -invariant. Therefore by irreducibility $T_{\max}(h_0) = T$, and the above condition implies $T = T_{\max}(f)$. \square

REMARK 8.7. In the proof of Theorem 8.6, we showed the following fact which needs no assumption on reducibility, and may be of independent interest: If (T, f) are so that first, T locally weak minimizes $d(T, \Phi T)$,³ and second T_{\max} is locally⁴ minimal,⁵ then T_{\max} is f -invariant.

LEMMA 8.8. *Suppose $T \in \mathcal{M}(\Phi)$ and suppose that $f : T \rightarrow T$ is a Lipschitz map with $\text{Lip}(f) = \lambda_\Phi$. Then f is optimal.*

Proof. First, consider $PL(f)$. Since $\text{Lip}(PL(f)) \leq \text{Lip}(f)$ and $\text{Lip}(f)$ is minimal, we have $\text{Lip}(PL(f)) = \lambda_\Phi$. Moreover, combining Lemma 7.4 and Theorem 8.6, we get that $PL(f)$ is optimal and $T_{\max}(PL(f)) = T$.

Since the maps $PL(f)$ have the property that $S_e(PL(f)) \leq S_e(f)$ with inequality being strict at some edge only if f is not PL , it follows that $f = PL(f)$. \square

THEOREM 8.9 ($\mathcal{M}(\Phi)$ is fold-invariant). *Let Φ be an \mathcal{O} -irreducible element of $\text{Aut}(G, \mathcal{O})$. Then, the set $\mathcal{M}(\Phi)$ is invariant under folding directed by optimal maps.*

More precisely, if $T \in \mathcal{M}(\Phi)$ and $f : T \rightarrow \Phi(T)$ is an optimal map, and if T_t is an isometric folding path from $T \rightarrow \Phi(T)$ directed by f , then we have:

- (a) $T_t \in \mathcal{M}(\Phi)$.
- (b) *The quotient maps $h_t : T_t \rightarrow \Phi(T_t)$, defined by the diagram in Figure 1, are optimal.*

In particular, any local fold directed by f stays in $\mathcal{M}(\Phi)$.

Proof. Claim (a) is a direct consequence of Lemma 7.4. Claim (b) follows from Lemma 8.8 because $\text{Lip}(h_t) = \text{Lip}(f) = \lambda_\Phi$. \square

8.3. Train track maps. Train track maps can be defined via topological properties as well as metric properties. In this section, we relate the two point of view.

Recall that we defined pre-train track and train track structures in Definition 6.1, and that in our notation a train track structure is required to have at least two gates at every vertex.

Definition 8.10 (Train track map). A PL-map $f : T \rightarrow T$ representing Φ is a *train track map* if there is a train track structure on T so that

- (1) f maps edges to legal paths (in particular, f does not collapse edges);

³ That is, $d(T, \Phi T) \leq d(T', \Phi T')$ for any T' in the same simplex of T sufficiently close to T in the Euclidean coordinates.

⁴ In the simplex of T .

⁵ With respect to the inclusion.

(2) If $f(v)$ is a vertex, then f maps inequivalent germs at v to inequivalent germs at $f(v)$.

Here a some remark is needed. First, we note that a part the PL requirement, this definition is topological and does not involves the metric on T . In fact if f is train track and we change the metric on edges of T , then up to re-PL-ize f it remains train track. For these reasons we have to distinguish between (topological) train track maps and (metric) optimal train track maps.

The second remark on the definition of train track map is that, given an optimal map f representing Φ , we can consider two pre-train track structures, namely that given by f and that generated by the iterates f^k (see Definition 6.3). We denote the two structures in the following way

$$\sim_f \quad \text{and} \quad \langle \sim_{f^k} \rangle.$$

So, two germs are \sim_f -equivalent if they are identified by f , they are \sim_{f^k} -equivalent if they are identified by f^k and they are $\langle \sim_{f^k} \rangle$ -equivalent if they are identified by some power of f .

In particular one may ask if f is a train track for \sim_f or for $\langle \sim_{f^k} \rangle$.

LEMMA 8.11. *Suppose $f : T \rightarrow T$ is a PL-map representing $\Phi \in \text{Aut}(G, \mathcal{O})$. If f is a train track map for some unspecified structure \sim , then $\sim \supseteq \langle \sim_{f^k} \rangle$.⁶ In particular, if f is a train track map for \sim_f , then $\sim_f = \langle \sim_{f^k} \rangle$.*

Proof. The last claim follows from the first because $\sim_f \subseteq \langle \sim_{f^k} \rangle$ by definition. Suppose $\tau = (e_1, e_2)$ is a turn and suppose that e_1 and e_2 are in the same gate for $\langle \sim_{f^k} \rangle$. Then there is some k so that $f^k(e_1) = f^k(e_2)$, choose the first $k \geq 1$ so that this happens. Either $f^{k-1}(\tau)$ is contained in an edge (that is to say, the image of the vertex of τ is not a vertex) or it is a turn. The first case is not allowed since f is a \sim -train track map. Therefore, $f^{k-1}(\tau)$ is an turn. Since f identifies the two germs of $f^{k-1}(\tau)$, by Condition (2) of Definition 8.10, $f^{k-1}(\tau)$ is \sim -illegal. Since f is a train track map, turns that are pre-images of illegal turns are illegal. It follows that τ is illegal, that is to say, $e_1 \sim e_2$. □

COROLLARY 8.12. *Is f is a train track map for some \sim , then it is train track for $\langle \sim_{f^k} \rangle$.*

Proof. Condition (2) of Definition 8.10 is automatically satisfied for $\langle \sim_{f^k} \rangle$. Lemma 8.11 tells us that \sim -legal implies $\langle \sim_{f^k} \rangle$ -legal. Thus, also condition (1) of Definition 8.10 is satisfied. □

Note that a priori f could be train track for $\langle \sim_{f^k} \rangle$ but not for \sim_f .

⁶ Hence, $\langle \sim_{f^k} \rangle$ is a train track structure and not only a pre-train track structure.

LEMMA 8.13 (Topological characterization of train track maps). *If Φ is irreducible, then for a map f representing Φ , to be a train track map is equivalent to the condition that there is a hyperbolic $g \in G$ with axis L so that $f^k|_L$ is injective $\forall k \in \mathbb{N}$.*

Proof. Let L be as in the hypothesis. The iterate images of L form a proper f -invariant sub-graph of T containing the axis of a hyperbolic element. Since Φ is irreducible, such sub-graph is the whole T . The pre-train track structure $\langle \sim_{f^k} \rangle$ is a train track structure and it is readily checked that f satisfies Conditions (1) and (2) of Definition 8.10 with respect to $\langle \sim_{f^k} \rangle$.

On the other hand, suppose that f is a train track map. We argue as in the proof of Theorem 6.11. As train track structures have at least two gates at every vertex, there is a hyperbolic element g with legal axis L , and Conditions (1) and (2) imply that this remains true under f -iterations.

By Corollary 8.12 it is a $\langle \sim_{f^k} \rangle$ -train track map, and $\langle \sim_{f^k} \rangle$ -legality of L implies injectivity of $f^k|_L$. □

Definition 8.14 (Train track bundle). Let Φ be an \mathcal{O} -irreducible element of $\text{Aut}(G, \mathcal{O})$. We define the train track bundle as

$$TT(\Phi) = \{T \in \mathcal{O} : \exists \text{ an optimal train track map } f : T \rightarrow T \text{ representing } \Phi\}.$$

We notice that the Axis bundle \mathcal{A}_Φ of Φ is defined in [23] as the closure of the union of all the sets $TT(\Phi^k)$.

Definition 8.15 (Strict train tracks). Let Φ be an \mathcal{O} -irreducible element of $\text{Aut}(G, \mathcal{O})$. We define the strict train track bundle as

$$TT_0(\Phi) = \{T \in \mathcal{O} : \exists f : T \rightarrow \Phi T \text{ optimal which is train track for } \sim_f\}.$$

LEMMA 8.16. *Suppose $f : T \rightarrow T$ is a train track map representing an \mathcal{O} -irreducible Φ . Then there is a rescaling of edges of T such that every edge is stretched the same, and f becomes optimal. In particular, if f is a train track for \sim_f , then the (closed) simplex of T contains a point $T' \in TT_0(\Phi)$.*

Proof. We argue as in the proof of Theorem 8.6. For any assignment of edge-lengths, that is to say for any point in the simplex of T , we have the corresponding stretching factors of $PL(f)$ (remember that f is PL for T but may need a reparametrization when we change edge-lengths). The reparametrization of f along edges does not affect the train track property of f . So $PL(f)$ is train track. As in the second part of the proof of Theorem 8.4, one can see that irreducibility implies that there is a point in the closed simplex of T that realizes the minimum, over the simplex, of $\text{Lip}(PL(f))$. We choose T' among such points so that it minimizes the tension graph with respect to the inclusion. As in Theorem 8.6, we see that the tension graph is Φ -invariant. If $T'_{\max} \setminus G$ is trivial (a forest whose trees contains at most one non-free vertex) we can collapse it obtaining a point in a facet of the simplex of T with strictly

smaller $\text{Lip}(PL(f))$ contradicting the minimality of T' . By irreducibility of Φ we have $T'_{\max} = T'$. Since the map $f' = PL_{T'}(f)$ is a train track map, and $T'_{\max} = T'$, then f' is optimal and $T' \in TT_0(\Phi)$. \square

Now, we want to prove that minimally displaced points and train tracks coincides. As a first observation, we have.

LEMMA 8.17. *For any $\Phi \in \text{Aut}(G, \mathcal{O})$, we have $TT_0(\Phi) \subseteq TT(\Phi) \subseteq \mathcal{M}(\Phi)$.*

Proof. If $T \in TT(\Phi)$ and f is an optimal train track map, then there is g such that $l_T(\Phi^n(g)) = \text{Lip}(f)^n l_T(g)$. On the other hand, if $Q \in \mathcal{M}(\Phi)$ then $l_Q(\Phi^n(g)) \leq \lambda_\Phi^n l_Q(g)$. Therefore, we have

$$\left(\frac{\text{Lip}(f)}{\lambda_\Phi}\right)^n \frac{l_T(g)}{l_Q(g)} \leq \Lambda_R(\Phi^n T, \Phi^n Q) = \Lambda_R(T, Q) < \infty$$

which implies $\text{Lip}(f) = \lambda_\Phi$. \square

THEOREM 8.18. *Let Φ be an irreducible element of $\text{Aut}(G, \mathcal{O})$. Then, $\emptyset \neq TT_0(\Phi)$ is dense in $\mathcal{M}(\Phi)$ with respect to the simplicial topology.*

Proof. Since Φ is hyperbolic there exists an element $T \in \mathcal{O}$ that realizes λ_Φ . Let $f : T \rightarrow \Phi(T)$ be an optimal map, which exists by Corollary 6.8. By Theorem 8.6, we know that $T_{\max}(f) = T$.

We consider the pre-train track structure \sim_f induced by f on T . As f is optimal the pre-train track structure is a train track structure (no one-gate vertex). Now, say that a vertex of T is *foldable* if it contains at least a gate with two elements.

By Theorem 8.9, up to perturbing T by as small an amount as required via a finite number of equivariant folds, we may assume that any foldable vertex has valence exactly 3. In particular, any foldable vertex has exactly two gates.

Moreover, again by Theorem 8.9 we may assume that T locally maximizes the number of orbits of foldable vertices.

We claim that in this situation f is a train track map with respect to \sim_f . First, we check Condition (1) of Definition 8.10.

Suppose that an edge e of T_{\max} has illegal image. Then it passes through an illegal turn τ . We equivariantly fold τ by a small amount t . The result is a new tree T_t and an induced map h_t . By Theorem 8.9 $T_t \in \mathcal{M}(\Phi)$ and h_t is optimal. But $e \notin T_{t \max}(h_t)$, so $T_{t \max} \neq T_t$, which is impossible by Theorem 8.6.

Now we check Condition (2). By definition of our pre-train track structure, inequivalent germs are mapped to different germs. Now, suppose that a legal turn η at a vertex v is mapped to an illegal turn τ at a vertex w . We equivariantly fold τ by a small amount getting a tree T_t and a map h_t . Theorem 8.9 guarantees that $T_t \in \mathcal{M}(\Phi)$ and h_t is optimal. Now, η became foldable. By

optimality, there are no one-gate vertices. Thus, we increased the number of foldable vertices in contradiction with our assumption of maximality. \square

In fact, more is true.

THEOREM 8.19. *Let Φ be an irreducible element of $\text{Aut}(G, \mathcal{O})$. Then, $TT(\Phi) = \mathcal{M}(\Phi)$.*

Proof. What we are going to prove is that if f is an optimal map representing Φ , then it is a train track map for $\langle \sim_{f^k} \rangle$. We need a couple of lemmas.

LEMMA 8.20. *Suppose $f : T \rightarrow T$ is a PL-map representing $\Phi \in \text{Aut}(G, \mathcal{O})$. If f is a train track map for \sim_f , then f^k , which represents Φ^k , is a train track map for \sim_{f^k} .*

Proof. By Lemma 8.11 $\sim_f = \langle \sim_{f^k} \rangle$, whence $\sim_f = \sim_{f^k}$ for any k . Since f is a train track for \sim_f , in particular \sim_f is a train track structure, so any vertex has at least two gates. Conditions (1) and (2) of Definition 8.10, that hold for f , imply that $f^k(e)$ is a legal path, hence Condition (1) for f^k .

If τ is a turn and $f^k(\tau)$ is \sim_{f^k} -illegal, then $f^k(\tau)$ is \sim_f -illegal, which implies that τ is \sim_f -illegal, and so \sim_{f^k} -illegal. \square

LEMMA 8.21. *If $TT_0(\Phi) \neq \emptyset$, then $\lambda_{\Phi^k} = (\lambda_\Phi)^k$.*

Proof. Let $T \in TT_0(\Phi) \subseteq \mathcal{M}(\Phi)$ and let $f : T \rightarrow T$ be an optimal train track map with respect to \sim_f . By Lemma 8.20 f^k is a train track map for \sim_{f^k} , in particular $\text{Lip}(f^k) = \text{Lip}(f)^k$, \sim_{f^k} is a train track structure and f^k is optimal by Remark 6.5. By Theorem 6.11 $(\lambda_\Phi)^k = \text{Lip}(f^k) = \lambda_{\Phi^k}$. \square

We can now conclude the proof of Theorem 8.19. Let $T \in \mathcal{M}(\Phi)$, and let $f : T \rightarrow T$ be an optimal map. By Lemmas 8.21 and 8.8, all the iterates f^k are optimal.

We claim that f is a train track map with respect to $\langle \sim_{f^k} \rangle$. First, note that since f^k is optimal, every vertex has at least two gates, hence $\langle \sim_{f^k} \rangle$ is a train track structure.

Now, we check Condition (1) of Definition 8.10. Suppose that an edge e is folded by some f^k , and choose the first k so that this happens. Let p be a point interior to e where a fold occurs.

By Theorem 8.18, there is T_t obtained from T by a finite number of as-small-as-required folds so that $T_t \in TT(\Phi)$ (in fact, in $TT_0(\Phi)$). Without loss of generality, we may suppose that T_t is obtained by T by a simple fold, and show that in this case $T \in TT(\Phi)$. Let $h_t : T_t \rightarrow T$ be the map induced by the fold as in Figure 1. We choose the fold small enough so that p remains in the

interior of the edge e_t corresponding to e . The following diagram commutes

$$\begin{array}{ccccccc}
 T & \xrightarrow{f} & \Phi(T) & \xrightarrow{f^2} & \Phi(T) & \xrightarrow{f^3} & \Phi(T) & \dots & \xrightarrow{f^k} & \Phi(T) \\
 \downarrow q_t & \nearrow f_t & \downarrow q_t & \nearrow f_t & \downarrow q_t & \nearrow f_t & \downarrow q_t & & & \downarrow q_t \\
 T_t & \xrightarrow{h_t} & \Phi(T_t) & \xrightarrow{(h_t)^2} & \Phi(T_t) & \xrightarrow{(h_t)^3} & \Phi(T_t) & \dots & \xrightarrow{(h_t)^k} & \Phi(T_t)
 \end{array}$$

and therefore e_t is folded at p by (h_t^k) . But in the proof of Theorem 8.18, we have seen that the h_t are train track maps, so edges are never folded.

As for Condition (2) of Definition 8.10, note that Condition (1) and the definition of $\langle \sim_{f^k} \rangle$ imply that Condition (2) is automatically satisfied. \square

REMARK 8.22. We notice that in [23] the authors ask if, given Φ , there is N such that \mathcal{A}_Φ is the closure of $\bigcup_{i=1}^N TT(\Phi^i)$. In the same work they provide an example of a Φ and point in $X \in \mathcal{A}_\Phi$ not supporting any train track for any Φ^i . This example, together with our Theorem 8.19, provides a negative answer to the question, since $TT(\Phi^i) = \mathcal{M}(\Phi)$ is closed, and hence any finite union of $TT(\Phi^i)$ is closed.

THEOREM 8.23 (Folding axis). *Let Φ be an \mathcal{O} -irreducible element of $\text{Aut}(G, \mathcal{O})$. Then $TT_0(\Phi)$ is invariant under folding directed by optimal train track maps.*

More precisely, if $f : T \rightarrow \Phi(T)$ is train track map with respect to \sim_f , and T_t is an isometric folding path from $T \rightarrow \Phi(T)$ directed by f , then the induced map $h_t : T_t \rightarrow \Phi(T_t)$ is a train track map with respect to \sim_{h_t} .

Proof. Let $\gamma_0(t) = T_t$ be a folding path from T to $\Phi(T)$, directed by f . Since f is an optimal train track map, there is a folding path from T to $\Phi^2(T)$ directed by f^2 , given by the concatenation of γ_0 and $\Phi(\gamma_0)$. Therefore, we can form a folding line directed by f (and its powers) by concatenating the paths $\Phi^k(\gamma_0)$.

Let g be an element such that $\text{axis}_T(g)$ is legal and $f^k(\text{axis}_T(g))$ is legal for any k . (Such an element exists because f is a train track map.) It follows that $\text{axis}_T(g)$ is never folded during the folding procedure, so $\text{axis}_{T_t}(g)$ is legal and $h_t^k(\text{axis}_{T_t}(g))$ is legal. Thus, h_t is a train track map as desired. \square

It is useful to have train track maps that respect the simplicial structure (i.e., that map vertices to vertices). The presence of non-free vertices is in this case an advantage with respect to the classical case ([3], [8]). We give a detailed proof of the following result in full generality by completeness. We notice that we make no use of Perron-Frobenius theory.

THEOREM 8.24 (Simplicial train track). *Let $\Phi \in \text{Aut}(G, \mathcal{O})$ be irreducible. Then there exists a simplicial optimal train track map representing Φ . More*

precisely, if $T \in TT(\Phi)$, then the closed simplex of T contains a point admitting a simplicial (optimal) train track map.

Proof. The idea is to “snap” images of vertices to nearest vertices, as suggested in [3]. Let $T \in TT(\Phi)$ and $f : T \rightarrow T$ be an optimal train track map (with respect to $\langle \sim_{f^k} \rangle$) representing Φ . Let $\lambda = \lambda_\Phi > 1$ be the Lipschitz constant of f .

We will argue by induction on the number of orbits of vertices whose image is not a vertex (note in particular that such vertices are free) making use of local surgeries for the inductive step.

First, we describe in details the local move that we use, and after we will adjust the map f . The moves can be interpreted as local isometric folds followed by local isometric “unfolds”. However, we describe them in terms of surgeries because this viewpoint helps in controlling the derivative of f . We remark that we are not working with covolume-one trees, thus no rescaling is needed.

Choose $\varepsilon > 0$ small enough so that:

- (1) $B(w, \varepsilon) \cap B(w', \varepsilon) = \emptyset \ \forall w, w' \in VT : w \neq w'$.
- (2) $f(B(w, \varepsilon)) \cap B(w', \varepsilon) = \emptyset \ \forall w, w' \in VT : f(w) \neq w'$.

Let $v \in VT$ be such that there is $k \geq 0 : f^k(v) \notin VT$. In particular, v is free and has two gates, that we label as positive and negative. We build an isometric model of $B(v, \varepsilon)$ as follows. By our choice of ε , $B(\varepsilon, v)$ is star-shaped with say n_- negative and n_+ positive strands. Therefore, $B(\varepsilon, v)$ is isometric to the space obtained from n_- copies of $(-\varepsilon, 0]$ and n_+ copies of $[0, \varepsilon)$ by gluing the 0’s. See Figure 2, left-hand side.

For $|t| < \varepsilon$, let $B_t(v)$ be the space obtained from n_- copies of $(-\varepsilon, t]$ and n_+ copies of $[t, \varepsilon)$ by gluing the points t ’s. See Figure 2, right-hand side.

$B_t(v)$ has a vertex, corresponding to the endpoints “ t ’s”, which we denote by v_t , and the boundary of $B_t(v)$ is naturally identified with that of $B(v, \varepsilon)$.

Now, we cut out from T the whole G -orbit of $B(v, \varepsilon)$ and we paste back copies of $B_t(v)$ using the natural identifications on the boundaries. We say that we equivariantly moved v by t .




FIGURE 2. Local models for $B(\varepsilon, v)$ and $B_t(v)$.

Now, choose $v \in VT$ with $f(v) \notin VT$ and move it by $t < \varepsilon$ in the direction given by the nearest vertex to $f(v)$. (If $f(v)$ is a midpoint of an edge, we chose a direction.)

Define $pre(v) = \{w \in VT : f^k(w) = v \text{ for some } k \geq 0 \text{ and } f^i(w) \in VT \text{ for all } 0 \leq i \leq k\}$. Thus, $pre(v)$ consists of the iterate f -pre-images of v .

First, we see that $pre(v)$ is finite. For any vertex $w \in VT$ consider the sequence $w, f(w), f^2(w), \dots$. If it happens that $f^i(w)$ is a vertex for $i = 0, \dots, k$ and $f^{k+1}(w) = gw$ for some $g \in G$, then by equivariance $f^i(w)$ is a vertex for every i . Since $f(v)$ is not a vertex, then the elements in $pre(v)$ are in different g -orbits. Since there are only finitely many orbits of vertices, $pre(v)$ is a finite set.

Moreover, any $w \in pre(v)$ is free and if $w \neq v$ then $w \notin pre(w)$. Note also that $Gv \cap pre(v) = v$. Moreover, since f is a train track map, any $w \in pre(v)$ has two gates, with positive and negative labels determined by that at v via f^k .

For $w \in pre(v)$, if k is the first power so that $f^k(w) = v$, we consider the ball $B(w, \varepsilon/\lambda^k)$ and we equivariantly move w by t/λ^k . Note that this is possible because such balls are all disjoint from the G -orbit of each other, and disjoint from $f(B(v, \varepsilon))$ because $(\lambda > 1)$ and our choice of ε .

We denote by T_t the tree obtained from T in such a way.

We are now left to define $f_t : T_t \rightarrow T_t$. Let N be the union of the G -orbits of the metric balls $B(w, \varepsilon/\lambda^k)$ for all $w \in pre(v)$ and k as above, and let N_t be the union of the corresponding sets B_t 's (see Figure 2). Thus, $T_t = (T \setminus N) \cup N_t$.

On the set $f^{-1}(T \setminus N) \cap (T \setminus N)$ we set $f_t = f$. Clearly, there $\dot{f}_t = \lambda$.

Let $\sigma \subset f^{-1}(N) \setminus N$ be a segment without vertex in its interior. As σ is connected, $f(\sigma)$ is contained in one of the balls $B(w, \varepsilon)$. Since f is a train track map, edges are mapped to legal paths. Therefore, $f(\sigma)$ is contained in the union of a negative and a positive strand of $B(w, \varepsilon)$. Such union is isometric to the union of the corresponding strands in $B_t(w)$. We define f_t on σ by composing f with such isometry. Clearly $\dot{f}_t = \lambda$.

It remains to define f_t on N_t . For any $w \in pre(v)$, any strand S of $B_t(w)$ corresponds isometrically to a legal path σ in $B(w, \varepsilon/\lambda^k)$, which is uniformly stretched by f by factor λ . Since we moved w by t/λ^k and $f(w)$ by t/λ^{k-1} , $f(\sigma) \subset B(f(w), \varepsilon/\lambda^{k-1})$ corresponds isometrically to a strand S_1 in $B_t(f(w))$ (or $f(B(v, \varepsilon))$ if $w = v$). Therefore, f_t is defined by pre- and post-composing f with such isometries. Clearly, $\dot{f}_t = \lambda$. See Figure 3.

It is clear that (T, f) is homeomorphic to (T_t, f_t) (that is to say, there is an homeomorphism from T to T_t that conjugates f to f_t) and therefore f_t still is an optimal train track representing Φ .

We remark that when we move v by t , then $f(v)$ moved toward its nearest vertex u by λt . On the other hand, even if u has been moved, that was by




FIGURE 3. The definition of f_t on $B_t(w)$.

an amount of t/λ^k for some $k \geq 0$. Therefore, $f(v)$ approaches u at speed at least $\lambda - 1 > 0$.

We can finally run the induction on the number of orbits of vertices whose image is not a vertex. Let v be such a vertex. We move v as described as long as we can. Since $d(f(v), VT)$ is strictly decreasing, the process must stop. The process stops when we cannot choose $\varepsilon > 0$ with the required properties. That is to say, when either $f(v)$ is a vertex or some moved vertex collided with another vertex v' . (It is readily checked that the resulting limit map is train track.) In both cases, we decreased by one our induction parameter. \square

The following is a direct corollary of the existence of train track maps for free products. It was proved in [8] for free groups and in [9] for free products.

COROLLARY 8.25 ([8], [9]). *Let G be a group acting co-compactly on a tree with trivial edge stabilisers and freely indecomposable vertex stabilisers. Then any $\Phi \in \text{Out}(G)$ has a representative which is a relative train track map. In particular, relative train tracks for free groups exist.*

Proof. We can write $G = G_1 * \dots * G_p * F_k$, where the G_i are freely indecomposable and non-free. By the Kurosh subgroup theorem, any subgroup of G can be written as a free product of conjugates of subgroups of the G_i and some free group. So if $H \leq G$, then $H \cong A_1 * \dots * A_m * F_l$, for some $A_i \neq 1$ which are conjugates of subgroups of the G_i and some free group F_l of rank l . Define the Kurosh rank of such a subgroup H to be $m + l$, denoted $\kappa(H)$. Note this number may be infinite in general, but will certainly be finite if H is a free factor (and in many other cases).

Now define the reduced Kurosh rank of H to be $\bar{\kappa}(H) = \max(0, \kappa(H) - 1)$. The Kurosh rank of a free factor system, $\mathcal{G} = \{[G_i]\}$ is then defined to be $\kappa(\mathcal{G}) = \sum \kappa(G_i)$ and the reduced Kurosh rank of \mathcal{G} is defined to be $\bar{\kappa}(\mathcal{G}) = \sum \bar{\kappa}(G_i)$.

These are finite numbers, and if $\mathcal{G} \sqsubseteq \mathcal{H}$ then $\bar{\kappa}(\mathcal{G}) \leq \bar{\kappa}(\mathcal{H})$ and $\kappa(\mathcal{G}) \leq \kappa(\mathcal{H})$. Moreover, if $\mathcal{G} \sqsubset \mathcal{H}$, then either $\bar{\kappa}(\mathcal{G}) < \bar{\kappa}(\mathcal{H})$ or $\kappa(\mathcal{G}) < \kappa(\mathcal{H})$.

Hence given $\Phi \in \text{Out}(G)$ there is a maximal Φ -invariant, proper free factor system, with corresponding space of trees \mathcal{O} . A simplicial train track map representing Φ for some tree in \mathcal{O} is a relative train track map in the sense of [9]. \square

9. Computing stretching factors

This section is devoted to prove that stretching factors are realized by a class of particularly simple elements. We generalize the line used in [20], taking in account possible pathologies coming from the presence of non-free vertices. We remark that even if this section is at the end of the paper, the results of this section are independent from those in Section 8 (where Theorem 9.10 is used).

Notation 9.1. Let x, y, z, t be vertices of a G -tree, not necessarily different from each other. We write

$$\begin{array}{c} x \\ \diagdown \quad \diagup \\ y \quad z \end{array} \Leftrightarrow x \in [y, z] \quad \text{and} \quad \begin{array}{c} x \\ | \\ y \quad z \end{array} \Leftrightarrow x \notin [y, z].$$

Note that $\begin{array}{c} x \\ | \\ y \quad z \end{array}$, if and only if the segments $[y, x]$ and $[x, z]$ intersect in a subsegment starting at x which is not a single point. We will use the following two inference rules, whose verification is immediate.

$$\begin{array}{c} x \\ | \\ y \quad z \end{array} + \begin{array}{c} x \\ | \\ z \quad t \end{array} \Rightarrow \begin{array}{c} x \\ | \\ y \quad t \end{array}$$

$$\begin{array}{c} x \\ | \\ y \quad z \end{array} + \begin{array}{c} x \\ \diagdown \quad \diagup \\ z \quad t \end{array} \Rightarrow \begin{array}{c} x \\ \diagdown \quad \diagup \\ y \quad t \end{array}$$

We will now be concerned in finding *good* elements g with a tight axis.⁷ In the subsequent discussion when we say “ g does not have a tight axis” we mean either that the axis of g is not tight or that g is elliptic.

LEMMA 9.2. *Let $A, B \in \mathcal{O}$ and $f : A \rightarrow B$ be a PL-map. Let $L \subset A$ be a tight sub-tree isomorphic to \mathbb{R} . Suppose there is $g \in G$ and $x \in L$ so that $x \neq gx$ and $[x, gx] \subset L$. Set $y = f(x)$. If g does not have tight axis, then we have*

$$\begin{array}{c} y \\ | \\ g^{-1}y \quad gy \end{array} .$$

⁷ See Definition 6.10.



FIGURE 4. The segment $[x, gx]$.

Proof. Suppose g has a fixed point v . Since A is a tree, v is the middle point of $[x, gx]$, and in particular we have $\begin{matrix} x \\ \swarrow \quad \searrow \\ g^{-1}x \quad gx \end{matrix}$ (where the center of the tripod is exactly v). The claim follows.

If g has no fixed point, then it has an axis. The union of the segments $[g^kx, g^{k+1}x]$, as k varies in \mathbb{Z} , is a g -invariant tree. Therefore, it contains the axis of g . Thus, $\text{axis}_A(g)$ is contained in the g -orbit of $[x, gx]$. In particular $\text{axis}_A(g) \subset A_{\max}$ because $[x, gx] \subset L \subset A_{\max}$ which is a G -invariant sub-set of A . By hypothesis the axis of g is not tight, thus $f|_{\text{axis}_A(g)}$ is not injective. Since L is tight, then $f|_L$ is injective. Therefore, f must overlap an initial segment of $[x, gx]$ with a terminal segment of $[g^{-1}x, x]$ which is exactly the claimed formula. \square

LEMMA 9.3 (No triple points). *Let $A, B \in \mathcal{O}$, $f: A \rightarrow B$ be a PL-map, and $g \in G$ be such that $\text{axis}_A(g)$ is tight. If there exists $x \in \text{axis}_A(g)$ such that $|Gx \cap [x, gx]| \geq 4$ then, there exists $h \in G$ with tight axis, such that $l_A(h) < l_A(g)$.*

Proof. By hypothesis, there exists $x \in A$ and $a, b \in G$ such that $\text{axis}_A(g)$ locally looks as depicted in Figure 4.

In other words, the segment $[x, gx]$ is the concatenation of segments $[x, ax]$, $[ax, bx]$ and $[bx, gx]$. Let $y = f(x)$. Since $\text{axis}_A(g)$ is tight, we have a similar situation in B . See Figure 5. In particular, we have

$$\delta := \begin{matrix} ay \\ \swarrow \quad \searrow \\ y \quad by \end{matrix} \quad b^{-1}\delta := \begin{matrix} b^{-1}ay \\ \swarrow \quad \searrow \\ b^{-1}y \quad y \end{matrix} \quad \text{and} \quad a^{-1}\delta := \begin{matrix} y \\ \swarrow \quad \searrow \\ a^{-1}y \quad a^{-1}by \end{matrix}$$

We look at a, ba^{-1}, gb^{-1} (corresponding to single steps in Figure 5). Clearly, all of them have A -length strictly smaller than that of g . Thus, if one of them has tight axis we are done. We can therefore suppose that none of them has



FIGURE 5. The segment $[y, gy]$.

tight axis. In particular, by Lemma 9.2 we know

$$\alpha := \begin{array}{c} y \\ \swarrow \quad \searrow \\ a^{-1}y \quad ay \end{array} \quad \beta := \begin{array}{c} ay \\ \swarrow \quad \searrow \\ ab^{-1}ay \quad by \end{array} .$$

Now we look at b (corresponding to a double step in Figure 5). As above, we have $l_A(b) < l_A(g)$. We argue by contradiction assuming that b has no tight

axis. Then by Lemma 9.2, we have $\begin{array}{c} y \\ \swarrow \quad \searrow \\ by \quad b^{-1}y \end{array}$. Moreover, by assumption we have $ay \in (y, by)$ whence $b^{-1}ay \in (b^{-1}y, y)$; from which we get

$$\chi := \begin{array}{c} y \\ \swarrow \quad \searrow \\ ay \quad b^{-1}ay \end{array} .$$

It follows that

$$\alpha + \chi + (a^{-1}\beta) = \begin{array}{c} y \\ \swarrow \quad \searrow \\ a^{-1}y \quad ay \end{array} + \begin{array}{c} y \\ \swarrow \quad \searrow \\ ay \quad b^{-1}ay \end{array} + \begin{array}{c} y \\ \swarrow \quad \searrow \\ b^{-1}ay \quad a^{-1}by \end{array} \Rightarrow \begin{array}{c} y \\ \swarrow \quad \searrow \\ a^{-1}y \quad a^{-1}by \end{array}$$

which contradicts $a^{-1}\delta = \begin{array}{c} y \\ \swarrow \quad \searrow \\ a^{-1}y \quad a^{-1}by \end{array}$. □

Definition 9.4. When the hypothesis of Lemma 9.3 are satisfied, we say that axis_A(g) has triple points.

LEMMA 9.5 (Four points lemma). *Let $A, B \in \mathcal{O}$ and $f: A \rightarrow B$ be a PL-map. Let $L \subset A$ be a tight sub-tree isomorphic to \mathbb{R} . Suppose there is $x \neq v \in L$ and $a, b \in G$ such that $ax \neq bv$ and $ax, v \in [x, bv] \subset L$ (see Figure 6). Let*

$y = f(x)$, $w = f(v)$. If $\begin{array}{c} y \\ \swarrow \quad \searrow \\ a^{-1}y \quad ay \end{array}$ and $\begin{array}{c} w \\ \swarrow \quad \searrow \\ b^{-1}w \quad bw \end{array}$, then $b^{-1}a$ has a tight axis, given by the iterates of $[x, b^{-1}ax]$. (Thus, by Lemma 9.2, we have that either a , or b , or $b^{-1}a$ has a tight axis.)





FIGURE 6. The two possibilities for the segment $[x, bv]$.

Proof. First of all, note that the tripod hypotheses imply $x \neq ax$ and $v \neq bv$ (and thus $x \neq bv$). We have to show that the tree formed by the iterates of the segment $[x, b^{-1}ax]$ is a tight line. Clearly it is contained in A_{\max} because L is, and $[x, b^{-1}ax]$ is contained in the union of the two segments $[x, v]$ and $[v, b^{-1}ax]$ (the latter may be a single point). Note that a priori, we may have



Since $[v, b^{-1}ax] = b^{-1}([bv, ax])$ and $[ax, bv] \subset L$, we know that f is injective on both $[x, v]$ and $[v, b^{-1}ax]$. If we prove

$$I := \begin{array}{c} w \\ | \\ y \quad b^{-1}ay \end{array} \quad \text{and} \quad II := \begin{array}{c} y \\ | \\ a^{-1}by \quad b^{-1}ay \end{array}$$

we are done. Indeed, I implies $\begin{array}{c} v \\ | \\ x \quad b^{-1}ax \end{array}$ and since $x \neq v$ we have $x \neq b^{-1}ax$. Now II and Lemma 9.2 conclude.

First, we prove I . Since both ax and v lie in $[x, bv]$, $ax \neq bv$ gives $\begin{array}{c} bv \\ | \\ ax \quad v \end{array}$.

Applying f we get $\begin{array}{c} bw \\ | \\ ay \quad w \end{array}$ whence, acting with b^{-1} , $\begin{array}{c} w \\ | \\ b^{-1}ay \quad b^{-1}w \end{array}$. By

hypothesis, we know $\begin{array}{c} w \\ | \\ b^{-1}w \quad bw \end{array}$. Summing up we get

$$\begin{array}{c} w \\ | \\ b^{-1}ay \quad b^{-1}w \end{array} + \begin{array}{c} w \\ | \\ b^{-1}w \quad bw \end{array} \Rightarrow \begin{array}{c} w \\ | \\ b^{-1}ay \quad bw \end{array}.$$

Now, if we had $\begin{array}{c} w \\ | \\ y \quad b^{-1}ay \end{array}$ we would get $\begin{array}{c} w \\ | \\ y \quad b^{-1}ay \end{array} + \begin{array}{c} w \\ | \\ b^{-1}ay \quad bw \end{array} \Rightarrow \begin{array}{c} w \\ | \\ y \quad bw \end{array}$

which is impossible because L is tight. Thus, $\begin{array}{c} w \\ | \\ y \quad b^{-1}ay \end{array}$ is not true and I is true. Note that this implies that $a \neq b$.

We now prove *II*. From tightness of L , we get $\begin{matrix} ay \\ bw \quad y \end{matrix}$ whence

$$a^{-1}bw \quad a^{-1}y \quad \text{and thus} \quad \begin{matrix} y \\ a^{-1}bw \quad a^{-1}y \end{matrix} + \begin{matrix} y \\ a^{-1}y \quad ay \end{matrix} \Rightarrow \begin{matrix} y \\ a^{-1}bw \quad ay \end{matrix}.$$

Acting on I by $a^{-1}b$ we get $\begin{matrix} a^{-1}bw \\ a^{-1}by \quad y \end{matrix}$ that, together with $\begin{matrix} y \\ a^{-1}bw \quad ay \end{matrix}$,

implies $\begin{matrix} y \\ a^{-1}by \quad ay \end{matrix}$ because by hypothesis $ax \neq bw$ and so $y \neq a^{-1}bw$. As

above, since both ax, v lie in $[x, bw]$ we have $\begin{matrix} x \\ ax \quad v \end{matrix}$ whence $\begin{matrix} y \\ ay \quad w \end{matrix}$. It

follows $\begin{matrix} y \\ a^{-1}by \quad w \end{matrix}$ that, together with $\begin{matrix} w \\ y \quad b^{-1}ay \end{matrix}$ gives

$$\begin{matrix} y \\ a^{-1}by \quad b^{-1}ay \end{matrix} = II. \quad \square$$

LEMMA 9.6 (No crossing points). *Let $A, B \in \mathcal{O}$, $f: A \rightarrow B$ be a PL-map, and $g \in G$ be such that $\text{axis}_A(g)$ is tight. Suppose that there exists points $x, v \in \text{axis}_A(g)$ and $a, b \in G$ such that $ax, bv \in \text{axis}_A(g)$ and, with respect to the linear order of $\text{axis}_A(g)$, we have $x < v < ax < bv < gx$ (see Figure 7) Then, there exists $h \in G$ with tight axis, such that $l_A(h) < l_A(g)$.*

Proof. Note that we may have $a = b$. Let $y = f(x)$ and $w = f(v)$. Since $[x, ax], [v, bv]$ lie in $\text{axis}_A(g)$, they lie in A_{\max} . Also $l_A(a), l_A(b) < l_A(g)$. If one of them has tight axis, then the claim is proved by letting $h = a$ or

$h = b$. Otherwise, by Lemma 9.2 we have $\begin{matrix} y \\ a^{-1}y \quad ay \end{matrix}$ and $\begin{matrix} w \\ b^{-1}w \quad bw \end{matrix}$ and by

Lemma 9.5 the element $h = b^{-1}a$ has a tight axis (in particular, in this case $a \neq b$). On the other hand $l_A([hx, x]) = l_A([ax, bx]) \leq l_A([ax, bv]) + l_A([x, v]) < l_A(g)$, therefore $l_A(h) < l_A(g)$. \square




FIGURE 7. The segment $[x, gx]_A$.





FIGURE 8. The local situation in $\text{axis}_A(g)$.

Definition 9.7. When the hypothesis of Lemma 9.6 are satisfied we say that $\text{axis}_A(g)$ has crossing points.

LEMMA 9.8 (No bad triangles). *Let $A, B \in \mathcal{O}$, $f: A \rightarrow B$ be a PL-map, and $g \in G$ be such that $\text{axis}_A(g)$ is tight. If there exists $x, v, t \in \text{axis}_A(g)$ and $a, b, c \in G$ such that $xa, bv, ct \in [x, gx] \subseteq \text{axis}_A(g)$ and, with respect to the linear order of $[x, gx]$, we have $x < ax < v < bv < t < ct < gx$ (see Figure 8) then there exists $h \in G$ with tight axis, such that $l_A(h) < l_A(g)$.*

Proof. Let $y = f(x), w = f(v)$ and $s = f(t)$. We first try $h = a, b, c$. Since $[x, ax], [v, bv]$ and $[t, ct]$ lie in A_{\max} and $l_A(x), l_A(y), l_A(z) < l_A(g)$, if one of them has a tight axis we are done. Otherwise, by Lemma 9.2 we have



that $b^{-1}a, c^{-1}b$ and $c^{-1}a$ all have tight axis. It suffices to show that one of them has length less than g in A . Let $\lambda_1 = l_A([x, ax]), \lambda_2 = l_A([ax, v]), \lambda_3 = l_A([v, bv]), \lambda_4 = l_A([bv, t]), \lambda_5 = l_A([t, ct])$.

$$l_A(b^{-1}a) \leq l_A([x, b^{-1}ax]) \leq l_A([x, v]) + l_A([v, b^{-1}ax]) = \lambda_1 + \lambda_2 + \lambda_2 + \lambda_3,$$

$$l_A(c^{-1}b) \leq l_A([bv, ct]) + l_A([v, t]) = \lambda_4 + \lambda_5 + \lambda_3 + \lambda_4.$$

Summing up

$$l_A(b^{-1}a) + l_A(c^{-1}b) < 2(\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4 + \lambda_5) \leq 2l_A(g).$$

So one of them has length strictly less than g in A . □

Definition 9.9. When the hypothesis of Lemma 9.8 are satisfied we say that $\text{axis}_A(g)$ has bad triangles.

THEOREM 9.10 (Sausage lemma and candidates). *Let $A, B \in \mathcal{O}$. Then the minimal stretching factor $\Lambda_R(A, B)$ is realized by an element g such that the projection of $\text{axis}_A(g)$ to $G \setminus A$ is of the form:*

- (i) *Embedded simple loop: O ;*
- (ii) *embedded figure-eight: ∞ (a bouquet of two copies of S^1);*
- (iii) *embedded barbel: $O - O$ (two simple loops joined by a segment);*
- (iv) *embedded singly degenerate barbell: $\bullet - O$ (a non-free vertex and a simple loop joined by a segment);*
- (v) *embedded doubly degenerate barbell: $\bullet - \bullet$ (two non-free vertices joined by a segment).*

The loops and segments above may contain free and non-free vertices.

Proof. By Corollary 6.8 and Theorem 6.11, we know that there is an optimal map $f : A \rightarrow B$ and an element g so that $\text{axis}_A(g)$ is tight. Translation lengths of hyperbolic elements form a discrete set, so we may assume that g has minimal translation length among those elements with tight axis. In order to prove our claim, it suffices to find an element with tight axis whose projection to $G \setminus A$ is of one of the types (i), ..., (v). By Lemmas 9.3, 9.6 and 9.8, we know that $\text{axis}_A(g)$ has no triple points, nor crossing points, nor bad triangles.

Choose $x_0 \in \text{axis}_A(g)$. Since there are no triple points in $\text{axis}_A(g)$ there cannot be three distinct points in $[x_0, gx_0)$ with the same image in $G \setminus A$. Hence, every point in $G \setminus A$ has at most two pre-images in $[x_0, gx_0)$. Pairs of points with the same image are exactly those of the form $\{x, ax\}$ with $a \in G$ and $x \neq ax$. Call such a pair a pair of double points.

If $[x_0, gx_0)$ has no pairs of double points then we are in case (i). If it has exactly a pair of double points we are in case (ii). Thus, we have reduced to the case where there are at least two pairs of double points.

Without loss of generality, we may assume that $x = x_0$ and $\{x, ax\}$ is a pair of double points. There is a second pair of double points $\{v, bv\}$ and we may assume that $v < bv$ with respect to the orientation of $[x, gx)$. Since there are no crossing points, we have either

$$x < ax < v < bv < gx$$

or

$$x < v < bv < ax < gx$$

by interchanging the role of x and v we may assume we are in the first case. Since the translation lengths of hyperbolic elements form a discrete set we may, after possibly replacing $\{x, ax\}$ by another pair of double points, assume that either there are no pairs of double points in $[x, ax)$ or a is elliptic. Similarly for $\{v, bv\}$. By minimality of the translation length of g neither a nor b can have tight axis. Therefore, Lemma 9.5 applies and $b^{-1}a$ has a tight axis formed by the iterated of $[x, b^{-1}ax]$.

In the order induced by the axis of $b^{-1}a$ we have Figure 9.

Since there are no bad triangles in $\text{axis}_A(g)$ the segment $[ax, v]$ projects injectively to $G \setminus A$, and the same is true for $[b^{-1}v, b^{-1}ax]$, which projects to same path in the quotient but with opposite orientation. If there are no pairs of double points in $[x, ax)$, then $[x, ax]$ projects to a simple closed curve. If a is elliptic, it must fix the middle point of $[x, ax]$. The similar picture holds true




FIGURE 9. The axis of $b^{-1}a$ in A .

for $[v, bv]$. Therefore, $\text{axis}_A(b^{-1}a)$ projects to a barbell, possibly degenerate depending on whether a, b are elliptic or not. \square

Acknowledgments. We are grateful to Yago Antolin Pichel and Camille Horbez for the many interesting discussions we had with them. Many thanks go to Sebastian Meinert for its helpful comments on a previous version of this manuscript. We want to thank the CRM of Barcelona and the LABEX of Marseille for having hosted great research periods in geometric group theory, where we had the occasion to discuss the subject of the present paper with so many great people. We are clearly in debt with the organizers of such events.

We would also like to thank the Istituto Nazionale di Alta Matematica for support during the period in which this research was carried out.

REFERENCES

- [1] Y. Agolm-Kfir and M. Bestvina, *Asymmetry of outer space*, *Geom. Dedicata* **156** (2012), 81–92. [MR 3329728](#)
- [2] R. Alperin and H. Bass, *Length functions of group actions on Λ -trees*. *Combinatorial Group Theory and Topology, Sel. Pap. Conf., alta/Utah 1984*, *Ann. of Math. Stud.* **111** (1987), 265–378. [MR 0895622](#)
- [3] M. Bestvina, *A Bers-like proof of the existence of train tracks for free group automorphisms*, *Fund. Math.* **214** (2011), no. 1, 1–12. [MR 2845630](#)
- [4] M. Bestvina, M. Feighn and M. Handel, *Laminations, trees, and irreducible automorphisms of free groups*, *Geom. Funct. Anal.* **7** (1997), no. 7, 215–244.
- [5] M. Bestvina, M. Feighn and M. Handel, *The tits alternative for $\text{Out}(Fn)$. I: Dynamics of exponentially-growing automorphisms*, *Ann. of Math. (2)* **151** (2000), no. 2, 517–623.
- [6] M. Bestvina, M. Feighn and M. Handel, *Solvable subgroups of $\text{Out}(Fn)$ are virtually Abelian*, *Geom. Dedicata* **104** (2004), 71–96.
- [7] M. Bestvina, M. Feighn and M. Handel, *The tits alternative for $\text{Out}(Fn)$. II: A Kolchin type theorem*, *Ann. of Math. (2)* **161** (2005), no. 1, 1–59. [MR 2150382](#)
- [8] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, *Ann. of Math. (2)* **135** (1992), no. 1, 1–51.
- [9] D. J. Collins and E. C. Turner, *Efficient representatives for automorphisms of free products*, *Michigan Math. J.* **41** (1994), no. 3, 443–464. [MR 1297701](#)
- [10] T. Coulbois and A. Hilion, *Ergodic currents dual to a real tree*, preprint; available at [arXiv:1302.3766](#).
- [11] T. Coulbois, A. Hilion and M. Lustig, *Non-unique ergodicity, observers' topology and the dual algebraic lamination for \mathbb{R} -trees*, *Illinois J. Math.* **51** (2007), no. 3, 897–911.
- [12] T. Coulbois, A. Hilion and M. Lustig, *\mathbb{R} -Trees and laminations for free groups I: Algebraic laminations*, *J. Lond. Math. Soc.* **78** (2008), no. 3, 723–736. [MR 2456901](#)
- [13] T. Coulbois, A. Hilion and M. Lustig, *\mathbb{R} -Trees and laminations for free groups II: The dual lamination of an \mathbb{R} -tree*, *J. Lond. Math. Soc.* **78** (2008), no. 3, 737–754.
- [14] T. Coulbois, A. Hilion and M. Lustig, *\mathbb{R} -Trees and laminations for free groups III: Currents and dual \mathbb{R} -tree metrics*, *J. Lond. Math. Soc.* **78** (2008), no. 3, 755–766.
- [15] T. Coulbois, A. Hilion and M. Lustig, *\mathbb{R} -Trees, dual laminations, and compact systems of partial isometries*, *Math. Proc. Cambridge Philos. Soc.* **147** (2009), no. 2, 345–368.
- [16] T. Coulbois, A. Hilion and P. Reynolds, *Indecomposable F_N -trees and minimal laminations*, preprint; available at [arXiv:1110.3506](#).

- [17] M. Culler and J. W. Morgan, *Group actions on \mathbb{R} -trees*, Proc. Lond. Math. Soc. (3) **55** (1987), no. 3, 57–604. [MR 0907233](#)
- [18] M. Foster, *Deformation and rigidity of simplicial group actions on trees*, Geom. Topol. **6** (2002), 219–267.
- [19] S. Francaviglia, *Geodesic currents and length compactness for automorphisms of free groups*, Trans. Amer. Math. Soc. **361** (2009), 161–176.
- [20] S. Francaviglia and A. Martino, *Metric properties of outer space*, Publ. Mat. **55** (2011), 433–473. [MR 2839451](#)
- [21] S. Francaviglia and A. Martino, *The isometry group of outer space*, Adv. Mat. (3-4) **231** (2012), 1940–1973.
- [22] A. Hadari, *Homological shadows of attracting laminations*, preprint; available at [arXiv:1305.1613](#).
- [23] M. Handel and L. Mosher, *Axis in outer space*, Mem. Am. Math. Soc., vol. 1004, 2011.
- [24] I. Kapovich, *Currents on free groups, topological and asymptotic aspects of group theory* (R. Grigorchuk, M. Mihalik, M. Sapir and Z. Sunik, eds.), AMS Contemporary Mathematics Series, vol. 394, 2006, pp. 149–176.
- [25] I. Kapovich, *Clusters, currents and Whitehead’s algorithm*, Exp. Math. **16** (2007), no. 1, 67–76. [MR 2312978](#)
- [26] I. Kapovich and M. Lustig, *The actions of $\text{Out}(F_k)$ on the boundary of outer space and on the space of currents: Minimal sets and equivariant incompatibili*, Ergodic Theory Dynam. Systems **27** (2007), no. 3, 827–847. [MR 2322181](#)
- [27] I. Kapovich and M. Lustig, *Intersection form, laminations and currents on free groups*, Geom. Funct. Anal. **19** (2010), no. 5, 1426–1467. [MR 2585579](#)
- [28] I. Kapovich and M. Lustig, *Invariant laminations for irreducible automorphisms of free groups*, Quarterly J. Math. (2010), to appear; published online Jan 30, 2014, Q. J. Math.
- [29] I. Kapovich and T. Nagnibeda, *The Patterson–Sullivan embedding and minimal volume entropy for outer space*, Geom. Funct. Anal. **17** (2007), no. 4, 1201–1236.
- [30] I. Kapovich and T. Nagnibeda, *Geometric entropy of geodesic currents on free groups*, Dynamical numbers: Interplay between dynamical systems and number theory, Contemporary Mathematics Series, American Mathematical Society, Providence, 2010, pp. 149–176.
- [31] I. Kapovich and T. Nagnibeda, *Subset currents on free groups*, Geom. Dedicata **166**, (2013), 307–348. pp. 307–348.
- [32] S. Meinert, *The Lipschitz metric on deformation spaces of G -trees*, preprint; available at [arXiv:1312.1829](#).
- [33] L. Mosher and C. Pfaff, *Lone Axes in Outer Space*, preprint; available at [arXiv:1311.6855](#).
- [34] M. Sykiotis, *Stable representatives for symmetric automorphisms of groups and the general form of the Scott conjecture*, Trans. Amer. Math. Soc. **356** (2004), no. 6, 2405–2441.
- [35] G. Vincente and G. Levitt, *Deformation spaces of trees*, Groups Geom. Dyn. **1** (2007), 135–181.
- [36] G. Vincente and G. Levitt, *The outer space of a free product*, Proc. Lond. Math. Soc. (3) **94** (2007), 695–714. [MR 2325317](#)

STEFANO FRANCAVIGLIA, DIPARTIMENTO DI MATEMATICA OF THE UNIVERSITY OF BOLOGNA, P.ZZA PORTA S. DONATO 5, 40126 BOLOGNA, ITALY.

E-mail address: stefano.francaviglia@unibo.it

ARMANDO MARTINO, MATHEMATICAL SCIENCES, UNIVERSITY OF SOUTHAMPTON, UNIVERSITY ROAD SOUTHAMPTON, SO17 1BJ, UNITED KINGDOM.

E-mail address: A.Martino@soton.ac.uk