



ANNALES DE L'INSTITUT FOURIER

Stefano FRANCAVIGLIA,
Armando MARTINO & Dionysios SYRIGOS

**On the action of relatively irreducible automorphisms on
their train tracks**

Article à paraître, mis en ligne le 26 janvier 2026, 92 p.

Article mis à disposition par ses auteurs selon les termes de la licence
CREATIVE COMMONS ATTRIBUTION-NoDERIVS (CC-BY-ND) 3.0
 <http://creativecommons.org/licenses/by-nd/3.0/>



Les *Annales de l'Institut Fourier* sont membres du
Centre Mersenne pour l'édition scientifique ouverte
www.centre-mersenne.org e-ISSN : 1777-5310

ON THE ACTION OF RELATIVELY IRREDUCIBLE AUTOMORPHISMS ON THEIR TRAIN TRACKS

by Stefano FRANCAVIGLIA,
Armando MARTINO & Dionysios SYRIGOS (*)

ABSTRACT. — Let G be a group and let \mathcal{G} be a free factor system of G , namely a free splitting of G as $G = G_1 * \dots * G_k * F_r$. In this paper, we study the set of train track points for \mathcal{G} -irreducible automorphisms ϕ with exponential growth. Such set is known to coincide with the minimally displaced set $\text{Min}(\phi)$ of ϕ , in the relative deformation space corresponding to the splitting. The theory of such relative spaces, even if it is more general by its own nature, is crucial to understanding reducible automorphisms of free groups, as any automorphism is relatively irreducible with respect to some free factor system \mathcal{G} .

Our main result is that $\text{Min}(\phi)$ is co-compact, under the action of the cyclic subgroup generated by ϕ .

Along the way we obtain other results that could be of independent interest. For instance, we prove that any point of $\text{Min}(\phi)$ is in uniform distance from $\text{Min}(\phi^{-1})$. We also prove that the action of G on the product of the attracting and the repelling trees for ϕ , is discrete. Finally, we get some fine insight about the local topology of relative outer space.

Some applications of co-compactness are discussed. In particular we generalise a classical result of Bestvina, Feighn and Handel for the centralisers of irreducible automorphisms of free groups, in the more general context of relatively irreducible automorphisms of a free product. From this, we deduce that centralisers of elements of $\text{Out}(F_3)$ are finitely generated, which was previously unknown. Finally, we mention that an immediate corollary of co-compactness is that the set $\text{Min}(\phi)$ is always quasi-isometric to a line.

Keywords: automorphisms of free groups, conjugacy problem.

2020 Mathematics Subject Classification: 20E06, 20E36, 20E08.

(*) The second and third authors were supported by Leverhulme Trust Grant RPG-2018-058 during the work of this paper. The first author was partially funded by INdAM group GNSAGA, by ERC starting grant “Definable Algebraic Topology”, and by the European Union – NextGenerationEU under the National Recovery and Resilience Plan (PNRR) – Mission 4 Education and research – Component 2 From research to business – Investment 1.1 Notice Prin 2022 – DD N. 104 del 2/2/2022, from title “Geometry and topology of manifolds”, proposal code 2022NMPLT8 – CUP J53D23003820001.

RÉSUMÉ. — Soit G un groupe et \mathcal{G} un système de facteurs libres de G , c'est à dire une décomposition de G en produits libres comme $G = G_1 * \dots * G_k * F_r$. Dans cet article, nous étudions l'ensemble de points *train track* pour les automorphismes ϕ qui sont \mathcal{G} -irréductibles et à croissance exponentielle. On sait qu'un tel ensemble coïncide avec l'ensemble $\text{Min}(\phi)$ des points qui sont minimalement déplacés par ϕ dans l'espace de déformation relatif correspondant à la décomposition \mathcal{G} . La théorie de tels espaces relatifs, même si elle est plus générale par nature, est essentielle pour comprendre les automorphismes réductibles des groupes libres, car tout automorphisme est relativement irréductible par rapport à *quelque* système de facteurs libres \mathcal{G} .

Notre résultat principal est que l'action sur $\text{Min}(\phi)$ du groupe cyclique engendré par ϕ est co-compacte.

En passant, nous obtenons d'autres résultats qui pourraient présenter un intérêt indépendant. Par exemple, nous prouvons que tout point de $\text{Min}(\phi)$ est à distance uniforme de $\text{Min}(\phi^{-1})$. Nous prouvons également que l'action de G sur le produit des arbres attractif et répulsif de ϕ est discrète. Enfin, nous obtenons un aperçu précis de la topologie locale de l'outre-espace relatif.

Quelques applications de la co-compacité sont discutées. En particulier, nous généralisons un résultat classique de Bestvina, Feighn et Handel pour les centralisateurs des automorphismes irréductibles du groupe libre, dans le contexte plus général des automorphismes relativement irréductibles de produits libres. On en déduit que les centralisateurs des éléments de $\text{Out}(F_3)$ sont finiment engendrés, ce qui était auparavant inconnu. Finalement, nous mentionnons qu'un corollaire immédiat de la co-compacité est que l'ensemble $\text{Min}(\phi)$ est toujours quasi-isométrique à une droite.

1. Introduction

Overview. Automorphisms of free groups play a central role in geometric group theory. Culler–Vogtmann outer space is one of the main methods that are currently used for the study of automorphisms of free groups. Irreducible automorphisms have been studied the most, as there are available many different tools for them (for instance, train tracks representatives [4]).

More recently, Guirardel and Levitt introduced in [20] the notion of a relative outer space of a group corresponding to a free factor system. These relative spaces have been used for the study of automorphisms of general free products but also for reducible automorphisms of free groups, as any automorphism is relatively irreducible in the appropriate relative outer space. Note that many of the classical tools that are available for irreducible automorphisms are also available in relative outer spaces; for instance, existence of train track representatives in the general context is proved in [11].

In this paper, together with our companion paper [14], we study relatively irreducible automorphisms. In particular, we focus on their minimally displaced set (with respect to the Lipschitz metric) which (by [11]) can be seen as the set of train track points (and train track properties will be crucially used in our arguments).

Main results of the paper. (See Section 2 for a more complete explanation of our terminology.)

Let G be a group, with a free factor system \mathcal{G} . Let $\text{Aut}(\mathcal{G})$ be the group of automorphisms of G which preserve \mathcal{G} (up to conjugacy) and $\text{Out}(\mathcal{G}) = \text{Aut}(\mathcal{G})/\text{Inn}(G)$. Denote by $[\phi] \in \text{Out}(\mathcal{G})$ the outer class of $\phi \in \text{Aut}(\mathcal{G})$.

We then let $\mathcal{O}(\mathcal{G})$ be the relative outer space corresponding to \mathcal{G} ; that is, the space of minimal edge-free actions of G on metric trees with stabilisers giving rise to \mathcal{G} , up to equivariant isometry. We denote by $\mathcal{O}_1(\mathcal{G})$ the co-volume one subspace of $\mathcal{O}(\mathcal{G})$. Both $\text{Aut}(\mathcal{G})$ and $\text{Out}(\mathcal{G})$ act on $\mathcal{O}(\mathcal{G})$ and $\mathcal{O}_1(\mathcal{G})$ by twisting the action. Inner automorphisms act trivially, hence the action of $\phi \in \text{Aut}(\mathcal{G})$ on $\mathcal{O}(\mathcal{G})$ depends only on $[\phi]$.

We denote by $\text{Min}(\phi)$ the set of points in $\mathcal{O}(\mathcal{G})$ which are minimally displaced by ϕ , with respect to the Lipschitz metric, and set $\text{Min}_1(\phi) = \text{Min}(\phi) \cap \mathcal{O}_1(\mathcal{G})$.

By [11], in the case where $[\phi]$ is irreducible, this is exactly the set of points which support train track maps representing $[\phi]$ (this is explained in more detail in Section 2 and Theorem 2.47).

The main result of this paper is that if $[\phi]$ is \mathcal{G} -irreducible with exponential growth, then $\text{Min}_1(\phi)$ is co-compact, under the action of the cyclic group generated by ϕ .

We first prove, in Section 7.2, our main result under the extra hypothesis that $[\phi]$ is primitive (see Theorem 7.11), i.e. it has a train track representative with primitive transition matrix. Then, in Section 7.3 we drop primitivity condition, proving the following result.

THEOREM (Theorem 7.20). — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) > 1$ (that is, a relatively irreducible automorphism with exponential growth). Then the action of $\langle \phi \rangle$ on $\text{Min}_1(\phi) = \text{Min}(\phi) \cap \mathcal{O}_1$ is co-compact.*

We notice that co-compactness was already known for genuine irreducible automorphisms of a free group (see [22], for the original proof of Handel and Mosher, and [15], for a recent elementary proof which was given by the authors). The main difference between the classical case and our general case is that the Culler–Vogtmann outer space of a free group has a locally finite simplicial structure, while our general relative deformation spaces are not even locally compact.

Remark. — In our companion paper [14], we prove that the minimally displaced set of an irreducible automorphism of exponential growth is locally finite. It may seem quite intuitive to the reader that as $\text{Min}(\phi)$ is

locally finite, its co-compactness is equivalent to the existence of a fundamental domain contained in the union of finitely many simplices.

In fact, our strategy, and the thrust of this paper, is to show that the action is co-bounded and then deduce co-compactness from the local finiteness result. This seems intuitive, but presents some challenges since our main tool is the Lipschitz metric which is an asymmetric metric and whose properties can sometimes fail to be well-behaved; for instance, a general relative outer space is usually a locally infinite space and so is not locally compact.

For the experts, one important question is — given we already have that $\text{Min}(\phi)$ is locally finite — why is this current paper so long? There are a couple of reasons. The first, and most important, is that the proofs of this fact for Culler–Vogtmann space divide into two main arguments according to whether the irreducible automorphism is nongeometric or geometric. The second case corresponds precisely to the existence of a closed periodic Nielsen path, and then it is a well-known result of [4] that the corresponding automorphism is induced from a surface homeomorphism. The arguments then proceed by appeals to surface theory in the latter case and to other arguments in the nongeometric case (which are not valid for the geometric case). A direct ancestor of our result for Culler–Vogtmann space can be found in [22], where they prove that the axis bundle is co-compact (amongst other results). As they note:

We shall do this only in the case where ϕ is nongeometric, meaning that it does not arise from a homeomorphism of a compact surface with boundary. The geometric case, while conceptually much simpler and more well understood, has some peculiarities whose inclusion in our theory would overburden an already well-laden paper.

However, in the free product case the appeal to surface theory is more delicate and less obviously valid. That is, realising a homotopy equivalence of a graph as a homeomorphism of a surface is a core part of the theory for free group automorphisms going back to [4], but it is not just that this analogue is absent in the free product or relative case, there is good reason to think that it isn't entirely valid. For instance, any irreducible automorphism of exponential growth of a free group has a train track representative whose transition matrix is primitive. (In [2] it is proved that any irreducible automorphism has a locally connected Whitehead graph — see also Section 5. In [8] this is called “weakly clean” and in that paper, Proposition B.2, it is shown that this implies clean, which means having a

primitive transition matrix. See also [28, Remark 2.10] for a discussion of this.)

But this is no longer true for free products, since one can write examples of train track representative of irreducible automorphisms which are not fully irreducible (that is, some power is reducible) and whose transition matrices are not primitive (this cannot happen in the absolute case).

It is possible that one can resolve these issues and introduce surfaces into the free product situation — using something like the improved relative train tracks of [3], especially when the underlying group is free, and the given automorphism is relatively irreducible with respect to a free factor system — which one can think of as the main case of interest for the free product theory. Taking that route would simplify our Section 6 a little, which is a generalisation of similar results in [2] and [7], but we have chosen not to do that since it seems to us that fully taking into account the subtleties of *relative surfaces* would add length and complexity in a different way. We would reiterate here that this process has not formally been done, and we feel that the subtleties and differences between free groups and free products require some caution in simply hand-waving through techniques which may not be entirely valid.

Instead, we actually deal with the geometric and nongeometric cases (where here the distinction being made is whether or not there is a closed periodic Nielsen path) at the same time. This is somewhat new, conceptually, and seems appropriate for the general free product case. (Also, Theorem 6.17 seems interesting and important to us, and proving it for any irreducible automorphism without adding technical conditions seems worthwhile. One of our goals was to prove [7, Lemma 2.13] but without the “no twinned subgroups” hypothesis, and in this we were successful.)

The second reason for the length of this current paper is that we do not pass to powers of the automorphism to make hypotheses and arguments simpler. For instance, we do not assume that our automorphisms are fully irreducible. (Proofs in the fully irreducible case are always easier, but the statements always seem to hold more generally.) This may seem a very minor difference, and in some situations, for instance in [7] where the goal is to prove relative hyperbolicity, passing to a power is a tame procedure. However, for more algebraic applications, like looking at centralisers, passing to powers is not benign. (Knowing something about the centraliser of a power does not yield strong information about the original one.) To take another example: in [5] there is a solution of the conjugacy problem for Dehn twists and in [26], this is extended to roots of Dehn twists. However,

one notes that the latter paper is longer than the former because the additional complexities of looking at roots are considerable. It is probably true that our arguments can be shortened by passing to suitable powers, but this would certainly invalidate our application and would definitely be a weaker result with less scope for further applications.

Ultimately, the argument proceeds as one would expect, but there are many places where the intuition one would get from the free group situation would be slightly incorrect. Thus some of our text is expository since experts familiar with the free group setting would tend to assume that many things which are true for free groups are also true for free products, and this is not always the case.

In the process of the proof of our main result, we obtain some new results that could be of independent interest. For instance, we show that $\text{Min}(\phi)$ is quasi-isometric to a line.

COROLLARY (Corollary 7.21). — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) > 1$. Then $\text{Min}_1(\phi)$, equipped with the symmetric Lipschitz metric, is quasi-isometric to a line.*

Remark. — Note that this is also true with respect to the path Euclidean metric, since the Svarc–Milnor Lemma also applies for that metric.

Moreover, we show that $\text{Min}(\phi)$ is undistorted in the relative outer space (Theorem 4.25).

We also show that for any \mathcal{G} -irreducible $[\phi]$ with exponential growth, any point of $\text{Min}(\phi)$ is at uniform distance from $\text{Min}(\phi^{-1})$.

THEOREM (Theorem 3.5). — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible, with $\lambda(\phi) > 1$. Then there is a D -neighbourhood (with respect to the symmetrised stretching factor, see Section 2.5) of $\text{Min}_1(\phi)$ containing $\text{Min}_1(\phi^{-1})$.*

More precisely, for any L there is a constant D (depending only on $[\phi]$ and L) such that for any volume-1 point, X with $\lambda_\phi(X) \leq L$, there is a volume-1 point, $Y \in \text{Min}(\phi^{-1})$ such that $\Lambda(X, Y)\Lambda(Y, X) < D$. In particular, for any $X \in \text{Min}(\phi)$ there is $Y \in \text{Min}(\phi^{-1})$ such that $\Lambda(X, Y)\Lambda(Y, X) < D$.

Another interesting result is the following. For a \mathcal{G} -irreducible automorphism ϕ with exponential growth, we can define the attracting and repelling trees (starting from a train track point X). We prove a discreteness result for the product of the limit trees.

THEOREM (Theorem 6.17). — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) > 1$ (that is, a relatively irreducible automorphism with exponential growth). Let $X \in \text{Min}(\phi)$ and $Y \in \text{Min}(\phi^{-1})$, and denote by $X_{+\infty}$*

and $Y_{-\infty}$ the corresponding attracting tree and repelling tree for ϕ , respectively (Definition 2.63). Then there exists an $\epsilon > 0$ such that for all $g \in G$, either

- $\ell_{X_{+\infty}}(g) = \ell_{Y_{-\infty}}(g) = 0$ or,
- $\max\{\ell_{X_{+\infty}}(g), \ell_{Y_{-\infty}}(g)\} \geq \epsilon$.

Applications. As a first application of co-compactness, we get Theorem 8.1, which describes the structure of centralisers of relatively irreducible automorphisms, in the spirit of the classical result of Bestvina, Feighn and Handel for irreducible automorphisms of free groups (see [2]).

Also, our results for general deformation spaces have applications to classical cases. An example is the following fact for centralisers of elements in $\text{Out}(F_3)$, which was previously unknown.

THEOREM (Theorem 8.2). — *Centralisers of elements in $\text{Out}(F_3)$ are finitely generated.*

Strategy of the proof of Theorem 7.20. Associated to any \mathcal{G} -irreducible automorphism class of exponential growth $[\phi]$ we can define the attracting (or stable) tree, $X_{+\infty}$, which is the forward limit of some point, $X_{+\infty} = \lim_{n \rightarrow \infty} \frac{X\phi^n}{\lambda(\phi)^n}$, where $X \in \text{Min}(\phi)$. Note that since X is minimally displaced, we get that the Lipschitz distance from X to $\frac{X\phi^n}{\lambda(\phi)^n}$ is also 1. In fact, the Lipschitz distance from X to $X_{+\infty}$ is again, 1, and this same calculation holds for any minimally displaced point and its forward limit. Thus any point in $\text{Min}(\phi)$ is uniformly close to its forward limit.

North-south dynamics then tells us that all points have the same forward limit, up to positive scaling constants, except for the *repelling* (or *unstable*) tree.

If the action of $\langle \phi \rangle$ on $\text{Min}_1(\phi)$ is not co-compact, we can then find a sequence of points $X_n \in \text{Min}(\phi)$ whose limit $T = \lim_{n \rightarrow \infty} X_n$ is very far (at infinite distance) from both the attracting and repelling trees for ϕ . This leads to the following contradiction:

- T is at finite distance from its forward limit $\lim_{n \rightarrow \infty} \frac{T\phi^n}{\lambda(\phi)^n}$, since T is a limit of minimally displaced points, each of which is uniformly close to its forward limit;
- T is at infinite distance from its forward limit $\lim_{n \rightarrow \infty} \frac{T\phi^n}{\lambda(\phi)^n}$, since the latter is $cX_{+\infty}$ for some $c > 0$, by north-south dynamics.

Implementing this argument requires a careful verification that one's geometric intuition concerning limits and distance are correct. We do this, by proving results that we believe are of independent interest along the

way. In more detail, we start with a \mathcal{G} -irreducible automorphism class $[\phi]$ with exponential growth rate $\lambda(\phi) > 1$.

- (1) We fix a “basepoint”, $X \in \text{Min}_1(\phi)$ and define the attracting tree, $X_{+\infty} = \lim_n \frac{X\phi^n}{\lambda(\phi)^n}$, which exists due to train track properties, Lemma 2.62.
- (2) We argue by contradiction, and suppose that $\text{Min}_1(\phi)/\langle\phi\rangle$ is not compact.
- (3) We thus produce a sequence — justified in Theorem 4.26 — $X_n \in \text{Min}_1(\phi)$, such that the distance from X to the ϕ -orbit of X_n tends to infinity. The distance we use here is the Lipschitz distance (where we can use either the symmetric or non-symmetric ones, since they are equivalent on the thick part, and any point of $\text{Min}_1(\phi)$ must be thick).
- (4) By replacing each X_n with a suitable element of its ϕ -orbit, we can assume that $1 \leq \Lambda(X_n, X_{+\infty}) \leq \lambda(\phi)$ and $\Lambda(X, X_n)$ is unbounded. (In fact, Theorem 4.26 has a long list of equivalent statements of co-compactness that includes this one.)
- (5) As $\overline{\mathcal{PO}(\mathcal{G})}$ is compact, we may find constants μ_n and a subsequence of X_n such that $\lim_n \frac{X_n}{\mu_n} \rightarrow T$ (this is convergence as length functions, and occurs in $\overline{\mathcal{O}(\mathcal{G})}$; in case G is not countable we can use ultralimits instead of classical limits).
- (6) Since T is the limit of points displaced by $\lambda(\phi)$, T itself is displaced by at most $\lambda(\phi)$ under ϕ , Lemma 7.5.
- (7) We then argue, in Proposition 7.8, that $\Lambda(T, X_{+\infty}) = \infty$, which in particular implies that T is not in the same homothety class as $X_{+\infty}$.
- (8) Symmetrically, we argue T is not in the same homothety class as the repelling tree. However, since many aspects of the theory are not symmetrical, this requires two important ingredients:
 - (i) Theorem 3.5 shows that there is a uniform distance between $\text{Min}_1(\phi)$ and $\text{Min}_1(\phi^{-1})$. That is, one is contained in a Lipschitz neighbourhood of the other, and so T is also a limit of points which are minimally displaced by ϕ^{-1} , even though $\text{Min}_1(\phi)$ and $\text{Min}_1(\phi^{-1})$ are different. (More precisely, T is bi-Lipschitz equivalent to a limit of such points.)
 - (ii) Theorem 6.17 shows that if we have a bound on the Lipschitz distance to the attracting tree, we also get a bound on the Lipschitz distance to the (in fact, any) repelling tree. Thus T is also a limit (or bi-Lipschitz equivalent to a limit) of points,

minimally displaced by ϕ^{-1} , whose distance to the repelling tree is bounded.

- (iii) This is enough symmetry to conclude — Corollary 7.10 — that T is not in the same homothety class as the repelling tree.
- (9) We then apply north-south dynamics to T (we need to know that T is not in the same homothety class as the repelling tree for this to work), which combined with the previous results says that $\lim_n \frac{T\phi^n}{\lambda(\phi)^n}$ is both at finite distance from T , and in the same homothety class as $X_{+\infty}$, which is a contradiction. Hence this contradiction implies that $\text{Min}_1(\phi)/\langle\phi\rangle$ is compact.
- (10) As north-south dynamics are not available for general irreducible automorphisms, in Section 7.3 we give an additional argument that is needed in order to deduce the co-compactness of a general irreducible automorphism, deducing it from the case of primitive irreducible automorphisms, where north-south dynamics are known to hold.

Some of the results stated here are dependent on others in unexpected ways. For instance, the equivalent formulations of co-compactness, Theorem 4.26, relies on the fact that $\text{Min}(\phi)$ is uniformly close to $\text{Min}(\phi^{-1})$, Theorem 3.5.

The organisation of the paper is as follows.

- Section 2 sets up terminology and recalls known results. While this is largely known to experts, we do have some minor proofs which appear to be new (Lemmas 2.58 and 2.59).
- Section 3 is a fairly short section showing that the minimally displaced for ϕ is uniformly close to that for ϕ^{-1} , using results from [12] and [13].
- Section 4 is devoted to proving the equivalent conditions for co-compactness, and also contains a discussion of the topologies on our deformation spaces.
- Section 5 is a short discussion on the north-south dynamics for primitive irreducible automorphisms. The material here is largely a verification, in this context, of results that are known in classical and/or less general cases.
- Section 6 is the most technical section. The goal of this section is the final “discreteness” Theorem 6.17. The proofs of this section are not used anywhere else, just the final result.

- Section 7 pulls everything together to prove co-compactness, first for the primitive irreducible case and then for the general irreducible case.
- Section 8 is devoted to applications, showing in particular that centralisers in $\text{Out}(F_3)$ are finitely generated.

Acknowledgments

We are grateful to the referee for their careful and thorough feedback along with many very useful comments.

2. Terminology and preliminaries

2.1. Relative outer space $\mathcal{O}(\mathcal{G})$

Let G be a group which decomposes as a finite free product

$$G = G_1 * \cdots * G_k * F_r$$

where F_r is the free group on $r \geq 0$ generators. We impose no restriction on the G_i 's (in particular we do not assume that the G_i 's are freely indecomposable nor non-cyclic, nor finitely generated, nor countable...). Any such free product decomposition is commonly referred to as a *free factor system* of G . More precisely:

Notation 2.1. — A free factor system of G is a pair

$$\mathcal{G} = (\{G_1, \dots, G_k\}, r)$$

such that $G = G_1 * \cdots * G_k * F_r$. We define the rank of \mathcal{G} as $\text{rank}(\mathcal{G}) = k + r$. With $[\mathcal{G}]$ we denote the set of conjugacy classes of the G_i 's, that is $[\mathcal{G}] = \{[G_1], \dots, [G_k]\}$. If $\mathcal{G}' = (\{G'_1, \dots, G'_s\}, m)$ is another free factor system, we say that \mathcal{G} is bigger than \mathcal{G}' if for any i there is j such that G'_i is a subgroup of some conjugate of G_j .

DEFINITION 2.2. — Let G be a group.

- A *G-tree* is a tree T together with an action of G . If the tree is simplicial (resp. metric), then the action is supposed to be simplicial (resp. isometric).
- A *G-tree* T is called *minimal*, if it has no proper G -invariant subtree.

- The action of G on a G -tree is called marking (and a marked tree is a tree equipped with a G -action).
- If T is a minimal simplicial metric G -tree, we denote by $\text{vol}(T)$ the co-volume of T , namely the sum of lengths of edges of the quotient graph $G \setminus T$ (a priori this number could be infinite).

In this paper the G -action on a G -tree will always be a left-action.

DEFINITION 2.3. — Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . A simplicial G -tree is called (simplicial) \mathcal{G} -tree, if:

- T has trivial edge stabilisers (that is to say, no $1 \neq g \in G$ pointwise fixes an edge), and no inversions (that is to say, no $g \in G$ maps an edge to its inverse).
- The non-trivial vertex stabilisers of T are exactly the conjugacy classes that are contained in $[\mathcal{G}]$. More precisely, for every i there is a unique vertex v_i with stabiliser G_i . The vertices with non-trivial stabiliser will be called non-free vertices; the other vertices will be called free vertices. We use the notation $G_{v_i} = \text{Stab}_G(v_i)$, and we often refer to factor groups G_i 's as vertex groups (since there are finitely many G_i 's and since F_r has finite rank, the volume of \mathcal{G} -trees is a finite number).

DEFINITION 2.4. — Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . The relative outer space of \mathcal{G} — denoted by $\mathcal{O}(\mathcal{G})$ — is the set of equivalence classes of minimal, simplicial, metric \mathcal{G} -trees, with no redundant vertices (i.e. any free vertex has valence at least 3), where the equivalence relation is given by G -equivariant isometries. We denote by $\mathcal{O}_1(\mathcal{G})$, the co-volume-1 subset of $\mathcal{O}(\mathcal{G})$.

There is a natural action of \mathbb{R}^+ on $\mathcal{O}(\mathcal{G})$ given by $a: T \mapsto aT$ where aT denotes the same marked tree as T , but with the metric scaled by $a > 0$. We denote by $\mathbb{P}\mathcal{O}(\mathcal{G})$ the projectivised relative outer space (that is, the quotient of $\mathcal{O}(\mathcal{G})$ by the \mathbb{R}^+ -action).

2.2. Simplicial structure of $\mathcal{O}(\mathcal{G})$

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G , and consider $X \in \mathcal{O}(\mathcal{G})$. The (open) simplex $\Delta(X)$ is the set of points of $\mathcal{O}(\mathcal{G})$ which are obtained from X by just changing the lengths of (orbits of) edges in such a way that any edge has positive length. Thus $\Delta(X)$ is parameterised by the positive cone of \mathbb{R}^n , where n is the number of orbits of edges

in X . Note that the positive cone of \mathbb{R}^n can be naturally identified with an open n -simplex.

If we work in $\mathcal{O}_1(\mathcal{G})$, then $\Delta(X)$ determines a standard open $(n-1)$ -simplex $\Delta(X)_1 = \Delta(X) \cap \mathcal{O}_1(\mathcal{G})$. We will often omit the subscript “1” and write just Δ or $\Delta(X)$ when it is clear from the context in which space we are working.

Remark 2.5. — The \mathbb{R}^+ action plus the parameterisations of $\Delta(X)$ and $\Delta(X)_1$ by convex subsets of \mathbb{R}^n , allow us to define Euclidean segments between pair of points X, Y in the same simplex by the usual formula $tX + (1-t)Y$.

Remark 2.6. — So far we have not mentioned topology, but all of the topologies we will consider induce the standard Euclidean topology on each simplex of $\mathcal{O}_1(\mathcal{G})$.

Simplicial faces of simplices of $\mathcal{O}(\mathcal{G})$ do not always live inside $\mathcal{O}(\mathcal{G})$, so the space is not a simplicial complex. Any face of a simplex $\Delta = \Delta(X)$ in $\mathcal{O}(\mathcal{G})$ is induced by collapsing a G -invariant sub-forest of X . Such a collapse produces a simplicial G -tree Y with trivial edge stabilisers, and there are two cases: either Y is a \mathcal{G} -tree or not (i.e. vertex stabilisers are not in $[\mathcal{G}]$). In the first case we say that $\Delta(Y)$ is a *finitary face*, in the second that it is a *face at infinity*. We notice that faces at infinity correspond to free factor systems strictly bigger than \mathcal{G} .

Remark 2.7. — Given $T \in \mathcal{O}(\mathcal{G})$, the quotient graph $G \setminus T$ comes endowed with a structure of graph of groups, and the choice of a marking of T corresponds to the choice of an isomorphism from G to $\pi_1(G \setminus T)$ (where the fundamental group is taken in the sense of graph of groups). The equivalence relation given by equivariant isometries of G -trees translates to a notion of equivalence of marked graphs, which is the usual one for the reader used to Teichmüller theory or classical Culler–Vogtmann outer space CV_n .

In this paper we will use only the tree-viewpoint, but in some case graphs are easier to visualise. For instance, one can easily see with graphs that if there is at least one G_i which is infinite, then the simplicial structure of $\mathcal{O}(\mathcal{G})$ is not locally finite.

Example 2.8. — Consider the simple case $G = G_1 * \mathbb{Z}$ where G_1 is an infinite group. The simplex corresponding to a graph of groups formed by a circle with a unique non-free vertex is a finitary face of infinitely



Figure 2.1. Graphs corresponding to open simplices.

many simplices corresponding to a graph formed by a circle with a segment attached, ending with the unique non-free vertex (See Figure 2.1).

This is because for any $g \in G_1$, if $\mathbb{Z} = \langle a \rangle$, then we can define an isomorphism $\phi_g: G \rightarrow G$ which is the identity on G_1 and maps a to ga . It is readily checked that all markings induced by all ϕ_g 's on the left-side graph are equivalent, while they are not equivalent on right-side graphs.

2.3. Action of the automorphism groups

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G .

DEFINITION 2.9. — *The group of automorphisms of G that preserve the set $[\mathcal{G}]$ (that is to say, $[f(G_i)] \in [\mathcal{G}]$ for all i) is denoted by $\text{Aut}(G; \mathcal{G})$, or simply $\text{Aut}(\mathcal{G})$. We set $\text{Out}(\mathcal{G}) = \text{Out}(G; \mathcal{G}) = \text{Aut}(G; \mathcal{G}) / \text{Inn}(G)$.*

There is a natural right action of $\text{Aut}(\mathcal{G})$ on $\mathcal{O}(\mathcal{G})$, given by twisting the marking. More specifically, given $T \in \mathcal{O}(\mathcal{G})$ and $\phi \in \text{Aut}(\mathcal{G})$, we define the point $T\phi$ as the same metric tree as T , but the G -action on $T\phi$ is given by

$$x \longmapsto \phi(g) \cdot x,$$

where \cdot denotes the G -actions on T . In terms of marked graphs this corresponds to precomposing the marking with ϕ .

If $\alpha \in \text{Inn}(G)$ and $T \in \mathcal{O}(\mathcal{G})$, then it is easy to see that there is a G -equivariant isometry between $T\alpha$ and T , i.e. they are equal as objects of $\mathcal{O}(\mathcal{G})$. It follows that $\text{Inn}(G)$ acts trivially on $\mathcal{O}(\mathcal{G})$. Thus there is an induced action of $\text{Out}(\mathcal{G})$ on $\mathcal{O}(\mathcal{G})$. Since the action of ϕ on $\mathcal{O}(\mathcal{G})$ depends only on $[\phi]$, we write simply $T\phi$ to denote both actions. Moreover, the action preserves the co-volume of trees, so we get induced actions on the co-volume-1 set $\mathcal{O}_1(\mathcal{G})$.

Remark 2.10. — Since F_r has finite rank, we have finitely many topological types of graphs $G \setminus T$, as T varies in $\mathcal{O}(\mathcal{G})$. As a consequence, there are finitely many orbits of simplices under the action of $\text{Out}(\mathcal{G})$.

2.4. Translation lengths, thickness, and boundary points

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G .

For any metric G -tree T (not necessarily in $\mathcal{O}(\mathcal{G})$) and for any $g \in G$, we define the translation length of g in T , which actually depends only on the conjugacy class $[g]$, by

$$\ell_T(g) = \ell_T([g]) = \inf\{d_T(x, gx) : x \in T\}.$$

It is well-known (see [6]) that the infimum is achieved by some $x \in T$. We have a dichotomy of elements in G . If $\ell_T(g) > 0$, then g is called *hyperbolic* (in T) or T -hyperbolic. In this case, the set of points achieving the minimum above is a line in T , on which g acts by translations by $\ell_T(g)$, and it is called the *axis* of g in T . Otherwise, g is called *elliptic* (in T) or T -elliptic.

If $T \in \mathcal{O}(\mathcal{G})$ then elliptic elements are exactly those belonging to some vertex group, and therefore hyperbolicity of elements does not depend on the tree $T \in \mathcal{O}(\mathcal{G})$ but only on $[\mathcal{G}]$. We denote the set of hyperbolic elements of \mathcal{G} by $\text{Hyp}(\mathcal{G})$, and we refer to them as \mathcal{G} -hyperbolic elements. Other elements are called \mathcal{G} -elliptic.

Let \mathcal{C} be the set of conjugacy classes of elements in G . We can define a map

$$\begin{aligned} L: \mathcal{O}(\mathcal{G}) &\longrightarrow \mathbb{R}^{\mathcal{C}}, \\ L(T) &= (\ell_T(c))_{c \in \mathcal{C}}. \end{aligned}$$

It is proved by Culler and Morgan in [6] that in our context that map is injective. Moreover, it induces an injective map $L: \mathbb{PO}(\mathcal{G}) \rightarrow \mathbb{PR}^{\mathcal{C}}$.

DEFINITION 2.11. — *The length function topology on $\mathcal{O}(\mathcal{G})$ and $\mathcal{O}_1(\mathcal{G})$ is that induced by the immersion $L: \mathcal{O}(\mathcal{G}) \rightarrow \mathbb{R}^{\mathcal{C}}$.*

Remark 2.12. — With respect to the length function topology, $T_n \rightarrow T$ if and only if for any $g \in G$ we have $\ell_{T_n}(g) \rightarrow \ell_T(g)$.

It is easy to check that the length function topology is Hausdorff, and agrees on each simplex with the Euclidean one. We alert the reader that the choice of the topology on $\mathbb{PO}(\mathcal{G})$ involves some subtlety, that will be discussed in Section 4.1. So far, the length function topology is the unique topology we have defined.

DEFINITION 2.13. — *We will denote by $\overline{\mathcal{O}(\mathcal{G})}$ the closure of $\mathcal{O}(\mathcal{G})$ as a sub-space of $\mathbb{R}^{\mathcal{C}}$, and by $\overline{\mathbb{PO}(\mathcal{G})}$ the closure of $\mathbb{PO}(\mathcal{G})$ as a sub-space of $\mathbb{PR}^{\mathcal{C}}$.*

It is known that $\overline{\mathcal{PO}(\mathcal{G})}$ is a compact space (see [6] for the classical case and [24] for the relative one). Moreover, there is a more detailed description of $\overline{\mathcal{PO}(\mathcal{G})}$ in terms of very small trees as follows.

DEFINITION 2.14. — *Let T be a metric G -tree such that every factor G_i fixes a unique point of T . Then T is called small if arc stabilisers in T are either trivial, or cyclic and not contained in any conjugate of some G_i . T is called very small if it is small, non-trivial arc stabilisers in T are closed under taking roots, and tripod stabilisers in T are trivial.*

THEOREM 2.15 (Horbez, [24]). — *Let \mathcal{G} be a free factor system of a countable group G , and let $\overline{\mathcal{O}(\mathcal{G})}$ be the corresponding relative outer space. Then $\overline{\mathcal{PO}(\mathcal{G})}$ is the space of projective length functions of minimal, very small trees (with respect to the free factor system \mathcal{G}).*

Remark 2.16. — In our Arc Stabiliser Lemma 2.58, we prove (for completeness) that non-trivial arc stabilisers in $\overline{\mathcal{O}(\mathcal{G})}$ are \mathcal{G} -hyperbolic, without assuming the group is countable.

In analogy with Teichmüller space, we can define thick and thin parts of outer spaces.

DEFINITION 2.17. — *For any $\epsilon > 0$ we define the thick part $\mathcal{O}(\mathcal{G}, \epsilon)$ as the set of all $T \in \overline{\mathcal{O}(\mathcal{G})}$ such that all elements in $\text{Hyp}(\mathcal{G})$ have translation length more than $\epsilon \text{vol}(T)$. Namely, $T \in \mathcal{O}(\mathcal{G}, \epsilon)$ if for all $g \in \text{Hyp}(\mathcal{G})$ we have $\ell_T(g)/\text{vol}(T) > \epsilon$. We denote also by $\mathcal{O}_1(\mathcal{G}, \epsilon) = \mathcal{O}_1(\mathcal{G}) \cap \mathcal{O}(\mathcal{G}, \epsilon)$, the thick part of $\mathcal{O}_1(\mathcal{G})$. We say that ϵ is the level of thickness (or simply the thickness) of $\mathcal{O}(\mathcal{G}, \epsilon)$.*

Remark 2.18. — It is immediate to see that for any simplex Δ , the closure of $\Delta \cap \mathcal{O}_1(\mathcal{G}, \epsilon)$ is compact. Hence, since we have finitely many $\text{Out}(\mathcal{G})$ -orbits of simplices, for any $\epsilon > 0$, the quotient space $\mathcal{O}_1(\mathcal{G}, \epsilon)/\text{Out}(\mathcal{G})$ is compact.

2.5. Stretching factors and Lipschitz metrics

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . For any $T \in \overline{\mathcal{O}(\mathcal{G})}$ and $S \in \overline{\mathcal{O}(\mathcal{G})}$, we define the (right) stretching factor as:

$$\Lambda(T, S) = \sup_{g \in \text{Hyp}(\mathcal{G})} \frac{\ell_S(g)}{\ell_T(g)}.$$

It is immediate from the definition that Λ is right-multiplicative and left-anti-multiplicative:

$$\lambda\Lambda(T, S) = \Lambda(T, \lambda S) = \Lambda\left(\frac{1}{\lambda}T, S\right).$$

The stretching factor is not symmetric, and in general fails to be quasi-symmetric. However, if it is restricted on any thick part $\mathcal{O}(\mathcal{G}, \epsilon)$ of $\mathcal{O}(\mathcal{G})$, it is quasi-symmetric.

THEOREM 2.19 ([33]). — *For any $\epsilon > 0$, there exists a constant $C = C(\epsilon)$ such that for all $X, Y \in \mathcal{O}_1(\mathcal{G}, \epsilon)$, we have*

$$\Lambda(X, Y) \leq \Lambda(Y, X)^C.$$

The stretching factor can be viewed as a multiplicative, non-symmetric, pseudo-metric. It comes with its left avatar and symmetrised version. All of them are generically referred to as “Lipschitz metrics” on $\mathcal{O}(\mathcal{G})$, and have been extensively studied, for instance in [9, 10, 11]. We list some of its basic properties.

THEOREM 2.20 ([11]). — *Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G , and let $\mathcal{O}(\mathcal{G})$ be its outer space. Then:*

- (1) *Λ is an asymmetric multiplicative pseudo-metric on $\mathcal{O}(\mathcal{G})$, which restricts to an asymmetric multiplicative metric on $\mathcal{O}_1(\mathcal{G})$:*
 - *for all $T \in \mathcal{O}(\mathcal{G})$, $\Lambda(T, T) = 1$;*
 - *for $T, S, Q \in \mathcal{O}(\mathcal{G})$, $\Lambda(T, S) \leq \Lambda(T, Q)\Lambda(Q, S)$;*
 - *for $T, S \in \mathcal{O}_1(\mathcal{G})$, we have $\Lambda(T, S) \geq 1$, and $\Lambda(T, S) = 1$ if and only if $T = S$.*
- (2) *For every $T \in \mathcal{O}(\mathcal{G})$ and $S \in \overline{\mathcal{O}(\mathcal{G})}$, there is a \mathcal{G} -hyperbolic element g_0 so that $\Lambda(T, S) = \frac{\ell_S(g_0)}{\ell_T(g_0)}$.*
- (3) *$\text{Out}(\mathcal{G})$ acts by Λ -isometries on $\mathcal{O}(\mathcal{G})$.*
- (4) *The symmetrised stretching factor $D(S, T) = \Lambda(S, T)\Lambda(T, S)$ satisfies the following: for all $T, S \in \mathcal{O}(\mathcal{G})$,*
 - *$D(T, S) \geq 1$, and $D(T, S) = 1$ if and only if there is $\lambda > 0$ such that $T = \lambda S$;*
 - *$D(T, S) = D(S, T)$;*
 - *for any $Q \in \mathcal{O}(\mathcal{G})$, $D(T, S) \leq D(T, Q)D(Q, S)$.*

In particular the function $\log D$ is a pseudo-metric on $\mathcal{O}(\mathcal{G})$ that restricts to a genuine metric on $\mathcal{O}_1(\mathcal{G})$.

Any of these metrics induces a topology on $\mathcal{O}(\mathcal{G})$, $\mathcal{O}_1(\mathcal{G})$, and on $\mathbb{P}(\mathcal{O}(G))$ as a quotient of $\mathcal{O}(\mathcal{G})$, whose relation with length function topology will

be discussed in Section 4.1. It is however readily checked that all such topologies induce the Euclidean one on each simplex of $\mathcal{O}_1(\mathcal{G})$.

2.6. Optimal maps and gate structures

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G .

DEFINITION 2.21. — Let $X \in \mathcal{O}(G)$, $Y \in \overline{\mathcal{O}(\mathcal{G})}$. A Lipschitz continuous and G -equivariant map $f: X \rightarrow Y$ is called an \mathcal{O} -map. $\text{Lip}(f)$ denotes the best Lipschitz constant for f .

The name “Lipschitz metric” when referring to stretching factor is motivated by the fact that $\Lambda(X, Y)$ can be viewed as the best Lipschitz constant of equivariant maps from X to Y .

THEOREM 2.22 ([11, 12]). — For any $X, Y \in \mathcal{O}(G)$ we have

$$\Lambda(X, Y) = \inf_f \text{Lip}(f)$$

where f runs over the set of \mathcal{O} -maps from X to Y . Moreover there is at least an \mathcal{O} -map $f: X \rightarrow Y$ realising the stretching factor, that is, such that $\Lambda(X, Y) = \text{Lip}(f)$.

DEFINITION 2.23. — Let $X \in \mathcal{O}(\mathcal{G})$, $Y \in \overline{\mathcal{O}(\mathcal{G})}$. An \mathcal{O} -map $f: X \rightarrow Y$ is called straight if it is linear on edges, i.e. for any edge e of X , there is non-negative number $\lambda_e(f)$ so that the edge e is uniformly stretched by $\lambda_e(f)$. Given a straight map, the tension graph of f is the set of maximally stretched edges:

$$X_{\max}(f) = \{\text{edges } e : \lambda_e(f) = \text{Lip}(f)\}.$$

DEFINITION 2.24. — Let $X \in \mathcal{O}(\mathcal{G})$, and let v be a vertex of X . A turn of X at v is the G_v -orbit of an unoriented pair of edges based at v .

DEFINITION 2.25. — A gate structure on a simplicial metric tree X is an equivalence relation on germs of edges at vertices of X . If $X \in \mathcal{O}(\mathcal{G})$, the gate structure is required to be G -invariant. Equivalence classes are called gates. Given a gate structure \sim , a turn on X is legal if its germs are not in the same gate. A path in X is legal if it crosses only legal turns. (Note that legality does depend on the chosen gate structure.)

Straight maps naturally induce gate structures.

DEFINITION 2.26. — Given a straight map $f: X \rightarrow Y$, the gate structure \sim_f is defined by declaring equivalent two germs of X that have the same non-collapsed image under f . A turn (or a path) is called f -legal if it is legal with respect to \sim_f .

In case $X = Y$ there is also a different natural gate structure, that takes in account iterates, and that will be discussed in Section 2.7.

DEFINITION 2.27. — Let $X \in \mathcal{O}(\mathcal{G})$, $Y \in \overline{\mathcal{O}(\mathcal{G})}$. A straight map is called optimal if $\Lambda(X, Y) = \text{Lip}(f)$ and the tension graph is at least two-gated at every vertex (with respect to \sim_f). Moreover, an optimal map is called minimal if its tension graph consists of the union of axes of maximally stretched elements it contains.

Remark 2.28. — For all $X \in \mathcal{O}(\mathcal{G})$ and $Y \in \overline{\mathcal{O}(\mathcal{G})}$, there is always an optimal map $f: X \rightarrow Y$ (and it is usually not unique). Moreover, there is always a minimal optimal map $f: X \rightarrow Y$. In [11, 12] these facts are proved for $Y \in \mathcal{O}(\mathcal{G})$, but the proofs clearly work without any change for trees in $\overline{\mathcal{O}(\mathcal{G})}$, as all technicalities take place on X .

2.7. Train tracks

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G .

We have already seen that straight maps $f: X \rightarrow Y$ induce a natural gate structure \sim_f on X . In case $X = Y$, we consider also a second natural gate structure, namely:

- \sim_f : two germs of X are \sim_f -equivalent, if they have the same non-collapsed image under f ;
- $\langle \sim_{f^i} \rangle$: two germs of X are $\langle \sim_{f^i} \rangle$ if there is some i , so that they have the same non-collapsed image under f^i .

Train tracks maps were introduced in [4]. The terminology we use here may sound different, but it is in fact equivalent. (See [11, 12]).

DEFINITION 2.29. — Given a gate structure \sim on a metric simplicial tree X , a train track map $f: X \rightarrow X$ with respect to \sim is a straight map such that:

- (1) f sends edges to \sim -legal paths;
- (2) if $f(v)$ is a vertex, then f maps \sim -inequivalent germs at v to \sim -inequivalent germs at v .

It turns out (see [11, Section 8]) that if f is train track for some gate structure \sim , then in fact the relation \sim is stronger than (i.e. it contains) $\langle \sim_{f^i} \rangle$. In fact, if f is \sim -train track, then f is $\langle \sim_{f^i} \rangle$ -train track. (Also, since \sim_f is always weaker than $\langle \sim_{f^i} \rangle$, if f is \sim_f -train track then $\sim_f = \langle \sim_{f^i} \rangle$.) In what follows, we generically refer to a train track map as a map f which is train track with respect to $\langle \sim_{f^i} \rangle$.

DEFINITION 2.30. — Let $\phi \in \text{Aut}(\mathcal{G})$ and $X \in \mathcal{O}(\mathcal{G})$. A topological representative of ϕ at X is just an \mathcal{O} -map $f: X \rightarrow X\phi$. In other words, a map $f: X \rightarrow X$ such that $f(gx) = \phi(g)x$. A topological representative of $[\phi] \in \text{Out}(\mathcal{G})$ is a topological representative of some $\psi \in [\phi]$. A (simplicial)⁽¹⁾ train track representative of $[\phi]$ is a (simplicial) topological representative f which is train track with respect to $\langle \sim_{f^i} \rangle$. Points admitting (simplicial) train track representatives of $[\phi]$ are called (simplicial) train track points of ϕ .

Remark 2.31. — Train track representatives are always optimal maps (see [11, 12]), and their Lipschitz constant, if bigger than one, is the exponential growth rate of the represented automorphism.

In the case where $[\phi]$ is irreducible, the tension graph of any train track representative (which always exists by Theorem 2.47) is the whole graph.

Remark 2.32. — It is well-known (see for instance [11]) that if X is a train track point of ϕ , then there is a simplicial train track point Y of ϕ such that either $Y \in \Delta(X)$ or, at worse, $\Delta(Y)$ is a simplicial face of $\Delta(X)$. In particular, given a train track point X of ϕ , there is a simplicial train track point Y of ϕ which is in (uniformly) bounded distance from X .

2.8. Bounded cancellation, critical constant, Nielsen paths

For any tree T and $a, b \in T$, we will denote by $[a, b]$ the unique directed reduced (i.e. without backtracks) path in T from a to b . For a path p in T , we denote by $[p]$ the reduced path with the same endpoints as p . In other words, $[p]$ is “ p pulled tight”. For any reduced path $\beta = [a, b]$ in T , we denote by $l_T(\beta)$ its length.

DEFINITION 2.33 (Bounded Cancellation Constant). — Given two trees T, S and a continuous map $f: T \rightarrow S$, the bounded cancellation constant of f , denoted by $\text{BCC}(f)$, is defined as the supremum of all real numbers B

⁽¹⁾I.e. mapping vertices to vertices and edges to edgepaths.

with the property that there exist $a, b, c \in T$ with $c \in [a, b]$, such that $d_S(f(c), [f(a), f(b)]) = B$.

In other words, $\text{BCC}(f)$ is the best number such that for any $a, b \in T$ and $c \in [a, b]$, the point $f(c)$ belongs to the $\text{BCC}(f)$ -neighbourhood of $[f(a), f(b)] = [f([a, b])]$.

LEMMA 2.34 (Bounded Cancellation Lemma [23, Proposition 4.12] and [17, Proposition 9.6]). — Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . Let $T \in \mathcal{O}(\mathcal{G})$, and $S \in \overline{\mathcal{O}(\mathcal{G})}$. Let $f: T \rightarrow S$ be an \mathcal{O} -map. Then $\text{BCC}(f) \leq \text{Lip}(f) \text{vol}(T)$. Moreover, if $S \in \mathcal{O}(\mathcal{G})$ we get the sharper inequality $\text{BCC}(f) \leq \text{Lip}(f) \text{vol}(T) - \text{vol}(S)$.

COROLLARY 2.35. — Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . Let $T \in \mathcal{O}(\mathcal{G})$, and $S \in \overline{\mathcal{O}(\mathcal{G})}$. Let $f: T \rightarrow S$ be a straight map, and suppose that there is $\mu > 0$ such that $\lambda_e(f) \geq \mu$ for any edge e . If $g \in G$ is such that its axis in T can be written as a g -periodic concatenation of at most c f -legal pieces (as, for instance, edges), then for any $B \geq \text{BCC}(f)$ we have

$$\ell_S(g) \geq \mu \ell_T(g) - cB.$$

In particular we can take $B = \Lambda(T, S) \text{vol}(T)$.

Proof. — This is an immediate application of Bounded Cancellation Lemma 2.34. \square

DEFINITION 2.36 (Critical constant). — Given two metric trees T, S and an expanding Lipschitz map $f: T \rightarrow S$ (i.e. with $\text{Lip}(f) > 1$), the critical constant of f is defined as $\text{cc}(f) = \frac{2\text{BCC}(f)}{\text{Lip}(f)-1}$.

LEMMA 2.37. — For any metric tree T and any expanding train track map $f: T \rightarrow T$, we have $\text{cc}(f^n) \leq \text{cc}(f)$.

Proof. — It is immediate to check by induction that $\text{BCC}(f^{n+1}) \leq \text{BCC}(f) \left(\sum_{i=0}^n \text{Lip}(f)^i \right)$, whence $\frac{\text{BCC}(f^{n+1})}{\text{Lip}(f)^{n+1}-1} \leq \frac{\text{BCC}(f)}{\text{Lip}(f)-1}$. The claim follows because, since f is a train track map, we have $\text{Lip}(f^{n+1}) = \text{Lip}(f)^{n+1}$. \square

LEMMA 2.38. — Let $f: T \rightarrow T$ be a train track map defined on a metric tree T , with $\text{Lip}(f) = \lambda > 1$. Let γ be a path in T , containing a legal subpath p , with $l_T(p) \geq \text{cc}(f)$. Then for all $n > 0$:

- (i) $[f^n(\gamma)]$ contains a legal subpath of length at least $l_T(p)$;
- (ii) $[f^n(\gamma)]$ contains a legal subpath of length at least $\lambda^n(l_T(p) - \text{cc}(f))$.

In particular if p is longer than $\text{cc}(f) + 1$, then $l_T(f^n(\gamma)) > \lambda^n$.

Proof. — Since p is legal, the length of the surviving part of $f(p)$ in $[f(\gamma)]$, after cancellations, is at least

$$\lambda l_T(p) - 2 \text{BCC}(f) = \lambda l_T(p) - \text{cc}(f)(\lambda - 1) > \lambda l_T(p) - l_T(p)(\lambda - 1) = l_T(p).$$

Thus we can iterate, and we get

$$\begin{aligned} l_T([f^n(\gamma)]) &> \lambda^n l_T(p) - \sum_{i=0}^{n-1} \lambda^i 2 \text{BCC}(f) \\ &= \lambda^n l_T(p) - 2 \text{BCC} \frac{\lambda^n - 1}{\lambda - 1} \\ &= \lambda^n l_T(p) - \text{cc}(f) \lambda^n + \text{cc}(f) \\ &> \lambda^n (l_T(p) - \text{cc}(f)). \end{aligned} \quad \square$$

DEFINITION 2.39. — Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . Let $X \in \mathcal{O}(\mathcal{G})$ and $f: X \rightarrow X$ be a G -equivariant simplicial map. A (non-trivial) simplicial path p in X is called:

- (1) Nielsen path (Np) if $[f(p)] = gp$ for some $g \in G$;
- (2) periodic Nielsen path (pNp) if $[f^n(p)] = gp$ for some $n > 0$;
- (3) pre-periodic Nielsen path ($ppNp$) if $[f^{n+m}(p)] = gf^m(p)$ for some $g \in G$ and integers $n, m > 0$; we say that the periodic behaviour of a $ppNp$ starts before n_0 iterates if in the above formula we have $n < n_0$;
- (4) trivial if $[p]$ is a point, and pre-trivial if $[f^n(p)]$ is trivial for some positive integer n .

2.9. Candidates

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . As we have seen (Theorem 2.20) the stretching factor between two trees is realised by some hyperbolic group element. In fact, more is true.

THEOREM 2.40 ([11, Theorem 9.10] and [12, Lemma 7.1]). — For every $T \in \mathcal{O}(\mathcal{G})$, there is a set of hyperbolic elements $\text{Cand}(T)$, called candidates, such that for every $S \in \overline{\mathcal{O}(\mathcal{G})}$ the stretching factor $\Lambda(T, S)$ is realised on a candidate, that is

$$\Lambda(T, S) = \max_{g \in \text{Cand}(T)} \frac{\ell_S(g)}{\ell_T(g)}.$$

Moreover, the possible projections of candidates to the graph $\Gamma = G \setminus T$ are finitely many. Specifically, the projection of the translation axis of any candidate has one of the following forms (possibly containing both free and non-free vertices):

- a simple loop (an embedded \bigcirc);
- a figure eight, i.e. two simple loops that intersect on a point (an embedded $\bigcirc\circ$);
- a non-degenerate bar-bell, i.e. a path formed by two separated simple loops, joined by an embedded arc (an embedded $\bigcirc-\bigcirc$);
- a simply degenerate bar-bell, i.e. a path formed by a simple loop with attached an embedded arc ending to a non-free vertex (an embedded $\bigcirc-\bullet$);
- a doubly degenerate bar-bell, i.e. an embedded arc whose endpoints are non-free vertices (an embedded $\bullet-\bullet$);

We will need also the following lemma.

LEMMA 2.41 ([23, Theorem 4.7], see also [14, Lemma 2.18]). — For every $T \in \mathcal{O}(\mathcal{G})$, we can extract a finite set from $H \subseteq \text{Cand}(T)$, so that for every $S \in \overline{\mathcal{O}(\mathcal{G})}$,

$$\Lambda(T, S) = \max_{g \in H} \frac{\ell_S(g)}{\ell_T(g)}.$$

Moreover H does depend only on the simplex that T belongs to, and not to the particular metric of T .

COROLLARY 2.42. — The stretching factor function

$$\Lambda: \mathcal{O}(\mathcal{G}) \times \overline{\mathcal{O}(\mathcal{G})} \longrightarrow \mathbb{R}^+$$

is continuous on the second variable and lower semi-continuous on the first one, with respect to length function topology.

Proof. — We start from lower semi-continuity on the first variable, which does not require Lemma 2.41. Let $T \in \overline{\mathcal{O}(\mathcal{G})}$, $X_n \in \mathcal{O}(\mathcal{G})$, with $X_n \rightarrow X \in \mathcal{O}(\mathcal{G})$.

$$\begin{aligned} \liminf_{n \rightarrow \infty} \Lambda(X_n, T) &= \liminf_{n \rightarrow \infty} \max_{g \in \text{Hyp}(\mathcal{G})} \frac{\ell_T(g)}{\ell_{X_n}(g)} \\ &\geq \max_{g \in \text{Hyp}(\mathcal{G})} \liminf_{n \rightarrow \infty} \frac{\ell_T(g)}{\ell_{X_n}(g)} \\ &= \max_{g \in \text{Hyp}(\mathcal{G})} \lim_{n \rightarrow \infty} \frac{\ell_T(g)}{\ell_{X_n}(g)} \\ &= \max_{g \in \text{Hyp}(\mathcal{G})} \frac{\ell_T(g)}{\ell_X(g)} \\ &= \Lambda(X, T). \end{aligned}$$

Now we prove the continuity on the second variable. Let $T \in \mathcal{O}(\mathcal{G})$, $T_n \in \overline{\mathcal{O}(\mathcal{G})}$, with $T_n \rightarrow T_\infty \in \overline{\mathcal{O}(\mathcal{G})}$. We will show that $\Lambda(T, T_n) \rightarrow \Lambda(T, T_\infty)$. Let

denote by H the finite set of candidates of T that we get from Lemma 2.41. Then the following equalities hold (as H is finite):

$$\begin{aligned} \lim_{n \rightarrow \infty} \Lambda(T, T_n) &= \lim_{n \rightarrow \infty} \max_{g \in H} \frac{\ell_{T_n}(g)}{\ell_T(g)} \\ &= \max_{g \in H} \lim_{n \rightarrow \infty} \frac{\ell_{T_n}(g)}{\ell_T(g)} \\ &= \max_{g \in H} \frac{\ell_{T_\infty}(g)}{\ell_T(g)} \\ &= \Lambda(T, T_\infty). \end{aligned}$$

□

It is easy to construct examples where continuity on first variable fails.

Example 2.43. — Consider graphs as in Example 2.8 (Figure 2.1), with edge-lengths as follows. On the left side of Figure 2.1, the unique edge has length 1. On the right side, all edges have length $1/3$. Then, any infinite sequence X_n of right-side graphs converges to the left-side graph X ; this is to say that for any $g \in G$, $\lim_{n \rightarrow \infty} \ell_{X_n}(g) = \ell_X(g)$. In fact for any g , $\ell_{X_n}(g) = \ell_X(g)$ for all but finitely many n . However, for every n we have $\Lambda(X_n, X) = 3 \neq 1 = \Lambda(X, X)$. Also, this example shows that the volume function in general is not continuous with respect to the length function topology, as $\text{vol}(X_n) = 2/3$ while $\text{vol}(\lim_n X_n) = \text{vol}(X) = 1$.

2.10. Displacement function and Min-Set

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . For any $\phi \in \text{Aut}(C, \mathcal{G})$ we define the displacement function with respect to $\mathcal{O}(\mathcal{G})$ as

$$\begin{aligned} \lambda_\phi: \mathcal{O}(\mathcal{G}) &\longrightarrow \mathbb{R} \\ \lambda_\phi(X) &= \Lambda(X, X\phi). \end{aligned}$$

We define also the minimal ϕ -displacement of a simplex Δ of $\mathcal{O}(\mathcal{G})$ as

$$\lambda_\phi(\Delta) = \inf \{ \lambda_\phi(X) : X \in \Delta \}$$

and the minimal displacement of ϕ as

$$\lambda(\phi) = \inf \{ \lambda_\phi(X) : X \in \mathcal{O}(\mathcal{G}) \}.$$

The set of minimally displaced points in $\mathcal{O}(\mathcal{G})$ or *Min-Set*, is defined as

$$\text{Min}(\phi) = \{ X \in \mathcal{O}(\mathcal{G}) : \lambda_\phi(X) = \lambda(\phi) \}.$$

Finally, the set of minimally displaced points with co-volume one is

$$\text{Min}_1(\phi) = \{X \in \mathcal{O}_1(\mathcal{G}) : \lambda_\phi(X) = \lambda(\phi)\}.$$

We remark that the displacement actually depends only on $[\phi] \in \text{Out}(\mathcal{G})$. We also remark that in case $[\phi]$ is reducible (see Section 2.11) the Min-Set has to be defined in the simplicial bordification of $\mathcal{O}(\mathcal{G})$, but we omit this point of view here because in this paper we are interested in irreducible automorphisms. We just say here that in the irreducible case, the Min-Set is always connected, and coincides with the set of points admitting train track representatives. We refer the interested reader to [11, 12, 13] for a detailed discussion on such properties.

2.11. Irreducible automorphisms

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G .

DEFINITION 2.44. — An element $[\phi] \in \text{Out}(\mathcal{G})$ is called \mathcal{G} -reducible (or simply reducible), if some $\psi \in [\phi]$ admits a topological representative $f: T \rightarrow T\psi$, $T \in \mathcal{O}(\mathcal{G})$, having a proper G -invariant, f -invariant sub-forest S which contains the axis of some \mathcal{G} -hyperbolic element. $[\phi]$ is \mathcal{G} -irreducible (or simply irreducible) if it is not reducible.

Remark 2.45. — We can define irreducibility in terms of free factor systems. An automorphism class $[\phi]$ is \mathcal{G} -irreducible if \mathcal{G} is a maximal ϕ -invariant free factor system (for more details, see [11]). Another viewpoint of the same fact is that $[\phi]$ is reducible if and only if there is some point X in some face at infinity of some simplex so that $\lambda_\phi(X) < \infty$.

Remark 2.46. — Let G be a finitely generated group, and let $[\phi] \in \text{Out}(G)$. First we note that if \mathcal{G} is the Grushko decomposition of G , then $[\phi] \in \text{Out}(\mathcal{G})$. Next, we observe that there always exists a free product decomposition \mathcal{G}' of G such that $[\phi]$ is irreducible as an element of $\text{Out}(\mathcal{G}')$. Note that in general \mathcal{G}' is not unique. In fact, there are examples where there are infinitely many different spaces for which $[\phi]$ is irreducible. An example is the identity outer automorphism.

We summarise below some well-known properties of irreducible automorphisms.

THEOREM 2.47 ([11, 12]). — Let $[\phi] \in \text{Out}(\mathcal{G})$ be irreducible. Then:

- (1) it admits a train track representative $f: T \rightarrow T\phi$ (see [11, Lemma 8.17, Theorem 8.18]);

- (2) the set of train track points of ϕ coincides with the set $\text{Min}(\phi)$ of minimally displaced points (see [11, Theorem 8.19]);
- (3) there is an $\epsilon > 0$ (that depends only $\text{rank}(\mathcal{G})$ and on $\lambda(\phi)$) for which $\text{Min}_1(\phi)$ is contained in the ϵ -thick part $\mathcal{O}_1(\mathcal{G}, \epsilon)$ (see [11, proof of Theorem 8.4] and [12, Propositions 5.5 and 5.6]).

The third item of Theorem 2.47 combined with Theorem 2.19, implies the following.

COROLLARY 2.48. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be irreducible. Then there exists some constant $C = C(\phi)$, for which for all $X, Y \in \text{Min}_1(\phi)$ we have*

$$\Lambda(X, Y) \leq \Lambda(Y, X)^C.$$

The next theorem is key for our present paper.

THEOREM 2.49 ([14]). — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be irreducible and with $\lambda(\phi) > 1$. Then the simplicial structure of $\text{Min}(\phi)$ (as a subset of $\mathcal{O}(\mathcal{G})$) is locally finite. In particular, the set $\text{Min}_1(\phi) = \mathcal{O}_1 \cap \text{Min}(\phi)$ is a locally finite simplicial complex.*

2.12. Primitive automorphisms

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G .

For any $T \in \mathcal{O}(\mathcal{G})$ and any simplicial \mathcal{O} -map $f: T \rightarrow T$, we can define the transition matrix M_f of f as follows. We label orbits of edges e_1, \dots, e_n and define the (i, j) coefficient of M_f as the number of times that $f(e_i)$ crosses the orbit of e_j (in either direction).

A matrix is called *non-negative* if all its entries are non-negative. A non-negative matrix is called *irreducible* if for any (i, j) it has a power for which the (i, j) entry is positive, and it is called *primitive* if it has a power so that all entries are positive.⁽²⁾ Clearly primitive implies irreducible. It is immediate to check that $[\phi] \in \text{Out}(\mathcal{G})$ is \mathcal{G} -irreducible if and only if any simplicial map representing $[\phi]$ has irreducible transition matrix.

DEFINITION 2.50 (Primitive automorphism). — *An automorphism which can be represented somewhere in $\mathcal{O}(\mathcal{G})$ by a simplicial train track map having primitive transition matrix is called \mathcal{G} -primitive (or simply primitive).*

⁽²⁾ We notice that in [3] the authors use the terminology *aperiodic* for primitive.

We recall that a train track representative does not need to be simplicial and in that case, the transition matrix is not even defined. However, as we have seen before, there are always simplicial train track representatives for irreducible automorphisms.

LEMMA 2.51. — *Suppose $[\phi] \in \text{Out}(\mathcal{G})$ is \mathcal{G} -primitive. Then for any $T \in \mathcal{O}(\mathcal{G})$, if $f: T \rightarrow T$ is any simplicial train track representative of $[\phi]$, then the transition matrix of f is primitive.*

Proof. — This is well-known in the free case and the proof is exactly the same in the general case (see [3, Lemma 3.1.14] for the proof). \square

Remark 2.52. — Note that a \mathcal{G} -primitive $[\phi]$ has always exponential growth, i.e. $\lambda(\phi) > 1$. Note also that for primitive automorphisms the condition on the transition matrix is only required for train track representatives. In particular, *a priori* it may happen that $[\phi]$ is primitive and reducible. We use \mathcal{G} -primitive irreducible (or simply primitive irreducible) to refer to those $[\phi]$ that are both \mathcal{G} -primitive and \mathcal{G} -irreducible.

We note that in the absolute case (when \mathcal{G} is trivial) that irreducibility and expanding implies primitivity. However this is not true for general free factor systems \mathcal{G} .

2.13. Arc stabiliser lemma

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . In this section we first describe points in $\overline{\mathcal{O}(\mathcal{G})}$ in terms of a chosen base point, and then prove a lemma about fixed points of elliptic elements of boundary points, that will be used in last step of the proof of Theorem 7.20.

The standing assumption in this subsection is that any $T \in \overline{\mathcal{O}(\mathcal{G})}$ is the limit point of a sequence of points in $\mathcal{O}(\mathcal{G})$. This is only certain when G is countable. The issue is that $\overline{\mathcal{O}(\mathcal{G})}$ is compact but, *a priori*, may not be sequentially compact and is a subspace of a Cartesian product, which is exactly the type of topological space which can be compact without being sequentially compact.

However, the results and proofs in this section remain true for any G , regardless of countability, by replacing sequences with nets (in particular, Lemmas 2.58 and 2.59 are true in general). In order to help the reader, we give the version of this argument for sequences and explain in Remark 2.61 how to extend it to the general case (see also [24] for the countable case).

Moreover, we only use Lemmas 2.58 and 2.59 in the case where the tree T is in fact the limit point of a sequence (Lemma 2.62 and Theorem 7.11).

Let $X \in \mathcal{O}(G)$ be a fixed reference point. Let $T \in \overline{\mathcal{O}}$ and let $Y_n \in \mathcal{O}(\mathcal{G})$ be a sequence that converges projectively to T , i.e. there is a sequence of positive numbers $\mu_n > 0$ so that:

$$\lim_{n \rightarrow \infty} \frac{Y_n}{\mu_n} = T$$

(with respect to the length function topology). Let $f_n: X \rightarrow Y_n$ be optimal maps. We define

$$d_n: X \times X \longrightarrow \mathbb{R}, \quad \text{where } d_n(x, y) = \frac{d_{Y_n}(f_n(x), f_n(y))}{\mu_n}.$$

As $\frac{Y_n}{\mu_n}$ converges to T , by Corollary 2.42 we get that the sequence $\frac{\Lambda(X, Y_n)}{\mu_n}$ converges to $\Lambda(X, T)$. In particular, this implies that $\frac{\Lambda(X, Y_n)}{\mu_n}$ is a bounded sequence and hence, since $\text{Lip}(f_n) = \Lambda(X, Y_n)$, we have

$$0 \leq d_n(x, y) \leq \frac{\Lambda(X, Y_n)}{\mu_n} d_X(x, y) \leq \sup_n \left\{ \frac{\Lambda(X, Y_n)}{\mu_n} \right\} d_X(x, y).$$

At this point we would like to take a limit of the d_n , but boundedness is not enough to guarantee a pointwise converging subsequence. The easiest way to do this is to take an ultralimit (or ω -limit) (see [16], for the definitions and the properties of ultralimits and ultrafilters). We briefly recap here.

DEFINITION 2.53. — *A non-principal ultrafilter on \mathbb{N} is a function from the power set of \mathbb{N} , $\omega: \mathcal{P}(\mathbb{N}) \rightarrow \{0, 1\}$, such that:*

- $\omega(\emptyset) = 0$, $\omega(\mathbb{N}) = 1$;
- $\omega(A \sqcup B) = \omega(A) + \omega(B)$ (i.e. ω is additive on disjoint subsets of \mathbb{N});
- ω is zero on any finite subset of \mathbb{N} .

Remark 2.54. — For the reader unfamiliar with this point of view, the second point above is crucial, and we emphasise that ω can only take values 0 and 1, so the additivity on disjoint subsets is a strong restriction. Informally, we think of the subsets A with $\omega(A) = 1$ as “big” and the others small. It is then straightforward to see that there do not exist two disjoint subsets that are both big, and that the intersection of any two big subsets is again big.

Limits are then defined as follows.

DEFINITION 2.55. — *Let ω be a non-principal ultrafilter on \mathbb{N} . For any sequence $(a_n)_{n \in \mathbb{N}}$ of real numbers, we say that $l \in \mathbb{R}$ is the ω -limit of a_n — and we write $\lim_{\omega} a_n = l$ — if for every $\varepsilon > 0$, the set $\mathbb{N}_{\varepsilon} = \{n \in \mathbb{N} : |a_n - l| < \varepsilon\}$ satisfies $\omega(\mathbb{N}_{\varepsilon}) = 1$.*

We then get the following facts whose proofs are left to the reader.

PROPOSITION 2.56. — *Let ω be a non-principal ultrafilter on \mathbb{N} . Let a_n and b_n be sequences of real numbers.*

- *If $\lim_n a_n = l$ then $\lim_\omega a_n = l$.*
- *If $\lim_\omega a_n$ exists, it is unique (it may depend on ω).*
- *If a_n is bounded, then $\lim_\omega a_n$ exists.*
- *The usual algebra of limits is valid for ultralimits: $\lim_\omega (a_n \pm b_n) = (\lim_\omega a_n) \pm (\lim_\omega b_n)$, $\lim_\omega a_n \cdot b_n = (\lim_\omega a_n) \cdot (\lim_\omega b_n)$.*
- *If $a_n \geq b_n$ then $\lim_\omega a_n \geq \lim_\omega b_n$.*
- *Ultralimits commute with finite max and min: $\lim_\omega (\max\{a_n, b_n\}) = \max\{\lim_\omega a_n, \lim_\omega b_n\}$.*

To proceed, we apply this in our situation. We let ω be a non-principal ultrafilter on \mathbb{N} and define:

$$d^+(x, y) = \lim_\omega d_n(x, y).$$

It is clear that this is an equivariant pseudo-metric on X , and we can study the associated quotient space X^+ . Here, elements of X^+ are “balls of radius 0”. That is, elements of X^+ are equivalence classes $[x] := \{y \in X : d^+(x, y) = 0\}$. (It may be worth to note here that X^+ is not a priori uniquely determined, as it depends on the chosen ultrafilter).

LEMMA 2.57. — *Let ω be a non-principal ultrafilter on \mathbb{N} . For all $x, y \in X$ let $d^+(x, y) = \lim_\omega d_n(x, y)$ and let $X^+ := (\frac{X}{(d^+=0)}, d^+)$ be the corresponding quotient space. Then:*

- (i) *the quotient map $X \rightarrow X^+$ is Lipschitz continuous and equivariant;*
- (ii) *(X^+, d^+) is an \mathbb{R} -tree with an isometric G -action, $g[x] = [gx]$;*
- (iii) *X^+ and T admit the same length function; $\ell_{X^+} = \ell_T$; in particular, X^+ is non-trivial and its minimal invariant subtree is equivariantly isometric to T .*

Proof. — For any $x, y \in X$, we have

$$d_n(x, y) \leq \frac{\Lambda(X, Y_n) d_X(x, y)}{\mu_n}.$$

It follows that $d^+([x], [y]) \leq K d_X(x, y)$, where $K = \Lambda(X, T)$ (by Corollary 2.42). Thus the quotient map $X \rightarrow X^+$ is K -Lipschitz, and equivariance is straightforward.

One can then see that $X^+ = \frac{X}{(d^+=0)}$ is a path connected 0-hyperbolic metric space, hence an \mathbb{R} -tree equipped with an isometric G -action.

It remains to show that $\ell_{X^+} = \ell_T$. For any tree Z , it is known that $\ell_Z(g) = \max\{0, d_Z(p, g^2p) - d_Z(p, gp)\}$ for any $p \in Z$ (see [6], for more details). Let $x \in X, g \in G$. Then

$$\begin{aligned} \ell_{X^+}(g) &= \max\{0, d^+(x, g^2x) - d^+(x, gx)\} \\ &= \max\{0, \lim_{\omega} (d_n(x, g^2x) - d_n(x, gx))\} \\ &= \max\left\{0, \lim_{\omega} \frac{d_{Y_n}(f_n(x), f_n(g^2x)) - d_{Y_n}(f_n(x), f_n(gx))}{\mu_n}\right\} \\ &= \max\left\{0, \lim_{\omega} \frac{d_{Y_n}(f_n(x), g^2f_n(x)) - d_{Y_n}(f_n(x), gf_n(x))}{\mu_n}\right\} \\ &= \lim_{\omega} \frac{l_{Y_n}(g)}{\mu_n} \\ &= \lim_n \frac{l_{Y_n}(g)}{\mu_n} \\ &= \ell_T(g). \end{aligned}$$

□

LEMMA 2.58 (Arc Stabiliser Lemma). — *Let $T \in \overline{\mathcal{O}(\mathcal{G})}$. If $1 \neq g \in G$ is \mathcal{G} -elliptic, then g fixes a unique a point of T . In particular no \mathcal{G} -elliptic, non-trivial element, point-wise fixes a non-trivial arc in T .*

Proof. — As above, X is our reference point. We let $Y_n \in \mathcal{O}(\mathcal{G})$ and $\mu_n > 0$ be such that $\frac{Y_n}{\mu_n} \rightarrow T$, and build X^+ as a ω -limit of metrics d_n on X . Since the minimal invariant subtree of X^+ is equivariantly isometric to T , it is sufficient to prove the theorem for arc stabilisers in X^+ .

Since edge stabilisers are trivial in $\mathcal{O}(\mathcal{G})$, g fixes a unique point p in X , and a unique point p_n in each Y_n . Thus

$$(2.1) \quad \forall w \in X, \quad d_X(w, gw) = 2d_X(w, p)$$

$$(2.2) \quad \forall w \in Y_n, \quad d_{Y_n}(w, gw) = 2d_{Y_n}(w, p_n).$$

We claim that $[p]$ is the unique point of X^+ fixed by g . Let's suppose that $[u] \in X^+$ is fixed by g . Thus $d^+(u, gu) = 0$.

By equivariance of the maps f_n , we get that $f_n(p) = p_n$. Then, by (2.2)

$$\begin{aligned} 2d_n(u, p) &= 2d_{Y_n}(f_n(u), p_n)/\mu_n \\ &= d_{Y_n}(f_n(u), gf_n(u))/\mu_n \\ &= d_n(u, gu) \\ &\longrightarrow \omega d^+(u, gu) = 0. \end{aligned}$$

As a consequence, $d^+(u, p) = 0$ and so $[u] = [p]$ in X^+ . □

LEMMA 2.59. — *Let $T \in \overline{\mathcal{O}(\mathcal{G})}$. Let $H \leq G$ be T -elliptic. Suppose that H contains a \mathcal{G} -elliptic subgroup $A \neq 1$. Then $\text{Fix}_T(H) = \text{Fix}_T(A) = \{v\}$, where v is a point of T .*

Proof. — Let $1 \neq a \in A$. The group element a fixes a unique point of T , by Lemma 2.58. Given any $h \in H$, the subgroup $\langle a, h \rangle$ fixes a point v of T (by a well-known result of Serre, see [31]). Therefore, $\text{Fix}_T(\langle a, h \rangle) = \text{Fix}_T(\langle a \rangle) = \{v\}$, for all $h \in H$, and hence $\text{Fix}_T(H) = \{v\}$. \square

Remark 2.60. — Note that here we are calling a subgroup elliptic if the restriction of the length function is zero on the subgroup. This is weaker than the definition — often used — that the subgroup fixes a point in the tree, although they coincide for finitely generated subgroups.

Remark 2.61. — In order to prove Lemma 2.58 when G is not countable, we argue via nets. The reason this is necessary is that we cannot guarantee that every point in $\overline{\mathcal{O}(G)}$ is the limit point of a sequence. We only really use it when T is the limit point of a sequence. However, it is true in general with substantially the same proof via nets.

Concretely, one takes the directed set Fin_G of all finite subsets of G ordered by inclusion. A net is then a function from Fin_G to our space and takes the place of sequences. One striking aspect of working with $\overline{\mathcal{O}(G)}$ is that we can use this set Fin_G as the universal indexing set. This is due to the fact that $\overline{\mathcal{O}(G)}$ is a subset of a Cartesian product whose indexing set is G , and so basic open subsets are described by finite subsets of G along with open subsets of \mathbb{R} in the corresponding coordinates.

For this reason, a net for us will always be a function from Fin_G to our space; either $\overline{\mathcal{O}(G)}$, or the composition of such a function with one of the projection maps, resulting in a function from Fin_G to \mathbb{R} .

A *tail* of Fin_G , is a subset $\text{Tail}_F = \{E \in \text{Fin}_G : F \subseteq E\}$ for some $F \in \text{Fin}_G$. A net to \mathbb{R} , $x: \text{Fin}_G \rightarrow \mathbb{R}$ then has limit l if every open set around l contains the image of a tail. Concretely this means that for every $\varepsilon > 0$, there exists an $F \in \text{Fin}_G$ so that $|x(E) - l| < \varepsilon$ for every $F \subseteq E \in \text{Fin}_G$. In this case we write $x(\cdot) \rightarrow l$. It is important to note here that the intersection of finitely many tails is again a tail; this is used repeatedly throughout.

Similarly, a net $Y: \text{Fin}_G \rightarrow \overline{\mathcal{O}(G)}$ has a limit T if $\ell_{Y(\cdot)}(g) \rightarrow \ell_T(g)$ for all $g \in G$. This is equivalent to saying that every open set in $\overline{\mathcal{O}(G)}$ containing T contains the image of a tail.

One can now see that every $T \in \overline{\mathcal{O}(G)}$ is the limit of a net of points — indexed by Fin_G — in $\mathcal{O}(G)$. For instance, one can do the following: for every $F \in \text{Fin}_G$, consider the basic open sets

$$B_F(T) = \{S \in \overline{\mathcal{O}(G)} : |\ell_S(g) - \ell_T(g)| < 1/n\},$$

where $n = |F| + 1$. For each F , choose any $Y_F \in B_F(T) \cap \mathcal{O}(\mathcal{G})$. Notice that if $F \subseteq E$, then $B_E(T) \subseteq B_F(T)$. Since any open set containing T contains some $B_F(T)$, we immediately see that $Y_F \rightarrow T$.

We then introduce a ultrafilter ω on Fin_G which, as above, is a function $\omega: \mathcal{P}(\text{Fin}_G) \rightarrow \{0, 1\}$. However, the condition we require this time is that ω is 1 on every tail — being zero on finite sets is not quite enough. Instead we require that $\omega(\text{Tail}_F) = 1$ for every Tail_F .

Concretely, the tails of Fin_G are intersection-closed and so form a filter-base. The set of supersets of tails then forms a filter, and we choose any maximal filter ω containing that filter; this will be an ultrafilter with the properties we require. This can be achieved via Zorn's Lemma.

One then defines ω limits as above, putting $\lim_{\omega} x(\cdot) = l$ if we have for every $\varepsilon > 0$, $\omega(\{F \in \text{Fin}_G : |x_F - l| < \varepsilon\}) = 1$. As before, $x(\cdot) \rightarrow l$ implies that $\lim_{\omega} x(\cdot) = l$ and every bounded sequence of reals has a unique ω limit (depending on ω). The rest proceeds in the same way.

2.14. Limit trees for irreducible automorphisms

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . The theory of attracting and repelling trees of a fully irreducible automorphism is well studied in the free case. We will see in north-south dynamics Theorem 5.2, that such trees exist and they are the unique fixed points, in $\overline{\mathcal{O}(\mathcal{G})}$, of a primitive irreducible automorphism.

In this section we recall the construction of the attracting tree for any irreducible automorphism with exponential growth (not necessarily primitive), starting from minimally displaced points. In fact, the only property of irreducibility that is used here is the existence of train track representatives (Theorem 2.47). We prove also some useful properties which can be proved using just the train track properties. Note that the repelling tree of ϕ will be just the attracting tree of ϕ^{-1} , so we will focus here only on the attracting tree.

LEMMA 2.62 (Attracting tree). — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be such that there exists $X \in \text{Out}(\mathcal{G})$ supporting a train track map $f: X \rightarrow X$ representing $[\phi]$, with $\text{Lip}(f) = \lambda > 1$ (for example if $[\phi]$ is \mathcal{G} -irreducible with exponential growth, and $X \in \text{Min}(\phi)$). Then the following limit exists:*

$$X_{+\infty} = \lim_n \frac{X\phi^n}{\lambda^n}.$$

That is to say, for any $g \in G$ the following limit exists:

$$\ell_{X_{+\infty}}(g) = \lim_n \frac{\ell_X(\phi^n(g))}{\lambda^n}.$$

Proof. — As in the construction of X^+ in Section 2.13, we take: $d^+(x, y) = \lim_n \frac{d_X(f^n(x), f^n(y))}{\lambda^n}$, and $X_{+\infty} = \frac{X}{d^+=0}$. Train track properties ensure convergence on the nose, without use of an ultralimit. \square

DEFINITION 2.63. — Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible with exponential growth, and let $X \in \text{Min}(\phi)$. We define the attracting tree or stable tree of ϕ (with respect to X), as:

$$X_{+\infty} = \lim_{n \rightarrow \infty} \frac{X\phi^n}{\lambda(\phi)^n}.$$

Similarly, we define the repelling tree or unstable tree of ϕ , with respect to $Y \in \text{Min}(\phi^{-1})$, as $Y_{-\infty} = \lim_{n \rightarrow +\infty} Y\phi^{-n}/\lambda(\phi^{-1})^n$.

For any $X \in \text{Min}(\phi)$, the following facts are straightforward:

- $\forall Y \in \mathcal{O}(\mathcal{G})$, $\Lambda(Y, X_{+\infty}) = \Lambda(Y\phi, X_{+\infty}\phi)$ (as ϕ acts by isometries);
- $X_{+\infty}\phi = \lambda(\phi)X_{+\infty}$ (easy application of train track properties);
- $\Lambda(X, X_{+\infty}) = 1$ (by continuity of Λ on the second variable, Corollary 2.42).

PROPOSITION 2.64 (Stable map). — Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible, of exponential growth rate $\lambda = \lambda(\phi) > 1$. Let $X, Y \in \text{Min}(\phi)$ and let $X_{+\infty}$ be the attracting tree with respect to X . Let $\varphi: Y \rightarrow Y$ be a train track representative of $[\phi]$. Then there is a minimal optimal map f_Y from Y to $X_{+\infty}$, called the stable map, with the following properties:

- any f_Y -legal periodic line γ in the tension graph of f_Y is φ -legal, and $\varphi(\gamma)$ is again f_Y -legal;
- the tension graph of f_Y is Y ;
- if e is an edge of Y then for any positive integer n , any subpath of $\varphi^n(e)$ is f_Y -legal.

Proof. — There is a minimal optimal map $f_Y: Y \rightarrow X_{+\infty}$ (Remark 2.28). In particular, $\text{Lip}(f_Y) = \Lambda(Y, X_{+\infty})$. Moreover, since the tension graph of optimal maps is everywhere at least two-gated, there is some $g \in \text{Hyp}(G)$ whose axis is f_Y -legal and contained in the tension graph of f_Y .

Let $g \in G$ be one of such \mathcal{G} -hyperbolic element. Then $\ell_{X_{+\infty}}(g) = \Lambda(Y, X_{+\infty})\ell_Y(g)$. On the other hand, since $Y \in \text{Min}(\phi)$, we have $\ell_Y(\phi(g)) \leq \lambda\ell_Y(g)$, with equality precisely when g is φ -legal. Combining these facts,

we have:

$$\begin{aligned}\ell_{(Y\phi)}(g) &\leq \lambda\ell_Y(g), \\ \ell_{(X_{+\infty}\phi)}(g) &= \lambda\ell_{X_{+\infty}}(g), \\ \frac{\ell_{X_{+\infty}}(g)}{\ell_Y(g)} &= \Lambda(Y, X_{+\infty}), \\ \frac{\ell_{(X_{+\infty}\phi)}(g)}{\ell_{(Y\phi)}(g)} &\leq \Lambda(Y\phi, X_{+\infty}\phi).\end{aligned}$$

Hence,

$$\begin{aligned}\Lambda(Y, X_{+\infty}) &= \frac{\ell_{X_{+\infty}}(g)}{\ell_Y(g)} \\ &= \frac{\lambda\ell_{X_{+\infty}}(g)}{\lambda\ell_Y(g)} \\ &\leq \frac{\ell_{(X_{+\infty}\phi)}(g)}{\ell_{(Y\phi)}(g)} \\ &\leq \Lambda(Y\phi, X_{+\infty}\phi) \\ &= \Lambda(Y, X_{+\infty}),\end{aligned}$$

whence we have equality throughout, and in particular $\ell_Y(\phi(g)) = \lambda\ell_Y(g)$. Hence f_Y -legal axes in the tension graph of f_Y are also φ -legal. Now, since the axis γ of g is φ -legal, and φ is train track, $\varphi^n(\gamma)$ remains φ -legal, and we have

$$\Lambda(Y, X_{+\infty}) = \frac{\ell_{X_{+\infty}}(g)}{\ell_Y(g)} = \frac{\lambda^n\ell_{X_{+\infty}}(g)}{\lambda^n\ell_Y(g)} = \frac{\ell_{X_{+\infty}}(\phi^n(g))}{\ell_Y(\phi^n(g))} \leq \Lambda(Y, X_{+\infty}),$$

whence the inequality is an equality and the axis of $\phi^n(g)$ — which is $\varphi^n(\gamma)$ because of φ -legality — is f_Y -legal and in the tension graph of f_Y . To prove that the tension graph of f_Y is the whole of Y , observe that $\bigcup_n \varphi^n(\gamma)$ is clearly ϕ -invariant, so it must be the whole Y .

The last claim now follows from the previous ones, as every edge can be extended to an f_Y -legal periodic line, which is φ -legal and all of whose iterates under φ are both f_Y -and φ -legal. \square

In the next proposition we prove that the homothety class of the attracting tree doesn't depend on the train track point that we choose as base point.

PROPOSITION 2.65. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible with exponential growth, and let $X, Y \in \text{Min}(\phi)$. Let $X_{+\infty}$ be the attracting tree for X , and $Y_{+\infty}$ be that for Y . Then*

$$X_{+\infty} = \Lambda(Y, X_{+\infty})Y_{+\infty}.$$

Proof. — Let $f_Y : Y \rightarrow X_{+\infty}$ be the stable map given by Proposition 2.64 (in particular f_Y stretches any edge by $\text{Lip}(f_Y) = \Lambda(Y, X_{+\infty})$). Let $g \in G$, and represent it as a path in $G \setminus Y$ with n_g edges. n_g can be zero, for instance if g is elliptic. Then by Proposition 2.64, $\phi^n(g)$ is represented as a concatenation of at most n_g f_Y -legal pieces. Hence, by Corollary 2.35,

$$\Lambda(Y, X_{+\infty})\ell_Y(\phi^n(g)) - n_g B \leq \ell_{X_{+\infty}}(\phi^n(g)) \leq \Lambda(Y, X_{+\infty})\ell_Y(\phi^n(g)),$$

where B is the bounded cancellation constant of f_Y , and the second inequality just follows from the definition of $\Lambda(Y, X_{+\infty})$. It follows that

$$\begin{aligned} l_{X_{+\infty}}(g) &= \lim_{n \rightarrow \infty} \frac{l_{X_{+\infty}}(\phi^n(g))}{\lambda^n} \\ &= \Lambda(Y, X_{+\infty}) \lim_{n \rightarrow \infty} \frac{\ell_Y(\phi^n(g))}{\lambda^n} \\ &= \Lambda(Y, X_{+\infty})l_{Y_{+\infty}}(g). \end{aligned} \quad \square$$

Note that the uniqueness of limit trees is a direct corollary of Theorem 5.2 under the extra assumption of primitivity of the matrix, but the previous proposition provides us with an exact description of the un-projectivised limits in the general irreducible case.

2.15. Relative boundaries and laminations

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . For any metric tree T , we agree that:

- a *half-line* in T is an isometric embedding $[0, \infty) \rightarrow T$;
- \overline{T} is the metric completion of T ;
- $\partial_\infty T$ is the Gromov boundary of T , i.e. the set of half-lines in T up to the equivalence relation \sim , where two half-lines $L \sim L'$ if and only if L, L' differ only on a compact set;
- $V_\infty(T)$ is the collection of vertices of T with infinite valence (if $T \in \mathcal{O}(\mathcal{G})$, this coincides with non-free vertices with infinite stabiliser);
- $\partial T = \partial_\infty T \cup V_\infty(T)$;
- $\widehat{T} = T \cup \partial T$;
- $\partial^2 T = \partial T \times \partial T \setminus \{(P, P) : P \in \partial T\}$;
- a *direction* based at a point P of \widehat{T} , is a connected component $\widehat{T} \setminus \{P\}$;
- the *observer's topology* of \widehat{T} is the topology generated by the set of directions.

It is easy to see that \widehat{T} is a compact set, equipped with the observer's topology. Moreover, ∂T is a closed subset of \widehat{T} and therefore compact. The following lemma shows that the boundary does only depend on \mathcal{G} and not on the chosen tree $T \in \mathcal{O}(\mathcal{G})$.

LEMMA 2.66 ([18, Lemma 2.2]). — *Let $T, S \in \mathcal{O}(\mathcal{G})$. Then any G -equivariant map $f: T \rightarrow S$ has a unique continuous extension $\widehat{f}: \widehat{T} \rightarrow \widehat{S}$. Moreover, the restriction map $h := \widehat{f}|_{\partial T}$ is a natural homeomorphism $\partial T \rightarrow \partial S$ (it does not depend on f) with $h(\partial_\infty T) = \partial_\infty S$ and $h(V_\infty(T)) = V_\infty(S)$.*

Therefore, the notions of $\partial(G, \mathcal{G})$, $\partial_\infty(G, \mathcal{G})$, $\partial^2(G, \mathcal{G})$, $V_\infty(G, \mathcal{G})$ can be naturally defined as ∂T , $\partial_\infty T$, $\partial^2 T$, $V_\infty(T)$ for a $T \in \mathcal{O}(\mathcal{G})$. Note that $\partial_\infty(G, \mathcal{G})$ can be identified with the set of simple infinite words in the free product length given by (G, \mathcal{G}) .

In particular, for any \mathcal{G} -hyperbolic group element $g \in G$, we can define the infinite words $g^{+\infty} = \lim_{n \rightarrow +\infty} g^n$ and $g^{-\infty} = \lim_{n \rightarrow +\infty} g^{-n}$. In this case, $(g^{-\infty}, g^\infty) \in \partial^2(G, \mathcal{G})$.

There is a natural \mathbb{Z}_2 -action on $\partial^2(G, \mathcal{G})$ given by flipping coordinates $(P, Q) \mapsto (Q, P)$.

DEFINITION 2.67. — *An algebraic lamination is a closed G -invariant, flip-invariant, subset of $\partial^2(G, \mathcal{G})$. Elements of $\partial^2(G, \mathcal{G})$ are called algebraic leaves. Given $T \in \mathcal{O}(\mathcal{G})$, a (bi)(infinite) line L in T represents an algebraic leaf $(P, Q) \in \partial^2(G, \mathcal{G})$ if its endpoints correspond to (P, Q) under the natural homeomorphism given by Lemma 2.66.*

2.16. Attracting and repelling laminations

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . Attracting and repelling laminations for irreducible automorphisms with exponential growth can be defined as in the classical case (see [2] for the free case). Classical proofs work also in the present case, as they are based only on the properties of train track maps.

More precisely, let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible with $\lambda(\phi) > 1$. Let $f: T \rightarrow T$ be a train track representative of $[\phi]$, and let e be an edge of T . Consider iterates $f^n(e)$ and group elements $g_n \in G$ such that $g_n f^n(e)$ intersects a fixed fundamental domain for the G -action on T . Then the limit of $g_n f^n(e)$ is a line in T , hence it represents an algebraic leaf $L \in \partial^2(G, \mathcal{G})$. The attracting (or stable) lamination Λ_ϕ^+ is defined as the closure of the G -orbit of L . Any line in the G -orbit of L is called a generic line of Λ_ϕ^+ .

This construction depends a priori on T, f, e, g_n . In fact, when $[\phi]$ is \mathcal{G} -primitive, it does not depend on the choices made (see [2, Section 1] for the proof in the free case). We define the repelling lamination of ϕ as the attracting lamination of ϕ^{-1} , and is denoted by $\Lambda_\phi^- := \Lambda_{\phi^{-1}}^+$.

DEFINITION 2.68. — We say that (the conjugacy class of) a subgroup $A < G$ carries Λ_ϕ^+ if there is $T \in \mathcal{O}(\mathcal{G})$ containing a minimal A -tree, which contains a line that realises a generic leaf of Λ_ϕ^+ .

3. The distance of points of $\text{Min}(\phi)$ from $\text{Min}(\phi^{-1})$ is uniform

Let's fix a free factor system $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ of a group G . In this section, we prove a result which could be of independent interest. More specifically, we show that if $[\phi]$ is irreducible and $\lambda(\phi) > 1$, then the distance of a point of $\text{Min}(\phi)$ from the set $\text{Min}(\phi^{-1})$ is uniformly bounded, by a constant depending only on $\lambda(\phi)$ (and on the dimension of the space).

3.1. Transition vectors and spectrum discreteness

Let Δ be a simplex of $\mathcal{O}(\mathcal{G})$. Let's denote by e_1, \dots, e_n the directed (orbits of) edges in Δ , and denote by E_i the inverse of e_i , $i = 1, \dots, n$.

Let $g \in \text{Hyp}(\mathcal{G})$. If $X \in \Delta$, then (the conjugacy class of) g can be written as a (cyclically) reduced loop $p(g)$ in the corresponding graph of groups $\Gamma = G \setminus X$. Note that the loop corresponding to g does depend only on Δ and not on the metric of X . To any $g \in \text{Hyp}(\mathcal{G})$, we assign a transition vector (a_1, a_2, \dots, a_n) , where a_i is the number of occurrences of e_i 's and E_i 's on the loop $p(g)$.

DEFINITION 3.1. — Let Δ be a simplex of $\mathcal{O}(\mathcal{G})$ and $g \in \text{Hyp}(\mathcal{G})$. The shape of g in Δ is the transition vector of g with respect to Δ .

Remarks 3.2.

- (1) Different (conjugacy classes) of group elements may have the same shape, and if so, these group elements have the same length with respect to any $X \in \Delta$. So $\ell_X(\gamma)$ is defined for any shape γ .
- (2) There are finitely many shapes of candidates in Δ (Theorem 2.40).

(3) For all $\epsilon > 0$ and $M > 0$, the set of shapes of hyperbolic elements whose length is bounded by M for some $X \in \mathcal{O}_1(\mathcal{G}, \epsilon) \cap \Delta$ is finite. This follows by the fact that each of the coefficients of the transition vector of such a g is bounded above by $M\Lambda(X, X_0)$, where X_0 is the centre of $\Delta \cap \mathcal{O}_1(\mathcal{G})$ (the point where all the edges have the same length). As X belongs to the ϵ -thick part and has co-volume one, the distance $\Lambda(X, X_0)$ is uniformly bounded above (for instance, from $\frac{2}{\epsilon}$) and the remark follows.

LEMMA 3.3 (Compare with [12, Theorem 7.2]). — *The simplex-displacement spectrum of any \mathcal{G} -irreducible $[\phi] \in \text{Out}(\mathcal{G})$ is discrete. That is to say⁽³⁾*

$$\text{spec}(\phi) = \{\lambda_\phi(\Delta) : \Delta \text{ a simplex of } \mathcal{O}(\mathcal{G})\}$$

is a closed discrete subset of \mathbb{R} .

Proof. — Let $\lambda = \lambda(\phi)$. We will prove the claim by showing that for any $C > \lambda$, $\text{spec}(\phi) \cap [\lambda, C]$ is finite (note that $\lambda = \inf(\text{spec}(\phi))$ just by definition).

Let Δ be a simplex of $\mathcal{O}(\mathcal{G})$. For any pair of shapes (γ, η) we consider $\ell_X(\eta)/\ell_X(\gamma)$, which is a function on Δ not depending on the marking, and for any family of pairs of shapes B we define

$$F_B(X) = \sup_{(\gamma, \eta) \in B} \frac{\ell_X(\eta)}{\ell_X(\gamma)}.$$

Again, $F_B(X)$ is a function on Δ that does not depend on marking, just on B .

Since $[\phi]$ is irreducible, for any $C > \lambda$ there is $\epsilon = \epsilon(C) > 0$ such that, for any $X \in \mathcal{O}(\mathcal{G})$, if $\lambda_\phi(X) \leq C$ then X is ϵ -thick (see for instance [12, Proposition 5.5]).

Let $S_2(C)$ be the set of shapes having length bounded by $2C \text{vol}(X)$ for some X in the $\epsilon(C)$ -thick part of Δ . The set $S_2(C)$ is finite (Remark 3.2).

By candidates Theorem 2.40, there is a finite set S_1 of candidate shapes so that for any $X \in \Delta$, $\lambda_\phi(X) = \Lambda(X, X\phi)$ is realised by $\ell_X(\phi(g))/\ell_X(g)$ for some g having shape in S_1 ; moreover all such shapes have length at most $2 \text{vol}(X)$. On the other hand, for any such g , the shape of $\phi(g)$ has length bounded by $\lambda_\phi(X)\ell_X(g)$, which is bounded by $2\lambda_\phi(X) \text{vol}(X)$. That is to say, if $\lambda_\phi(X) \leq C$, then for any g with shape in S_1 , the shape of $\phi(g)$ is in $S_2(C)$. (We remark that the sets S_1 and $S_2(C)$ do not depend on the

(3) The displacement of simplices is defined in Section 2.10.

marking. That is to say, two simplices with the same unmarked underlying graph exhibit the same sets S_1 and $S_2(C)$.)

It follows that there exists a family of pairs $B \subseteq S_1 \times S_2(C)$ such that $\lambda_\phi(X) = F_B(X)$ for any $X \in \Delta$. Note that B may depend on the marking of Δ . However, since $S_1 \times S_2(C)$ is finite, there are only finitely many choices for B . It follows that the possible displacement functions on Δ run over a finite set, hence so do their minima. \square

3.2. Distance between Min-Sets of an automorphism and its inverse

LEMMA 3.4 ([13, Theorem 5.3, and Lemmas 8.4, 8.5, 8.6]). — *Given $[\psi] \in \text{Out}(\mathcal{G})$ and any $X, Y \in \mathcal{O}(\mathcal{G})$ with $\lambda_\psi(X) \geq \lambda_\psi(Y)$, there is a simplicial path from X to Y — that is to say, a sequence of adjacent simplices $\Delta_0, \Delta_1, \dots, \Delta_m$ with $X \in \Delta_0$ and $Y \in \Delta_m$ — such that there exists i_0 such that the sequence $\lambda_\psi(\Delta_i)$ is strictly monotone decreasing from 0 to i_0 , and constant from i_0 to m .*

THEOREM 3.5. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible, with $\lambda(\phi) > 1$. Then there is a D -neighbourhood (with respect to the symmetrised stretching factor, see Section 2.5) of $\text{Min}_1(\phi)$ containing $\text{Min}_1(\phi^{-1})$.*

More precisely, for any L there is a constant D (depending only on $[\phi]$ and L) such that for any volume-1 point, X with $\lambda_\phi(X) \leq L$, there is a volume-1 point, $Y \in \text{Min}(\phi^{-1})$ such that $\Lambda(X, Y)\Lambda(Y, X) < D$. In particular, for any $X \in \text{Min}(\phi)$ there is $Y \in \text{Min}(\phi^{-1})$ such that $\Lambda(X, Y)\Lambda(Y, X) < D$.

Proof. — Let $X \in \mathcal{O}_1(\mathcal{G})$ so that $\Lambda(X, X\phi) \leq L$. By Theorem 2.19, the right- and left- Lipschitz distances are comparable on the thick part. Since $[\phi]$ is irreducible, X is ϵ -thick (with ϵ depending on L but not on X , see for instance [12, Proposition 5.5]) and hence there is a constant C_1 , not depending on X , such that $\Lambda(X, X\phi^{-1}) = \Lambda(X\phi, X) < C_1$. Now we apply Lemma 3.4 with $\psi = \phi^{-1}$ and any $Y \in \text{Min}(\phi^{-1})$ (which, in particular, implies $\lambda_{\phi^{-1}}(X) \geq \lambda_{\phi^{-1}}(Y)$). Let (Δ_i) be the sequence of simplices provided by Lemma 3.4. Up to replace Y with an element of $\text{Min}(\phi^{-1}) \cap \Delta_{i_0}$, we may assume that the sequence $\lambda_{\phi^{-1}}(\Delta_i)$ is strictly monotone decreasing. By Lemma 3.3 there are only finitely many values in $\text{spec}(\phi^{-1}) \cap [\lambda(\phi^{-1}), C_1]$, whose cardinality depends only on $[\phi]$. This implies that there is a uniform bound on the length of the sequence of Δ_i 's joining X to Y . And this implies that $\Lambda(X, Y)$ is uniformly bounded depending only on $[\phi]$. Since both X and Y are ϵ -thick, for some ϵ depending only on $[\phi]$ and L , then (by Theorem 2.19) also $\Lambda(Y, X)$ is uniformly bounded. \square

4. Equivalent conditions for co-compactness of $\text{Min}(\phi)$

As customary, let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G . In this section, we discuss equivalent conditions of co-compactness of the Min -Set of \mathcal{G} -irreducible automorphisms classes $[\phi]$ with exponential growth.

There are several topologies for our deformation spaces, but our co-compactness result is in the strongest sense; we actually prove that any irreducible $[\phi]$ acts on $\text{Min}_1(\phi)$ (whence on $\text{Min}(\phi)$) with finitely many orbits of simplices, and in this sense the topology doesn't matter, since we have a fundamental domain which is compact with respect to any of the topologies.

However, our strategy is to show that the action is co-bounded with respect to the Lipschitz metric, and deduce co-compactness from there and the fact that $\text{Min}_1(\phi)$ is locally finite (Theorem 2.49). In this section we show how to make that reduction. We start with some observations about some of the topologies commonly used for our spaces.

4.1. Topology on deformation spaces

We have (among others) the equivariant Gromov topology (see for instance [30]); the length space topology (defined in Section 2.4); and the Lipschitz metric defines three topologies (see Theorem 2.20), where the basis is given by either symmetric balls, in-balls or out-balls:

(i) the symmetric or bi-Lipschitz ball of centre T and radius R :

$$B_{\text{sym}}(T, R) = \{S \in \mathcal{O}(\mathcal{G}) : \Lambda(T, S)\Lambda(S, T) \leq R\};$$

(ii) the Lipschitz out-ball of centre T and radius R :

$$B_{\text{out}}(T, S) = \{S \in \mathcal{O}(\mathcal{G}) : \Lambda(T, S) \leq R\};$$

(iii) the Lipschitz in-ball of centre T and radius R :

$$B_{\text{in}}(T, R) = \{S \in \mathcal{O}(\mathcal{G}) : \Lambda(S, T) \leq R\}.$$

Remark 4.1. — By Theorem 2.20 all three Lipschitz metrics are actually (multiplicative, asymmetric) metrics only when restricted to $\mathcal{O}_1(\mathcal{G})$. However, the three topologies are well-defined also in $\mathcal{O}(\mathcal{G})$.

Remark 4.2. — Since the Lipschitz metric is multiplicative, one should really say that the radii of these balls is $\log R$. This doesn't cause any problems in $\mathcal{O}_1(\mathcal{G})$, as the Lipschitz metric is 1 exactly when the points

are equal, and is never less than that. It does cause problems in the non-symmetric case in $\mathcal{O}(\mathcal{G})$, since non-symmetric Lipschitz metrics change with scale, so one can get any positive real number as a value for $\Lambda(T, S)$. The symmetric Lipschitz metric is a well-defined multiplicative *pseudo metric* on the whole $\mathcal{O}(\mathcal{G})$.

Let us start by proving that all topologies agree in the co-volume-one slice $\mathcal{O}_1(\mathcal{G})$.

LEMMA 4.3. — *Let $T \in \mathcal{O}_1(\mathcal{G})$. For any $\delta > 0$ there exists an $\epsilon > 0$ such that for any $S \in \mathcal{O}_1(\mathcal{G})$, if $\Lambda(T, S) \leq 1 + \epsilon$, then $\Lambda(S, T) \leq 1 + \delta$.*

Proof. — Consider an optimal map, $f: T \rightarrow S$. Then, by Lemma 2.34,

$$\text{BCC}(f) \leq \text{vol}(T) \text{Lip}(f) - \text{vol}(S) = \text{Lip}(f) - 1 = \Lambda(T, S) - 1 \leq \epsilon.$$

Hence the BCC of f is bounded above by ϵ . Let a be the length of the smallest edge in T . Now, for any edge of T , if ℓ is its length in T , and μ is how it is stretched by f , by looking at volumes, we get

$$1 = \text{vol}(S) \leq (1 + \epsilon)(1 - \ell) + \mu\ell,$$

giving

$$\mu \geq 1 - \frac{\epsilon(1 - \ell)}{\ell} \geq 1 - \frac{\epsilon(1 - a)}{a}.$$

Thus, f stretches all edges at least by $1 - \epsilon(1 - a)/a$. By Corollary 2.35, for any g ,

$$\ell_S(g) \geq \left(1 - \frac{\epsilon(1 - a)}{a}\right) \ell_T(g) - \frac{\ell_T(g)}{a} \epsilon = \ell_T(g) \left(\frac{a - 2\epsilon + a\epsilon}{a}\right),$$

where $\frac{\ell_T(g)}{a}$ is an estimate of the number of edges crossed by g in T , and the ϵ is just the above bound on $\text{BCC}(f)$. Therefore, for any $g \in \text{Hyp}(\mathcal{G})$,

$$\frac{\ell_T(g)}{\ell_S(g)} \leq \frac{a}{a - 2\epsilon + a\epsilon}.$$

As the upper bound tends to 1 as ϵ tends to 0, we have proved the result. \square

Remark 4.4. — Lemma 4.3 remains true if we replace $T, S \in \mathcal{O}_1(\mathcal{G})$ with $T, S \in \mathcal{O}(\mathcal{G})$, modified as follows: $\forall T \forall \delta \exists \epsilon : \Lambda(T, S) \text{vol}(T) < \text{vol}(S) + \epsilon \Rightarrow \Lambda(S, T) \text{vol}(S) < \text{vol}(T) + \delta$. So the lemma is basically true for trees with almost the same co-volume.

We also have the reverse.

LEMMA 4.5. — *Let $T \in \mathcal{O}_1(\mathcal{G})$. For any $\delta > 0$ there exists an $\epsilon > 0$ such that for any $S \in \mathcal{O}_1(\mathcal{G})$, if $\Lambda(S, T) \leq 1 + \epsilon$, then $\Lambda(T, S) \leq 1 + \delta$.*

Proof. — Any given T is tautologically in the thick part for some appropriate level of thickness. Next, for any g ,

$$\Lambda(S, T) \leq K \implies \ell_S(g) \geq \frac{\ell_T(g)}{K},$$

implying that if $\Lambda(S, T) \leq K$, then S will also be thick (where the thickness is divided by K). We can then invoke quasi-symmetry Theorem 2.19 to immediately get the result. \square

Remark 4.6. — As in Remark 4.4, also Lemma 4.5 remains true if we replace $T, S \in \mathcal{O}_1(\mathcal{G})$ with $T, S \in \mathcal{O}(\mathcal{G})$, by suitably modifying constants.

LEMMA 4.7. — *The following topologies on $\mathcal{O}_1(\mathcal{G})$ are the same:*

- (i) *the equivariant Gromov topology;*
- (ii) *the length function topology;*
- (iii) *the symmetric Lipschitz topology;*
- (iv) *the out-ball Lipschitz topology;*
- (v) *the in-ball Lipschitz topology.*

Proof. — By [30], the equivariant Gromov topology and the length function topology are the same. (Paulin has a standing assumption that the group is finitely generated, but this is not used for this result.) Lemmas 4.3 and 4.5 imply that all three Lipschitz topologies are the same.

Next, if we take a sub-basic open set in the length function topology, this involves picking a hyperbolic group element g , and taking all $T \in \mathcal{O}_1(\mathcal{G})$ such that $\ell_T(g)$ belongs to some open interval. Since for any g the function $\ell_T(g)$ is continuous with respect to Lipschitz metrics on $\mathcal{O}_1(\mathcal{G})$, such a set is open in the Lipschitz topology.

Conversely, by Corollary 2.42, Lipschitz out-balls are open with respect to the length function topology, and so Lipschitz-open sets are open in that topology. \square

Remark 4.8. — One can also consider other topologies. An obvious one is the path metric obtained after giving each (open) simplex in $\mathcal{O}_1(\mathcal{G})$ the Euclidean metric. This also turns out to be the same as the previous ones.

One also has the coherent topology, which is the finest topology (on $\mathcal{O}_1(\mathcal{G})$ and also $\mathcal{O}(\mathcal{G})$) which makes the inclusion maps of the simplices continuous. Care needs to be taken here, since our spaces are only a union of *open* simplices, but we can take any open simplex and add all the faces that we are allowed, then insist that these inclusions are all continuous (topologising each simplex in the standard way). This is a very different topology to the one above, and we mention it only for interest.

Remark 4.9. — We will always endow $\mathcal{O}_1(\mathcal{G})$ with the topology given by Lemma 4.7.

Now, we move to $\mathcal{O}(\mathcal{G})$ and $\mathbb{P}\mathcal{O}(\mathcal{G})$.

LEMMA 4.10 ($\mathcal{O}_1(\mathcal{G}) \simeq \mathbb{P}\mathcal{O}(\mathcal{G})$). — *Let $\mathcal{O}(\mathcal{G})$ be endowed with the bi-Lipschitz topology, and consider on $\mathbb{P}\mathcal{O}(\mathcal{G})$ the quotient topology. Then $\mathbb{P}\mathcal{O}(\mathcal{G})$ is homeomorphic to $\mathcal{O}_1(\mathcal{G})$.*

Proof. — Let $\pi: \mathcal{O}(\mathcal{G}) \rightarrow \mathbb{P}\mathcal{O}(\mathcal{G})$ be the natural projection, which is continuous by definition of quotient topology. Since $\mathcal{O}_1(\mathcal{G})$ is a sub-space of $\mathcal{O}(\mathcal{G})$, then restriction $\pi: \mathcal{O}_1(\mathcal{G}) \rightarrow \mathbb{P}\mathcal{O}(\mathcal{G})$ is continuous. Also, it is clearly bijective. It remains to prove that it is open. This is equivalent to say that for any open set $U \subseteq \mathcal{O}_1(\mathcal{G})$, the cone \mathbb{R}^+U is open in $\mathcal{O}(\mathcal{G})$ for the Lipschitz topology. Clearly it suffices to prove it when U is an open symmetric ball, say centered at Y and radius ε . But from $\Lambda(Z, Y)\Lambda(Y, Z) = \Lambda\left(\frac{Z}{\text{vol}(Z)}, Y\right)\Lambda\left(Y, \frac{Z}{\text{vol}(Z)}\right)$ we get that the symmetric ball of $\mathcal{O}(\mathcal{G})$, centered at Y and of radius ε , is just \mathbb{R}^+U . \square

Remark 4.11. — The bi-Lipschitz topology on $\mathcal{O}(\mathcal{G})$ is not Hausdorff, because the symmetric metric is only a pseudo-metric. One can naturally use on $\mathcal{O}(\mathcal{G}) = \mathcal{O}_1(\mathcal{G}) \times \mathbb{R}^+$ the product topology, which is Hausdorff, agrees with the Euclidean one on simplices, and for which $\mathbb{P}\mathcal{O}(\mathcal{G})$ is tautologically homeomorphic to $\mathcal{O}_1(\mathcal{G})$. The following lemma shows in particular that both the bi-Lipschitz and the product one are different from the length function topology.

LEMMA 4.12 ($\mathcal{O}_1(\mathcal{G}) \not\simeq \mathbb{P}\mathcal{O}(\mathcal{G})$). — *Let $\mathcal{O}(\mathcal{G})$ be endowed with the length function topology, and consider on $\mathbb{P}\mathcal{O}(\mathcal{G})$ the quotient topology. Then $\mathbb{P}\mathcal{O}(\mathcal{G})$ is not homeomorphic to $\mathcal{O}_1(\mathcal{G})$ in general. In other words, the restriction to $\mathcal{O}_1(\mathcal{G})$ of the natural projection $\pi: \mathcal{O}(\mathcal{G}) \rightarrow \mathbb{P}\mathcal{O}(\mathcal{G})$ is continuous, bijective, but in general is not open (for the projective length function topology).*

Proof. — Let X, X_n be as in Example 2.43. The points $Y_n = \frac{3}{2}X_n$ belong to $\mathcal{O}_1(\mathcal{G})$ and projectively converge, with respect to the length function topology, to X . However $\Lambda(Y_n, X) = 2$ for all n , in particular $\Lambda(Y_n, X)$ does not converge to 1. In other words, there are sets U in $\mathcal{O}_1(\mathcal{G})$ that are open for the length function topology, such that \mathbb{R}^+U is not open in $\mathcal{O}(\mathcal{G})$ (again for the length function topology). An example of such set is when U is the in-ball centered at X and of radius $11/10$. \mathbb{R}^+U does not contain any of the X_n , while any open neighbourhood of X in the length function topology contains infinitely many of them. \square

Remark 4.13. — The words “in general” in Lemma 4.12 really means “if at least one of the free factor groups is infinite”, as the used example is based only on this fact.

Remark 4.14. — Lemma 4.12 can be rephrased by saying that the quotient length function topology on $\mathbb{PO}(\mathcal{G})$ is coarser than the subspace length function topology on $\mathcal{O}_1(\mathcal{G})$.

Remark 4.15. — Example 2.43 shows that $\mathcal{O}_1(\mathcal{G})$ is not closed in $\mathcal{O}(\mathcal{G})$ with respect to the length function topology, as the co-volume of the limit of points in $\mathcal{O}_1(\mathcal{G})$ can be different from 1.

More precisely, Example 2.43 gives a sequence of points, X_n , such that $X_n \rightarrow X$ in $\mathcal{O}(\mathcal{G})$ and $\text{vol}(X_n) = 2/3$, whereas $\text{vol}(X) = 1$. By simply rescaling, we can set $Y_n = 3/2X_n$ and $Y = 3/2X$ so that $\text{vol}(Y_n) = 1$ and $\text{vol}(Y) = 3/2$. We again get that $Y_n \rightarrow Y$ in $\mathcal{O}(\mathcal{G})$ and now Y_n is a sequence in $\mathcal{O}_1(\mathcal{G})$ which converges in $\mathcal{O}(\mathcal{G})$ but whose limit does not belong to $\mathcal{O}_1(\mathcal{G})$. Hence $\mathcal{O}_1(\mathcal{G})$ is not a closed subspace of $\mathcal{O}(\mathcal{G})$.

We comment that $\mathcal{O}_1(\mathcal{G})$ will fail to be closed in $\mathcal{O}(\mathcal{G})$ if and only if some vertex group is infinite, as can be seen by appropriately tweaking the example above.

Remark 4.16. — As a consequence of Lemma 4.12, we get that in general with respect to the length function topology, the closure $\overline{\mathcal{O}_1(\mathcal{G})}$ of $\mathcal{O}_1(\mathcal{G})$, is not the same as the closure $\overline{\mathbb{PO}(\mathcal{G})}$ of $\mathbb{PO}(\mathcal{G})$. More explicitly, $\overline{\mathcal{O}_1(\mathcal{G})}$ is exactly the simplicial closure of $\mathcal{O}_1(\mathcal{G})$, which can be identified with the free splitting simplex (relative to the fixed free factor system) and it is exactly the space of edge-free actions on simplicial trees with elliptic subgroups, containing \mathcal{G} . On the other hand, in [6, 24] it is proven that $\overline{\mathbb{PO}(\mathcal{G})}$ is a compact space which contains non-simplicial trees and trees with non-trivial edge stabiliser.

Another caveat is about local compactness, as clarified by the following facts.

PROPOSITION 4.17. — *In general $\mathcal{O}_1(\mathcal{G})$ and $\mathcal{O}(\mathcal{G})$ are not locally compact.*

Proof. — We will present slightly different arguments in each case but with the same underlying goal; to produce inside of any neighbourhood of a specially chosen point, a sequence which converges to a point *outside* the space. This will show that the space is not locally compact, since our spaces are Hausdorff with respect to the length function topology.

We base our arguments on a modification of Example 2.43. Referring to Figure 2.1, we let X denote the graph in Example 2.43 which is topologically a circle of length 1. This is going to be the point which does not admit a compact neighbourhood for both $\mathcal{O}_1(\mathcal{G})$ and $\mathcal{O}(\mathcal{G})$.

Let us deal with $\mathcal{O}_1(\mathcal{G})$ first. Build points $X_{n,\epsilon}$, as in Example 2.43, by assigning length ϵ to the horizontal edge ending at the non-free vertex, and $1 - \epsilon$ to the loop.

Then, for each n , all the $X_{n,\epsilon}$ belong to the same simplex in $\mathcal{O}_1(\mathcal{G})$. Hence, for each n , $\lim_{\epsilon \rightarrow 0} X_{n,\epsilon} = X$. However, just as in Example 2.43, $\lim_{n \rightarrow \infty} X_{n,\epsilon} = (1 + \epsilon)X$. (In fact, this convergence is quite strong; for any group element the lengths with respect to the $X_{n,\epsilon}$ are eventually constant).

Let V be any neighbourhood of X in $\mathcal{O}_1(\mathcal{G})$. Since $\lim_{\epsilon \rightarrow 0} X_{n,\epsilon} = X$ there is an $\epsilon > 0$ such that, for all but finitely many n , V contains the sequence $X_{n,\epsilon}$. But now $\lim_{n \rightarrow \infty} X_{n,\epsilon} = (1 + \epsilon)X \notin \mathcal{O}_1(\mathcal{G})$, hence V cannot be compact as it admits a sequence with a limit point outside of $\mathcal{O}_1(\mathcal{G})$ (and hence no limit point in $\mathcal{O}_1(\mathcal{G})$ as our spaces are Hausdorff). Hence $\mathcal{O}_1(\mathcal{G})$ is not locally compact.

Now we deal with $\mathcal{O}(\mathcal{G})$, using the same point X . This time, define points $Y_{n,\epsilon}$ as in Example 2.43 by assigning the length ϵ to the horizontal edge and $1 - 2\epsilon$ to the loop — here $0 < \epsilon < 1/2$. (These are similar to the $X_{n,\epsilon}$ above, but the volume is no longer 1 so we are not looking at sequences in $\mathcal{O}_1(\mathcal{G})$). Each $Y_{n,\epsilon}$ is obtained from X by an isometric fold of length ϵ . This means that, for any group element g , the length of g in X will be equal to the length of g in $Y_{n,\epsilon}$ for all but finitely many n and for all ϵ .

Hence, since a basic open set in the length function topology only imposes restrictions on finitely many group elements, we have that if V is a neighbourhood of X in $\mathcal{O}(\mathcal{G})$, then V contains $Y_{N,\epsilon}$ for some fixed value of N and for all $0 < \epsilon < 1$. However, if we let $\epsilon \rightarrow 1/2$ we get an element (the loop of X) which is hyperbolic in $\mathcal{O}(\mathcal{G})$ whose length tends to 0. Thus we have obtained a sequence in V which converges to a point in $\overline{\mathcal{O}(\mathcal{G})}$ which is not in $\mathcal{O}(\mathcal{G})$. Therefore, as above, $\mathcal{O}(\mathcal{G})$ is not locally compact. \square

Finally, note that we have no hypothesis on our vertex groups, in particular G may be not countable. However, when G is countable we get the following result.

Remark 4.18. — When G is countable, bounded sets in $\overline{\mathcal{O}(\mathcal{G})}$ are sequentially pre-compact. Namely, if $(X_n) \subset \overline{\mathcal{O}(\mathcal{G})}$ is a sequence such that for any $g \in G$ there is $K(g)$ for which $\ell_{X_n}(g) \leq K(g)$, then X_n has a converging subsequence. This follows from the fact that the product of countably many closed, bounded real intervals, is sequentially compact.

4.2. Equivalent formulations of co-compactness for $\text{Min}_1(\phi)$

Let $[\phi] \in \text{Out}(\mathcal{G})$ be irreducible and with exponential growth, that is, $\lambda(\phi) > 1$. We know that in this case the simplicial structure of $\text{Min}_1(\phi)$ is locally finite (Theorem 2.49). Our aim here is to show that co-compactness of this space is actually equivalent to co-boundedness. We first recall some terminology.

DEFINITION 4.19. — *A simplicial path between $T, S \in \mathcal{O}(\mathcal{G})$ (or $\mathcal{O}_1(\mathcal{G})$) is given by:*

- (1) *a finite sequence of points $T = X_0, X_1, \dots, X_n = S$, called vertices, such that for every $i = 1, \dots, n$ there is a simplex Δ_i such that the simplices $\Delta_{X_{i-1}}$ and Δ_{X_i} are both faces (not necessarily proper) of Δ_i ;*
- (2) *Euclidean segments $\overline{X_{i-1}X_i} \subset \Delta_i$ called edges (the simplicial path is then the concatenation of these edges).*

The simplicial length of such a path is just the number n .

Remark 4.20. — The notion of *calibrated* path (with respect to ϕ) is introduced in [13]. In the case of irreducible automorphisms, this simply amounts to asking that any non-extremal vertex X_i realises $\lambda_\phi(\Delta)$ on its simplex, that is to say, $\lambda_\phi(X_i) = \lambda_\phi(\Delta(X_i))$. In particular, if $[\phi]$ is \mathcal{G} -irreducible, then any simplicial path may be assumed to be calibrated by just replacing any X_i with a suitable point in the closure of $\Delta(X_i)$ (see [13] for more details).

It is proved in [13] that any two points in $\text{Min}(\phi)$ (or $\text{Min}_1(\phi)$) can be joined by a simplicial path lying entirely within $\text{Min}(\phi)$ (or $\text{Min}_1(\phi)$, respectively). This is done by a peak-reduction argument. We use here the following quantitative version of that.

PROPOSITION 4.21 ([13, Remark 8.7]). — *Any calibrated simplicial path Σ connecting two points in $\mathcal{O}(\mathcal{G})$ can be peak-reduced to a new calibrated simplicial path, by removing a local maximum (peak) of the function λ_ϕ , via a peak-reduction surgery that increases the simplicial length of Σ by at most a uniform amount K depending only on $\text{rank}(\mathcal{G})$.*

Remark 4.22. — In [13], the authors are interested in the function λ_ϕ , which is scale invariant, and results are stated and proved for $\mathcal{O}(\mathcal{G})$, but it is readily checked that the whole peak-reduction can be carried on $\mathcal{O}_1(\mathcal{G})$.

Recall that if $[\phi]$ is \mathcal{G} -irreducible, then its simplex-displacement spectrum $\text{spec}(\phi)$ is discrete (Theorem 3.3). Thus for any $x > \lambda(\phi)$ the set $\text{spec}(\phi) \cap [\lambda(\phi), x]$ is finite.

COROLLARY 4.23. — Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) > 1$. Let $X, Y \in \mathcal{O}(\mathcal{G})$ and let $X = X_0, \dots, X_L = Y$ be a simplicial path. Let $D = \max_i \lambda_\phi(\Delta(X_i))$ and $D_0 = \max(\lambda_\phi(X), \lambda_\phi(Y))$. Let N be the cardinality of $\text{spec}(\phi) \cap [\lambda(\phi), D]$.

Then there exists a simplicial path $X = X'_0, \dots, X'_{L'} = Y$ such that $\lambda_\phi(X'_i) \leq D_0$ for any i , and such that $L' \leq L(K+1)^N$, where K is the constant of Proposition 4.21.

Proof. — By Remark 4.20 we may assume that our starting simplicial path is calibrated. List possible simplex displacements less than D , $\lambda(\phi) = \lambda_1 < \lambda_2 < \dots < \lambda_N \leq D$. To any calibrated simplicial path we assign a triple (λ_i, m, L) , where λ_i is the maximum displacement of vertices along the path, m is the number of vertices in the simplicial path which realise the maximum displacement, and L is the simplicial length. Note that $m \leq L$.

The peak reduction process (Proposition 4.21) allows us to reduce the value of m by 1, at a cost of increasing the value of L by K . Hence, after at most L peak reductions, we have transformed our simplicial path to one where the maximum peak has displacement at most λ_{i-1} . The effect on the triple is to replace it with (λ_j, m', L') , where $j < i$ and $L' \leq L + KL = L(K+1)$. Inductively, we see that we eventually arrive at a path with the requested bound on displacement, and of simplicial length at most $L(K+1)^N$. \square

Given $X \in \mathcal{O}(\mathcal{G})$, let X^c denote the “centre” of the open simplex containing X in $\mathcal{O}(\mathcal{G})$. That is, X^c has the same action as X , but the edges are rescaled to all have length 1. (Hence X^c does not have co-volume 1, but its co-volume is uniformly bounded, since there is an upper bound on the number of orbits of edges.)

The following proposition shows that symmetric Lipschitz balls can be connected via simplicial paths of uniform length, if we allow ourselves to enlarge the ball slightly.

PROPOSITION 4.24. — For all $\epsilon > 0$ there exist constants M, κ, α such that for any $T \in \mathcal{O}_1(\mathcal{G}, \epsilon)$ and all $R > 0$, any two points $S_1, S_2 \in B_{\text{sym}}(T, R) \cap \mathcal{O}_1(\mathcal{G})$ are connected by a simplicial path entirely contained in $B_{\text{sym}}(T, \kappa R^\alpha) \cap \mathcal{O}_1(\mathcal{G})$, and crossing at most MR simplices. Moreover, all points of such path are $\frac{\epsilon}{\kappa R^\alpha}$ -thick.

Proof. — It is sufficient to prove the claim when $S_2 = T$ and $S_1 = S$ is any other co-volume-one point in $B_{\text{sym}}(T, R)$. Also, we observe that since T is ϵ -thick, any point in any $B_{\text{sym}}(T, \rho)$ is ϵ/ρ -thick. In particular the last claim follows from the first one.

By triangular inequality, $\Lambda(S, T^c) \leq \Lambda(S, T)\Lambda(T, T^c)$ and since $\text{vol}(S) = \text{vol}(T) = 1$, we have $\Lambda(S, T) \leq \Lambda(S, T)\Lambda(T, S) < R$. Moreover $\Lambda(T, T^c)$ is uniformly bounded by a constant depending only on ϵ (and the maximal number of edge-orbits of trees in $\mathcal{O}(\mathcal{G})$).

So $\Lambda(S, T^c)$ is bounded by a uniform multiple of R . Let $f: S \rightarrow T^c$ be an optimal map. Subdivide the edges of S by “pulling back” the edge structure on T^c ; that is, subdivide S so that every new edge maps to an edge of T^c under f . The number of new edges created by this subdivision is at most $\Lambda(S, T^c)$ times the number of edges in S . More concretely, for each edge e in S , $\Lambda(S, T^c)$ is an upper bound for the number of edges crossed by $f(e)$ because each edge in T^c has length 1.

Now — as in [11, Definition 7.6] (see also [9, Theorem 5.6]) — construct a folding path directed by f , from S to T^c ; it is a simplicial path in $\mathcal{O}(\mathcal{G})$. The simplicial length of this path is bounded above by the number of subdivided edges of S . That is, it will be bounded above by a uniform multiple of R .

Moreover such path, say $S = X_0, X_1, \dots, X_n = T^c$, has the following properties (see [9, 11]):⁽⁴⁾

$$\Lambda(X_i, X_j) = 1 \text{ for any } 0 < i < j \leq n, \quad \Lambda(S, X_1) = \Lambda(S, T^c),$$

$$\text{vol}(X_i) \leq \text{vol}(X_1) \text{ for any } 0 < i \leq n, \quad \text{vol}(X_1) \leq \Lambda(S, T^c).$$

We now “correct” this path by tracing a simplicial path which goes through the same simplices, but whose vertices are uniformly thick and with co-volume 1. More precisely for any $i \geq 1$, we replace each X_i with $X_i^c / \text{vol}(X_i^c)$. Now, since $T^c = X_n$, we have:

$$\Lambda(X_i^c, T) \leq \Lambda(X_i^c, X_i)\Lambda(X_i, T^c)\Lambda(T^c, T) = \Lambda(X_i^c, X_i)\Lambda(T^c, T);$$

$\Lambda(X_i^c, X_i)$ is bounded by $\text{vol}(X_i) \leq \text{vol}(X_1) \leq \Lambda(S, T^c) \leq \Lambda(S, T)\Lambda(T, T^c)$. It follows

$$\begin{aligned} \Lambda\left(\frac{X_i^c}{\text{vol}(X_i^c)}, T\right) &= \text{vol}(X_i^c)\Lambda(X_i^c, T) \\ &\leq \text{vol}(X_i^c)\Lambda(S, T)\Lambda(T, T^c)\Lambda(T^c, T) \\ &\leq R\text{vol}(X_i^c)\Lambda(T, T^c)\Lambda(T^c, T). \end{aligned}$$

The factor $\Lambda(T, T^c)\Lambda(T^c, T)$ is a priori bounded by a constant depending on ϵ (see for example [14, Lemma 6.7]), and the claim follows from quasi-symmetry (Theorem 2.19) — because both T and all X_i^c are thick — and

⁽⁴⁾ We remind that the first step of such construction is to build X_1 by changing lengths to edges of S so that $\Lambda(X_1, T^c) = 1$, and all edges are maximally stretched. Then we proceed, as the name suggests, by isometrically folding edges identified by f .

from the fact that the centres X_i^c have uniformly bounded co-volume. Note that we can take $\alpha = 1 + C$ where C is the constant of Theorem 2.19. \square

We next show that $\text{Min}_1(\phi)$ is not overly distorted, in the following sense.

THEOREM 4.25. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible with $\lambda(\phi) > 1$. For any $T \in \mathcal{O}(\mathcal{G})$ there are constants C, C' , depending only $\lambda_\phi(T)$ (and on $[\phi]$ and $\text{rank}(\mathcal{G})$), such that for all $R > 0$:*

- (1) *any two points in $\text{Min}_1(\phi) \cap B_{\text{sym}}(T, R)$ are connected by a simplicial path entirely contained in $\text{Min}_1(\phi)$ (but not necessarily in $B_{\text{sym}}(T, R)$), and whose simplicial length is bounded above by CR ;*
- (2) *the ball $B_{\text{sym}}(T, R)$ intersects at most $C'R$ simplices of $\text{Min}_1(\phi)$.*

Proof. — The first claim implies in particular that if Δ_0 is a simplex intersecting $\text{Min}_1(\phi) \cap B_{\text{sym}}(T, R)$, then any other simplex with the same property stays at bounded simplicial distance from Δ_0 . Since the simplicial structure of $\text{Min}_1(\phi)$ is locally finite by Theorem 2.49, the second claim follows.

Let us now prove the first claim. Since the symmetric Lipschitz pseudo-metric is scale invariant, we may assume $\text{vol}(T) = 1$. Moreover, since $[\phi]$ is irreducible, then T is ϵ -thick for some $\epsilon > 0$ depending on $\lambda_\phi(T)$ (but not on T , see for instance [12, Proposition 5.5]).

By Proposition 4.24, any two points $S_1, S_2 \in B_{\text{sym}}(T, R) \cap \mathcal{O}_1(\mathcal{G})$ can be joined by a simplicial path of simplicial length at most MR , and lying inside $B_{\text{sym}}(T, R')$ (with $R' = \kappa R^\alpha$ and constants M, κ, α as in Proposition 4.24).

For any $S \in B_{\text{sym}}(T, R')$ and any hyperbolic group element g we have

$$\frac{\ell_S(\phi(g))}{\ell_T(\phi(g))} \leq R' \quad \text{and} \quad \frac{\ell_T(g)}{\ell_S(g)} \leq R',$$

which implies

$$\frac{\ell_S(\phi(g))}{\ell_S(g)} \leq (R')^2 \frac{\ell_T(\phi(g))}{\ell_T(g)} \implies \lambda_\phi(S) \leq (R')^2 \lambda_\phi(T).$$

Hence the displacements of points in $B_{\text{sym}}(T, R')$ are uniformly bounded by $(R')^2 \lambda_\phi(T)$.

Let N be the cardinality of $\text{spec}(\phi) \cap [\lambda(\phi), (R')^2 \lambda_\phi(T)]$. By Corollary 4.23, if $S_1, S_2 \in \text{Min}(\phi)$, they can therefore be connected by a simplicial path in $\text{Min}(\phi)$ whose length is bounded by $MR(K+1)^N$ (where K is the constant of Corollary 4.23) and by scaling volumes we can get such path in $\text{Min}_1(\phi)$. \square

Recall (see Section 2.14) that the stable or attracting tree of an $X \in \text{Min}(\phi)$ exists and is given by $X_{+\infty} = \lim_n \frac{X\phi^n}{\lambda(\phi)^n}$.

THEOREM 4.26. — Consider $\mathcal{O}_1(\mathcal{G})$ endowed with one of the equivalent topologies given by Lemma 4.7. Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) > 1$, let $X \in \text{Min}_1(\phi)$ and $X_{+\infty}$ be the stable tree. Then the following are equivalent:

- (1) $\langle \phi \rangle$ acts on $\text{Min}_1(\phi)$ with finitely many orbits of simplices;
- (2) $\text{Min}_1(\phi)/\langle \phi \rangle$ is compact;
- (3) there exists a compact set in $\mathcal{O}_1(\mathcal{G})$ whose $\langle \phi \rangle$ -orbit covers $\text{Min}_1(\phi)$;
- (4) there exists a closed symmetric Lipschitz ball B whose $\langle \phi \rangle$ -orbit covers $\text{Min}_1(\phi)$;
- (5) there exists a closed Lipschitz out-ball B whose $\langle \phi \rangle$ -orbit covers $\text{Min}_1(\phi)$;
- (6) $\exists C$ such that $\forall Y \in \text{Min}_1(\phi)$, $1 \leq \Lambda(Y, X_{+\infty}) \leq \lambda(\phi) \Rightarrow \Lambda(X, Y) \leq C$;
- (7) $\forall D_0 > 1 \ \exists C_0$ such that $\forall Y \in \text{Min}_1(\phi)$, $\frac{1}{D_0} \leq \Lambda(Y, X_{+\infty}) \leq D_0 \Rightarrow \Lambda(X, Y) \leq C_0$;
- (8) $\forall V_1 < V_2 \ \exists C'$ such that $\forall Y \in \text{Min}(\phi)$, if $V_1 \leq \text{vol}(Y) \leq V_2$ and $\Lambda(Y, X_{+\infty}) = 1$, then $\Lambda(X, Y) \leq C'$.

Proof. — Since λ_ϕ is continuous, then for every simplex Δ of $\mathcal{O}(\mathcal{G})$, the set $\text{Min}_1(\phi) \cap \Delta$ is compact. Moreover, as a consequence of Theorem 4.25, we see that (even if $\mathcal{O}_1(\mathcal{G})$ is not locally compact) $\text{Min}_1(\phi)$ is a locally compact space, whose compact subsets meet finitely many of its simplices and are contained in a closed symmetric Lipschitz ball. This gives immediately the equivalence between (1), (3) and (4).

Moreover, since ϕ acts by homeomorphisms on $\text{Min}_1(\phi)$, from local compactness we get also that (2) is equivalent to (3).

Uniform thickness of $\text{Min}_1(\phi)$ (Theorem 2.47) and quasi-symmetry (Theorem 2.19) give the equivalence of (4) and (5).

It is clear that (7) and (8) are equivalent. It is also easy to see that (7) implies (6), by taking $D_0 = \lambda(\phi)$.

We see now that (6) implies (5). Notice that $X_{+\infty}\phi = \lambda(\phi)X_{+\infty}$. Hence for any integer n (positive or negative),

$$\Lambda(Y\phi^n, X_{+\infty}) = \frac{\Lambda(Y, X_{+\infty})}{(\lambda(\phi))^n}.$$

In particular, for every $Y \in \text{Min}_1(\phi)$, there exists a n such that $1 \leq \Lambda(Y\phi^n, X_{+\infty}) \leq \lambda(\phi)$, and (5) follows from (6).

To summarise, we have that (1), (2), (3), (4) and (5) are equivalent, that (7) and (8) are equivalent, that (7) implies (6) and (6) implies (4).

Thus, in order to complete the proof, it suffices to show that (4) implies (7). Let B be a closed symmetric Lipschitz ball of radius R whose translates cover $\text{Min}_1(\phi)$. Without loss of generality, we can assume that its centre is X . Since $\Lambda(X, X_{+\infty}) = 1$, from multiplicative triangle inequality, for all Y we get the following inequalities:

$$\frac{1}{\Lambda(X, Y)} \leq \Lambda(Y, X_{+\infty}) \leq \Lambda(Y, X).$$

If also Y has co-volume 1, then we have both $\Lambda(X, Y), \Lambda(Y, X) \geq 1$ (Theorem 2.20). So for any $Y \in B \cap O_1(\mathcal{G})$ we get $1 \leq \Lambda(X, Y), \Lambda(Y, X) \leq R$, whence

$$\frac{1}{R} \leq \Lambda(Y, X_{+\infty}) \leq R.$$

Now suppose that we are given D_0 and $Y \in \text{Min}_1(\phi)$ such that

$$\frac{1}{D_0} \leq \Lambda(Y, X_{+\infty}) \leq D_0.$$

Since we know that the translates of B cover $\text{Min}_1(\phi)$, we get that $Y\phi^n \in B$ for some integer n . Hence, for this n ,

$$\frac{1}{R} \leq \Lambda(Y\phi^n, X_{+\infty}) = \frac{\Lambda(Y, X_{+\infty})}{(\lambda(\phi))^n} \leq R.$$

Therefore,

$$\frac{1}{RD_0} \leq (\lambda(\phi))^n \leq RD_0,$$

and we get a bound on $|n|$ depending only on D_0 . But now

$$\Lambda(X, Y) \leq \Lambda(X, Y\phi^n)\Lambda(Y\phi^n, Y) \leq R\Lambda(Y\phi^n, Y)$$

where $\Lambda(X, Y\phi^n) \leq R$ follows since $Y\phi^n \in B$. The following claim will conclude the proof. \square

CLAIM. — $\Lambda(Y\phi^n, Y) \leq \max\{(\lambda(\phi))^{|n|}, D(\lambda(\phi^{-1}))^{|n|}\}$, where D is the constant from Theorem 3.5.

Proof of the claim. — If n is negative, since $Y \in \text{Min}_1(\phi)$ we have

$$\Lambda(Y\phi^n, Y) = \Lambda(Y, Y\phi^{-n}) = (\lambda(\phi))^{|n|}.$$

Whereas, if n is positive, then by Theorem 3.5 Y is uniformly close a point $Z \in \text{Min}_{\phi^{-1}}$ and hence

$$\begin{aligned} \Lambda(Y\phi^n, Y) &= \Lambda(Y, Y\phi^{-n}) \\ &\leq \Lambda(Y, Z)\Lambda(Z, Z\phi^{-n})\Lambda(Z\phi^{-n}, Y\phi^{-n}) \\ &= \Lambda(Y, Z)\Lambda(Z, Z\phi^{-n})\Lambda(Z, Y) \\ &\leq D(\lambda(\phi^{-1}))^n. \end{aligned} \quad \square$$

5. North-south dynamics for primitive irreducible automorphisms

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G , where in addition we require that $\text{rank}(\mathcal{G}) \geq 3$.

5.1. Statement of north-south dynamics

The so-called north-south dynamics for iwip automorphisms in the classical Culler–Vogtmann outer space CV_n is a well-known fact established in [27], and generalised in [21] in greater generality (in [21] the hypothesis $\text{rank}(\mathcal{G}) \geq 3$ is required). Here we need a north-south dynamics result for our case, which is slightly more general.

The proof of Theorem 5.1 below is essentially exactly the same as the proof of [21, Theorem C], where the author assumes that G is free and that the automorphism is iwip (or fully irreducible). However, her proof applies for general groups, where some missing point can be filled using results of [18]. Finally, we note that the iwip property is not really used anywhere in the proof. More specifically, what is needed is an irreducible automorphism which can be represented by some simplicial train track map with primitive transition matrix (whence every simplicial train track representative of $[\phi]$ has this property). Irreducibility is needed because it implies connectedness of local Whitehead graphs — which is in fact used — but irreducibility of powers is not used.

For these reasons, we decided to not include all details of the proof here, but for the rest of the section, we just mention the main steps of the proof of [21, Theorem C], and we explain why the proofs of relevant statements still hold on our context, by giving appropriate references when needed.

THEOREM 5.1. — *Any \mathcal{G} -primitive irreducible $[\phi] \in \text{Out}(\mathcal{G})$ acts on $\mathcal{O}(\mathcal{G})$ with projectivised uniform north-south dynamics: there are two fixed projective classes of trees $[T_\phi^+]$ and $[T_\phi^-]$, such that for every compact set K of $\overline{\mathcal{PO}(\mathcal{G})}$ that does not contain $[T_\phi^-]$ (resp. $[T_\phi^+]$) and for every open neighbourhood U (resp. V) of $[T_\phi^+]$ (resp. $[T_\phi^-]$), there exists an $N \geq 1$ such that for all $n \geq N$ we have $(K)\phi^n \subseteq U$ (resp. $(K)\phi^{-n} \subseteq V$).*

To begin with, we define T_ϕ^\pm : they are the attracting and repelling tree (Definition 2.62) of a given train track point with primitive transition matrix, which exists by hypothesis. The proof of Theorem 5.1 starts with pointwise convergence, that we state separately in Theorem 5.2 and whose

proof is contained in the proof of [21, Theorem C], even if it is not written as a separate statement there (compare also with Proposition 2.65).

THEOREM 5.2. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -primitive irreducible. For any $X \in \overline{\mathcal{O}(\mathcal{G})}$, $X \notin [T_\phi^-]$ (resp. $X \notin [T_\phi^+]$), we have that for $n \rightarrow \infty$:*

$$\begin{aligned} \frac{X\phi^n}{\lambda(\phi)^n} &\longrightarrow cT_\phi^+, \text{ for some } c > 0 \\ (\text{resp. } \frac{X\phi^{-n}}{\lambda(\phi^{-1})^n} &\longrightarrow dT_\phi^-, \text{ for some } d > 0). \end{aligned}$$

Before going into the proof of Theorem 5.2 we show how it allows to extend the results of Section 2.14 to general trees.

DEFINITION 5.3. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -primitive irreducible. Let $X \in \overline{\mathcal{O}(\mathcal{G})}$, which does not belong to the projective class $[T_\phi^+]$ (resp. $[T_\phi^-]$). We define the tree:*

$$X_{+\infty} = \lim_n \frac{X\phi^n}{\lambda(\phi)^n} \quad (\text{resp. } X_{-\infty} = \lim_n \frac{X\phi^{-n}}{\lambda(\phi^{-1})^n}).$$

It is very important to note here that these limits do exist and they are \mathbb{R} -trees because of the previous Theorem 5.2. Note that this extends (when $[\phi]$ is \mathcal{G} -primitive) the definition given in Section 2.14 for points $X \in \text{Min}(\phi)$.

The assumption that our irreducible automorphism is primitive is crucial in order to apply Theorem 5.2. For a general irreducible automorphism (without the extra assumption of the primitive property), we cannot ensure that the limits do exist for general points of $\overline{\mathcal{O}(\mathcal{G})}$, but only for points of $\text{Min}(\phi)$, and we also may lose uniqueness.

Remark 5.4. — Given a point $X \in \overline{\mathcal{O}(\mathcal{G})}$ which does not belong to $[T_\phi^+]$ (resp. to $[T_\phi^-]$), by Theorem 5.2 we have $X_{+\infty} = cT_\phi^+$ (resp. $X_{-\infty} = dT_\phi^-$) for some $c > 0$ (resp. $d > 0$). If $X \in \text{Min}(\phi)$ (resp. $X \in \text{Min}(\phi^{-1})$), then by continuity of stretching factor (Corollary 2.42) we have

$$\begin{aligned} 1 = \Lambda\left(X, \frac{X\phi^n}{\lambda(\phi)^n}\right) &\longrightarrow \Lambda(X, cT_\phi^+) = c\Lambda(X, T_\phi^+) \\ (\text{resp. } 1 = d\Lambda(X, T_\phi^-)). \end{aligned}$$

(Compare also Proposition 2.65, where the existence of T_ϕ^+ is not used.)

We are now ready to analyse the main steps of the proof of [21, Theorem C], and adapt them to prove Theorems 5.1 and 5.2.

5.2. The attracting tree does not depend on the chosen train track

A key step for proving that the attracting tree does not depend on the chosen train track is the following proposition.

PROPOSITION 5.5. — *Let $T \in \overline{\mathcal{O}(\mathcal{G})}$. Suppose there exists a tree $T_0 \in \mathcal{O}(\mathcal{G})$, an equivariant map $h: T_0 \rightarrow T$, and a bi-infinite geodesic γ_0 of T_0 , representing a generic leaf γ of Λ_ϕ^+ , such that $h(\gamma_0)$ has diameter greater $2 \text{BCC}(h)$. Then:*

- (1) $h(\gamma_0)$ has infinite diameter in T ;
- (2) there exists a neighbourhood V of T such that $(V)\phi^p$ converges to T_ϕ^+ , uniformly as $p \rightarrow \infty$.

Proof. — The proof goes as in the classical case (see [2, Lemma 3.4] and [27, Proposition 6.1]) and it is the same as that of [21, Proposition 3.3]. The iwip property is used there, just in order to ensure the existence of a train track representative of the automorphism with primitive transition matrix. The rest of the argument uses this primitive matrix in applying Perron–Frobenius theory. We conclude that our assumption that the automorphism is \mathcal{G} -primitive is enough. Note that the assumption in [21] that the group G is free is never used for the proof. \square

5.3. Infinite-index subgroups do not carry the attracting lamination

Another key step in the proof [21, Theorem C] is the following.

LEMMA 5.6. — *Let $T \in \mathcal{O}(\mathcal{G})$ and let $f: T \rightarrow T$ be a train track representative of a \mathcal{G} -primitive irreducible $[\phi] \in \text{Out}(\mathcal{G})$. Let C be a subgroup of G such that for every $[G_i] \in [\mathcal{G}]$, either $C \cap G_i$ is trivial or equal to G_i up to conjugation. Suppose moreover that $[\mathcal{G}]$ induces on C a free decomposition of finite rank. If C carries Λ_ϕ^+ , then C has finite index in G .*

Proof. — Again, no particular patch is needed, and the proof is exactly the same as that of [21, Lemma 3.9(c)]. It relies on the fact that there is one (and so every) leaf of the lamination which crosses (the orbit of) every edge. Here we use the fact that the relative Whitehead graphs are connected at each vertex, and this can be ensured under our assumptions because $[\phi]$ is \mathcal{G} -irreducible. The fact that the group G is free is not used for this proof at all. \square

5.4. Q -map and dual laminations of trees

In this section we give the definition and some results about the statements on dual laminations of trees, which are well-known for free groups and they have been recently generalised for the context of free products in [18].

PROPOSITION 5.7 ([18, Lemma 4.18]). — *Let $T \in \overline{\mathcal{O}(\mathcal{G})}$ be a minimal \mathcal{G} -tree with dense orbits and trivial arc stabilisers. Given $\epsilon > 0$, there exists a tree $T_0 \in \mathcal{O}(\mathcal{G})$ with co-volume $\text{vol}(T_0) < \epsilon$, and an equivariant map $h: T_0 \rightarrow T$ whose restriction to each edge is isometric, and with $\text{BCC}(h) < \epsilon$.*

The so-called Q -map, which was defined in [27] for free groups, can also be generalised for general free products. Any $X \in \partial_\infty(G, \mathcal{G})$ can be represented as the “point at infinity” of a half-line in a \mathcal{G} -tree $T \in \mathcal{O}(\mathcal{G})$. Almost the same happens for trees $T \in \overline{\mathcal{O}(\mathcal{G})}$, the difference is that in this case, the path representing X in T could have finite length. If this happens, X is called T -bounded.

PROPOSITION 5.8 ([18, Q -map, Proposition 6.2], [27, Proposition 3.1]). — *Let $T \in \overline{\mathcal{O}(\mathcal{G})}$ be a minimal \mathcal{G} -tree with dense orbits and trivial arc stabilisers. Suppose $X \in \partial_\infty(G, \mathcal{G})$ is T -bounded. Then there is a unique point $Q_T(X) \in \overline{T}$ such that for any $T_0 \in \mathcal{O}(\mathcal{G})$, any half-line ρ representing X in T_0 , and equivariant map $h: T_0 \rightarrow T$, the point $Q_T(X)$ belongs to the closure of $h(\rho)$ in T . Also, every $h(\rho)$ is contained in a $2\text{BCC}(h)$ -ball centered at $Q_T(X)$, except for an initial part.*

DEFINITION 5.9. — *Let $T \in \overline{\mathcal{O}(G)}$. We define the following.*

- (1) *The algebraic lamination dual to the tree T is defined as $L(T) = \bigcap_{\epsilon > 0} L_\epsilon(T)$ where $L_\epsilon(T)$ is the closure of set of pairs $(g^{-\infty}, g^\infty)$ where $\ell_T(g) < \epsilon$ and g does not belong to some free factor of $[\mathcal{G}]$.*
- (2) *Let’s further assume that T has dense orbits. We define $L_Q(T) = \{(X, X') : Q_T(X) = Q_T(X')\} \subset \partial^2(G, \mathcal{G})$.*

These definitions are equivalent, in the case of trees with dense orbits, by [18, Proposition 6.10]. Moreover, in [21, Remark 3.1] is shown that the leaves of $L_Q(T_\phi^-)$ are either leaves of Λ_ϕ^+ or concatenation of two rays, based at a non-free vertex, obtained as iterated images of an edge via a train track map. The latter are called *diagonal leaves* (and do not arise in the classical case).

PROPOSITION 5.10 ([21, Proposition 3.22]). — *If $Y \in \overline{\mathcal{O}(\mathcal{G})}$ is a minimal \mathcal{G} -tree with dense orbits and trivial arc stabilisers, then at least one of the following is true:*

- (1) *there exists a generic leaf (X, X') of Λ_ϕ^+ or of Λ_ϕ^- such that $Q_Y(X) \neq Q_Y(X')$;*
- (2) *there exists a diagonal leaf (i.e. the concatenation of two half-lines) (X, X') of $L_Q(T_\phi^-)$ or $L_Q(T_\phi^+)$ such that $Q_Y(X) \neq Q_Y(X')$.*

Proof. — The proof is again the same as that of [21, Proposition 3.22], using the generalised version of the Q -map given in [18] (Propositions 5.8 and 5.7). All the intermediate steps still hold in our context. Connectedness of Whitehead graphs is used here, which is safe because we are assuming $[\phi]$ is \mathcal{G} -irreducible. \square

We also need to ensure that limit trees have dense orbit, but this is already part of the literature.

LEMMA 5.11 ([21, Lemma 4.5]). — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -primitive irreducible. Then the trees T_ϕ^+ and T_ϕ^- have dense orbits.*

5.5. At least one of the laminations is long in any tree of the boundary

The key lemma here is the following.

LEMMA 5.12 ([21, Lemma 3.26]). — *Let $T \in \overline{\mathcal{O}(\mathcal{G})}$. Then there exists a tree $T_0 \in \mathcal{O}(\mathcal{G})$, an equivariant map $h: T_0 \rightarrow T$, and a bi-infinite geodesic γ_0 representing a generic leaf of Λ_ϕ^+ or of Λ_ϕ^- , such that $h(\gamma_0)$ has diameter greater than $2\text{BCC}(h)$.*

Proof. — As in the proof of [21, Lemma 3.26] (see also [27]), we distinguish three cases. We just give a sketch of the proof for each case and we refer to the original proof for the details.

- Suppose that T has dense orbits. First, we note that arc stabilisers of T are trivial (this is true by [24, Proposition 5.17]). In this case, the conclusion is a consequence of Propositions 5.10 and 5.7 exactly as in the proof of [21, Lemma 3.26].
- Suppose that T does not have dense orbits and that it is not simplicial. This sub-case can be reduced to the first case (of a tree with dense orbits), by collapsing the simplicial part, exactly as in [21].

- Suppose that T is simplicial. In this case, we have to show that a generic leaf of the attracting lamination cannot be contained in the boundary ∂B of some vertex stabiliser B in T . In other words, we want to prove that the lamination is not contained in any vertex stabiliser of a (non-trivial) tree in the boundary of $\mathcal{O}(\mathcal{G})$. By [24, Corollary 5.5], point stabilisers of trees in the boundary have finite rank and, more specifically, their rank is bounded above by $\text{rank}(\mathcal{G})$. It follows that they have infinite index and so they cannot carry the lamination, by 5.6. \square

5.6. Proof of Theorem 5.1

Everything flows as in the proof of [21, Theorem C]. The point-wise convergence of Theorem 5.2 follows directly from Proposition 5.5 and Lemma 5.12. The locally uniform convergence then follows, because of the compactness of $\overline{\mathbb{PO}(\mathcal{G})}$.

6. Discreteness of the product of limit trees of an irreducible automorphism

6.1. Dynamics of train track maps

Let $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ be a free factor system of a group G .

In this section we prove the discreteness of the G -action on the product of the two limit trees of irreducible automorphisms with exponential growth. We do not assume primitivity here, so powers of the automorphism may be reducible. Similar results in the free case have been proved in [2] and in the free product case in [7].

In particular, both those papers have a precise analogue of Proposition 6.4; the argument in [7], which also deals with free products, relies on a technical hypothesis of no twinned subgroups. Effectively, this allows that paper to argue that the “angles” (the vertex group elements one encounters) remain bounded, and hence one observes similar behaviour to that seen in [2]. However, we obtain finiteness conditions in a slightly different way by observing that there are finitely many orbits of paths which occur as the train track image of an edge. However, while this idea is straightforward, it is somewhat more difficult to implement.

We also observe that a version of Theorem 6.17, in the free group case, is proved in [2], but in a slightly different way. There, the main argument deals with the case where there is no “closed INP”, whose analogue is that no \mathcal{G} -hyperbolic element becomes elliptic in the limit tree. (A version of Theorem 6.17 is also proved in [7], again with the same technical assumption of no twinned subgroups.) The other case — where there is a closed INP — is dealt with in [2] via surface theory.

Here we need to face the fact that our deformation spaces have a locally infinite simplicial structure. For this reason we need different arguments. However, our approach allows us to deal with both the above cases at the same time. We recall that we are using the square-bracket notation for reduced paths (see the start of Section 2.8).

DEFINITION 6.1. — *Let $f: X \rightarrow X$ be a train track map representing some $[\phi] \in \text{Out}(\mathcal{G})$. Let L be a periodic line in X . The number of turns of L is the number of turns appearing in a fundamental domain. We say that L splits as a concatenation of paths, if we can write a fundamental domain of L as $\rho_1 \cdots \rho_n$ such that for any i we have*

$$[f^i(\rho_1 \cdots \rho_n)] = [f^i(\rho_1)] \cdots [f^i(\rho_n)]$$

as a cyclically reduced path.

DEFINITION 6.2. — *Let $f: X \rightarrow X$ be a train track map representing some $[\phi] \in \text{Out}(\mathcal{G})$. An f -piece, or simply a piece is an edge path p which appears as subpath of $f(e)$, with e edge, or $f(e_1 e_2)$ with e_i edges meeting at a legal free turn (i.e. a turn at a free vertex).*

DEFINITION 6.3. — *For a not necessarily simplicial path p in a simplicial tree, its simplicial closure is the smallest simplicial path containing p . In other words, the simplicial closure of p is obtained by prolonging the extremities of p till the next vertex.*

We recall that we defined the critical constant $\text{cc}(f)$ of a map f in Definition 2.36, and Nielsen paths in Definition 2.39.

PROPOSITION 6.4. — *Let $f: X \rightarrow X$ be a simplicial train track map representing some $[\phi] \in \text{Out}(\mathcal{G})$, with $\text{Lip}(f) = \lambda > 1$. Let $C = \text{cc}(f) + 1$. Then there exist explicit positive constants $N, M \in \mathbb{N}$ (with $M = 5N^2 + N$), such that for any finite edge path, or periodic line L in X , one of the following holds true:*

- (1) $[f^M(L)]$ has less illegal turns than L ;
- (2) $[f^M(L)]$ has a legal subpath of length more than C ;

(3) L splits (not necessarily at vertex-points) as a concatenation of paths $\rho_1 \cdots \rho_k$ so that each ρ_i is a pre-periodic Nielsen path with at most one illegal turn.

Moreover, the periodic behaviour of ppNp's starts before N iterates, and with period less than N .

The same conclusion holds true for finite paths whose endpoints are not necessarily vertices, with (3) above replaced by:

(3') there exists L' , contained in the simplicial closure of L , such that:

- (i) L' splits (not necessarily at vertex-points) as a concatenation of paths $\rho_1 \cdots \rho_k$ so that each ρ_i is pre-periodic Nielsen path with at most one illegal turn;
- (ii) endpoints of L are at distance less than C/λ^M to those of L' ;
- (iii) if the initial point x of L is in L' , then $f^M(x)$ is in the same edge as the image of the initial point of L' (with a corresponding statement for terminal points).

Proof. — Since powers of train track maps are also train track, by replacing f with some power, we may freely assume that $l_X(f(e)) > C$ for any edge. (Note that if f represents an irreducible $[\phi]$, then f^n represents $[\phi^n]$ which might not be irreducible. However, it will still be the case that f^n is a train track map). We now set constants (whose role will become clear along the proof):

- M_0 is one plus the number of orbit of pieces (which is finite);
- $m = (M_0^2 + M_0)^2$;
- Q_m is the number of orbit of turns at non-free vertices that appears in iterates $f^n(p)$ where p runs over the set of pieces, and $n = 1, \dots, m$ (Q_m is a finite number);
- $N_0 = m(Q_m + 2)^2 + 1$;
- $N = mN_0$;
- $M = 5N^2 + N$.

We give the proof in case L is a finite edge path, by analysing what happen to maximal legal subpaths of L ; the case where L is a periodic line follows by applying our reasoning to a fundamental domain (and make cyclic reductions). Also, note that the case where L is legal easily reduces to (3') with no ppNp appearing (hence L is just in the neighbourhood of a point), so we may assume L contains at least one illegal turn.

In case L is not simplicial, we refer to non-simplicial maximal legal subpaths of L (that possibly arise only at its extremities), as *tails*.

First, we observe that if (1) holds for some iterate f^n with $n \leq M$, then by train track properties, it holds also for f^M . The same is true for (2) by Lemma 2.38.

Now, we suppose that we have a path L for which (1) fails (in particular it fails for any $n \leq M$). Then the number of illegal turns in $f^n(L)$ remains constant. This implies that, in calculating $[f(L)]$, we apply f to each maximal legal subpath of L , then cancel, and we are assured that some portion of the image of that path survives, and that the new turns formed after cancellations are illegal ones.

For any α maximal legal subpath of L , which is not a tail, we denote by α_n the corresponding maximal subpath in $[f^n(L)]$, i.e. the portion of $f^n(\alpha)$ that survives after cancellations (for $1 \leq n \leq M$). If α is a tail, then we define α_n to be the simplicial closure of the surviving portion. So α_n is a simplicial path in any case. Note that since f is simplicial and expanding, then

$$f^{-1}(\alpha_n) \subseteq \alpha_{n-1}$$

also in case of tails.

Now we assume that also (2) fails, and prove that in that case (3) is true. Since f -images of edges are longer than C , the f -preimages of legal paths we see in $[f^n(L)]$ (for $1 \leq n \leq M$) cross at most two edges. In particular any maximal legal subpath of $[f^n(L)]$ consists of at most two pieces. Note that a legal path may a priori be divided in pieces in different ways. Here we consider the subdivision of α_n given by $f^{-1}(\alpha_n)$. Note that from the definition of piece it follows that if α_n consists of two pieces, then they meet at a non-free vertex.

For each $1 \leq n \leq M$, and any α_n maximal legal subpath of $[f^n(L)]$, we define $\text{surv}(\alpha_n) = f^{-M+n}(\alpha_M)$ (the portion of α_n that survives all M iterates).

Any such α_n therefore splits in three (not-necessarily simplicial) subpaths

$$\alpha_n = \text{left}(\alpha_n) \text{ surv}(\alpha_n) \text{ right}(\alpha_n).$$

Remark 6.5. — We observe the following.

- (1) Since f -images of edges have length more than C , then any $\text{surv}(\alpha_n)$ contains at most one vertex.
- (2) If $\text{surv}(\alpha_n)$ contains a vertex v , then $\text{surv}(\alpha_{n+j}) = f^j(\text{surv}(\alpha_n))$ contains the vertex $f^j(v)$.
- (3) If $\alpha\beta$ are consecutive maximal legal subpaths, hence forming an illegal turn, then cancellations between $f(\alpha_n)$ and $f(\beta_n)$ occur in the subpath $f(\text{right}(\alpha_n) \text{ left}(\beta_n))$.

- (4) $\text{surv}(\alpha_n)$ is never involved in cancellations.
- (5) $\text{right}(\alpha_n)$ and $\text{left}(\alpha_n)$ are eventually cancelled by f^M , unless α_n is a tail.
- (6) If α is a non-tail extremal maximal legal subpath of L , say on the left side (the start of L), then $\text{left}(\alpha_n)$ is empty, because no cancellations occur on its left side (same for the end of L).
- (7) If α is a tail, say on the left side, then $\text{left}(\alpha_n)$ is just the portion of the edge containing the beginning of $[f^n(L)]$, but which is not in $[f^n(L)]$.

Next we focus our attention on iterates till N . Pick two consecutive such maximal legal subpaths α, β and look at α_n, β_n (for $1 \leq n \leq N$).

CLAIM 6.6. — *There exist $1 \leq s < t \leq N$ and points $a_s \in \alpha_s$, $a_t \in \alpha_t$, $b_s \in \beta_s$, $b_t \in \beta_t$, such that*

- $f^{t-s}(a_s) = a_t, f^{t-s}(b_s) = b_t$ (hence $[a_t, b_t] = [f^{t-s}([a_s, b_s])]$);
- there is $h \in G$ so that $[a_t, b_t] = h[a_s, b_s]$ (so $[a_s, b_s]$ is a pNp of period $t-s < N$, containing a single illegal turn: that formed at the concatenation point of $\alpha_s \beta_s$);
- a_t is the unique fixed point of the restriction to α_t of hf^{s-t} ; b_t is the unique fixed point of the restriction to β_t of hf^{s-t} .
- a_s is not internal to $\text{right}(\alpha_s)$ and b_s is not internal to $\text{left}(\beta_s)$.

Proof. — The proof is based on pigeon hole principle. As mentioned, any α_n consists of either one or two pieces. In case α_n consists of two pieces, we denote by v_n the non-free vertex separating the pieces of α_n , and similarly we define w_n as the vertex separating the pieces of β_n , if any.

By definition of constants, we have $M_0 - 1$ orbits of pieces. So the possible configurations of orbit of pieces that we read in a maximal legal subpath are less than $M_0^2 + M_0$. Consequently, the configuration of orbit pieces that we read in paths $\sigma_n = \alpha_n \beta_n$, runs over a set of cardinality strictly less than $m = (M_0^2 + M_0)^2$. Let \mathcal{T} be the set of orbit of turns at non-free vertices that appears in iterates of pieces up to power m (the cardinality of this family is Q_m by definition).

Now, we subdivide the family $\Sigma = \{\sigma_n, 1 \leq n \leq N\}$ in N_0 subfamilies Σ_ν each made of m consecutive elements. By pigeon hole principle any such Σ_ν contains a pair of paths σ_i, σ_j (with $i < j$) with the same configuration of orbit of pieces. To any such pair we associate a tag $(\text{Conf}, \text{Turn}_\alpha, \text{Turn}_\beta)$ as follows: Conf is just the configuration of orbit of pieces. We define now Turn_α , the other being defined in the same way.

- $\text{Turn}_\alpha = 1$ if α_i consists of a single piece.

- $\text{Turn}_\alpha = \text{Per}$ if α_i is made of two pieces and $f^{j-i}(v_i) = v_j$ (Per is just a label).
- Finally, if none of the above occur, we set Turn_α to be the orbit of τ_α , the turn we read in α_j at v_j . Note that if α_i consists of two pieces and $v_j \neq f^{j-i}(v_i)$, then τ_α belongs to \mathcal{T} .

In the last case, the possibilities for the orbit of τ_α are at most Q_m (same for τ_β). It follows that the cardinality of the set of possible tags $(\text{Conf}, \text{Turn}_\alpha, \text{Turn}_\beta)$ is $m(Q_m + 2)^2$. Since $N_0 = m(Q_m + 2)^2 + 1$ we must have at least one repetition. That is to say, we find two pairs $(\sigma_{i_0}, \sigma_{j_0})$, $(\sigma_{i_1}, \sigma_{j_1})$ (with $1 \leq i_0 \leq j_0 < i_1 < j_1 \leq N$) with same tags. Now we have three cases.

Case 1 (Both Turn_α and Turn_β are different from Per). — In this case we set $s = j_0$ and $t = j_1$. Let's focus on α -paths. If $\text{Turn}_\alpha = 1$ then α_s and α_t both consist of single pieces, and in the same orbit. If $\text{Turn}_\alpha = \tau_\alpha$, then α_s and α_t both consist of two pieces in the same respective orbit, and whose middle turns are also in the same orbit. So, also in this case we have that α_s and α_t are in the same orbit. The same reasoning shows that β_s and β_t are in the same orbit.

Thus, there exist $h, h' \in G$ such that $\alpha_t = h\alpha_s$ and $\beta_t = h'\beta_s$. Both turns we read at (the concatenation points of) $\alpha_s\beta_s$ and $\alpha_t\beta_t$ are illegal. Since legality of turns is invariant under the action of G , we have that the turn we read in $(h\alpha)(h\beta)$ is illegal. On the other hand the illegal turn at $\alpha_t\beta_t$ is $(h\alpha_s)(h'\beta_s)$. This forces $h = h'$, and in particular the whole path $\alpha_s\beta_s$ is in the same orbit of $\alpha_t\beta_t$.

Now, we set a_t to be the unique fixed point of the restriction to α_t of the contraction hf^{s-t} and set $a_s = f^{s-t}(a_t)$. Similarly we define b_t and b_s . Thus $a_t = f^{t-s}(a_s) = ha_s$, and the same holds for b -points.

Case 2 ($\text{Turn}_\alpha = \text{Turn}_\beta = \text{Per}$). — In this case we set $s = i_0, t = j_0$ (note that this choice is different from that of Case 1). The paths $[v_s, w_s]$ and $[v_t, w_t]$ both consist of two pieces in the same respective orbit, meeting at illegal turns. As in Case 1, we deduce that in fact the whole $[v_s, w_s]$ is in the same orbit of $[v_t, w_t]$. In this case we set $a_s = v_s, a_t = v_t, b_s = w_s, b_t = w_t$. Note that if $[v_t, w_t] = h[v_s, w_s]$ then v_t is the unique fixed point of the restriction to α_t of hf^{t-s} (and similarly for w_t).

Case 3 (One of Turn's, say Turn_α , is Per and the other, Turn_β , is different). — In this case we set $s = i_0, t = j_0$ (as in Case 2). As above, we see that there is $h \in G$ so that the concatenation of the right-side piece of α_t with the left-side piece of β_t is the h -translate of the concatenation of corresponding pieces in α_s, β_s . If $\text{Turn}_\beta = 1$, then as above we see

that $\beta_t = h\beta_s$, and we define b_t as the unique fixed point of hf^{s-t} in β_t , $b_s = f^{s-t}(b_t)$, $a_s = v_s$, and $a_t = v_t$.

So we are left with the case $\text{Turn}_\beta = \tau_\beta$. Let $\tau_s = (e, e')$ be the turn that we read in β_s at w_s , and let τ_t be the turn we read at w_t . Since the configurations of pieces are the same at iterates s, t , we know that there is $h' \in G$ so that $\tau_t = (he, h'e')$ (note that $h^{-1}h'$ is in the stabiliser of w_s). Now we define $H: \beta_s \rightarrow \beta_t$ to be h on the left-side piece, and h' on the right-side one; and set b_t to be the unique fixed point of contraction $Hf^{s-t}: \beta_t \rightarrow \beta_t$, and $b_s = f^{s-t}(b_t)$.

In order to have $[a_t, b_t] = h[a_s, b_s]$, we have to prove that if w_s is in $[a_s, b_s]$, that is to say if w_s is on the left side of b_s , then $h = h'$. In this case, since $f^{t-s}(w_s) \neq w_t$ and since f is expanding, then $f^{t-s}(w_s)$ is on the left side of w_t , possibly on the cancelled region. Now we iterate f for $(t-s)$ more times (note that since $t = j_0$ we have enough room to iterate $(t-s)$ times).

If $f^{t-s}(w_s)$ is in β_t (that is to say, it is not in a cancelled region), then $f^{t-s}(w_t)$ is in $\beta_{t+(t-s)}$, and from $[v_t, w_t]$ to $[v_{t+(t-s)}, w_{t+(t-s)}]$ we see the same cancellations we had from $[v_s, w_s]$ to $[v_t, w_t]$. It follows that $f^{t-s}(\tau_t)$ is in the same orbit of $f^{t-s}(\tau_s)$ and this forces $h = h'$.

Similarly, if $f^{t-s}(w_s)$ is cancelled, then $f^{t-s}(w_t)$ must also be cancelled — overlapping a turn in the image of α_t being in the same orbit as $f^{t-s}(\tau_s)$ — otherwise the turn we read at concatenation point of $\alpha_{t+(t-s)}\beta_{t+(t-s)}$ would become legal, contradicting the fact that the number of illegal turns stay constant (and cancellations on the right side of β and on the left side of α never touch the illegal turn between α and β , see Remark 6.5(4)). Again, $f^{t-s}(\tau_t)$ and $f^{t-s}(\tau_s)$ are in the same orbit and thus $h = h'$.

In all three cases, we proved the first three properties. We check now the last one. We prove that a_s is not in the interior of $\text{right}(\alpha_s)$, the same reasoning proving that b_s is not in the interior of $\text{left}(\beta_s)$.

We already proved that $[a_s, b_s]$ is a pNp. A priori a_s could belong to $\text{right}(\alpha_s)$. If $b_s \in \text{left}(\beta_s) \cup \text{surv}(\beta_s)$ then it is clear (Remark 6.5(3)) that the cancellations we see in subsequent iterations are the same we see from $[a_s, b_s]$ to $[a_t, b_t]$ and in particular a_n is never cancelled, so a posteriori a_s would belong to $\text{surv}(\alpha_s)$. But now note that the very same holds true also if $b_s \in \text{right}(\beta_s)$. Indeed, in this case $f^n(b_s)$ may, a priori, eventually disappear from β_{s+n} ; but still, cancellations with α -paths arise in a subpath which is in the image of $f^{s+n}[a_s, b_s]$ because b_s is on the right side of $\text{surv}(\beta_s)$ which is never involved in cancellations. The claim is proved. \square

If $\alpha\beta$ are as in the claim, then by pulling back a_s, b_s to (the simplicial closures of) α, β we find a ppNp in (the simplicial closure of) L . We set $a = f^{-s}(a_s)$ and $a_n = f^n(a_s)$ for any $1 \leq n \leq M$. Similarly we define b -points. The path $[a_n, b_n]$ evolves till $n = s < N$, then starts with a (orbit) periodic behaviour with period $p = t - s < N$. The idea is that this provides the requested splitting of L .

Let γ be the maximal legal subpath of L on the left side of α (if any). Claim 6.6 can then be applied to the subpath $\gamma\alpha$. We wish to show that the point in α obtained from that process is the same as the one obtained by applying Claim 6.6 to $\alpha\beta$.

Let c, a' in γ, α respectively, be the points provided by Claim 6.6, so that $[c, a']$ is ppNp. As above we denote $c_n = f^n(c)$ and $a'_n = f^n(a')$.

CLAIM 6.7. — *We have $a = a'$. That is, applying Claim 6.6 locally results in well-defined points globally.*

Proof. — Let s' be the iterate where periodicity of $[c_n, a'_n]$ starts, and let p' be the period. Without loss of generality we may assume $s' \leq s$. Since $[c_{s'}, a'_{s'}]$ is a pNp, then also $[c_s, a'_s]$ is a pNp with the same period p' . Let $P = pp'$, note that $P \leq N^2$. Both $[c_s, a'_s]$ and $[a_s, b_s]$ are P -periodic. At time s the segment $[a_s, b_s]$ is contained in $[f^s(L)]$ but a priori its extremities may get cancelled from $[f^n(L)]$ in subsequent iterations. The same holds for $[c_{s'}, a'_{s'}]$.

Suppose that images of a, a' are not cancelled till the next three iterations of f^P (note that $s + 3P \leq N + 3N^2 \leq M$). Since f -images of edges have length more than C and L contains no legal subpath of that length, then for $n = s, s + P, s + 2P$ the segment $[a_n, a'_n]$ — which is the preimage of $[a_{s+3P}, a'_{s+3P}]$ — contains at most one vertex. Therefore there are two iterates in the first three steps so that $[a_n, a'_n]$ contains the same number of vertices which is either zero or one. Now, since f is expanding and a_n and a'_n are orbit-periodic, this forces $a_n = a'_n$, so $a = a'$ (because $[a, a']$ is a legal path). (To be precise here we don't use only the periodicity of a, a' but the periodicity of the pNp's $[c_s, a'_s]$ and $[a_s, b_s]$ because we need the orbit periodicity of the oriented edges containing a_n and a'_n .)

We end the proof by proving that a is not cancelled in iterations till $s + 3P$ (the same reasoning will work for a'). Suppose the contrary. Let τ_n be the illegal turn of $[c_n, a'_n]$ (τ_n is in the orbit of τ , but $\tau_n \neq f^n(\tau)$). Since a is cancelled, then $a \in \text{left}(\alpha)$ (because we know it is not in $\text{right}(\alpha)$) and in particular $[a_{s+3P}, b_{s+3P}]$ contains τ_{s+3P} on the left side of $\text{surv}(\alpha_{s+3P})$. By periodicity $[a_{s+4P}, b_{s+4P}]$ contains the segment $[f^P(\tau_{s+3P}), \tau_{s+4P}]$, still in the left side of $\text{surv}(\alpha_{s+4P})$. In particular it contains a whole edge, hence

$[a_{s+5P}, b_{s+5P}]$ contains a legal segment of length more than C on the left side of $\text{surv}(\alpha_{s+5P})$. By periodicity, $[a_s, b_s]$ contains a legal subpath of length more than C which contradicts the fact that (2) fails (note that $s + 5P \leq N + 5N^2 = M$). \square

Now, if L is simplicial, then we have provided a splitting of L in ppNp's, as required. In the general case, we have a splitting of the simplicial closure of L so that interior paths are ppNp's.

Now let's focus on tails. Let α be a tail, say the starting one, and let a be the point in the simplicial closure of α given by Claim 6.6. So a is the starting point of our L' . Let x be the starting point of L (note that x is never cancelled till M iterates because (1) fails). Point a may lie either on the left or on the right side of x . If a is on the left side of x , i.e. $x \in L'$, then the images $f^M(x)$ and $f^M(a)$ are in the same edge (and edges have length less than C because (2) fails), and in particular at distance less than C apart. If x is not in L' , then the segment $[f^M(x), f^M(a)]$ is shorter than C because $[x, a]$ is legal and never affected by cancellations, and (2) fails. In both cases $d_T(x, a) \leq C/\lambda^M$.

The proof of Proposition 6.4 is now complete. \square

The following is now immediate, since we may iterate Proposition 6.4, bearing in mind Lemma 2.38.

COROLLARY 6.8. — *In the hypotheses of Proposition 6.4, for any $C_1 > 0$ there exists an $M_1 \in \mathbb{N}$ such that, for any finite edge path or periodic line L in X , one of the following holds true:*

- (1) $[f^{M_1}(L)]$ has less illegal turns than L ;
- (2) $[f^{M_1}(L)]$ has a legal subpath of length more than C_1 ;
- (3) L splits (not necessarily at vertex-points) as a concatenation of paths $\rho_1 \cdots \rho_\kappa$ so that each ρ_i is a pre-periodic Nielsen path with at most one illegal turn.

The same conclusion holds true for finite paths whose endpoints are not necessarily vertices, with (3) above replaced by:

- (3') (i) L splits as a concatenation of paths $\delta_0 \rho_1 \cdots \rho_\kappa \delta_1$ so that each ρ_i is pre-periodic Nielsen path with at most one illegal turn;
- (ii) δ_0, δ_1 each cross at most one illegal turn;
- (iii) δ_0, δ_1 each have length at most $2\text{cc}(f)$.

DEFINITION 6.9. — *Let $f: X \rightarrow X$ be a train track map. A path β in X is called pre-legal if for some $n \in \mathbb{N}$, $[f^n(\beta)]$ is legal.*

LEMMA 6.10 (The 2/3-lemma). — *Let $f: X \rightarrow X$ be a simplicial train track map representing some $[\phi] \in \text{Out}(\mathcal{G})$, with $\text{Lip}(f) = \lambda > 1$. Let L*

be either a finite path or a periodic line in X . Let $M = 5N^2 + N$ be the constants of Proposition 6.4. Suppose that no legal subpath of length more than $C = \text{cc}(f) + 1$ appears in iterates $[f^n(L)]$ for $1 \leq n \leq M$, but that $[f^n(L)]$ becomes eventually completely legal for some $n > M$ (that is, L is pre-legal). Then:

$$\#\{\text{illegal turns of } [f^M(L)]\} \leq \frac{2}{3}(\#\{\text{illegal turns of } L\} + 1).$$

In particular, if $\#\{\text{illegal turns of } L\} > 3$, then⁽⁵⁾

$$\#\{\text{illegal turns of } [f^M(L)]\} < \frac{8}{9}\#\{\text{illegal turns of } L\}.$$

Proof. — We order L and $f^M(L)$ accordingly. Let $\sigma_0 \dots \sigma_h$ be the subdivision of $[f^M(L)]$ in maximal legal subpaths. Let S_i be the starting point of σ_i , which coincides with the ending point of σ_{i-1} . Let $V_1 = \min f^{-M}(S_1)$ and $W_1 = \max f^{-M}(S_1)$. (We may have $V_1 = W_1$ if for example $f^{-M}(S_1) = V_1$ and the turn at V_1 is legal, which can happen because our maps are train track for $\langle \sim_{f^i} \rangle$, not necessarily for \sim_f .) Define $\gamma_1 = [V_1, W_1]$. The set $f^{-M}(S_2) \cap \{x > W_1\}$ is nonempty just because $f^M(\gamma_1)$ can be retracted to S_1 in $[f^M(L)]$ (also, note that $f^M(\gamma_1)$ is a tree which is rooted at S_1 , but a priori it could contain the segment σ_2 , so one could possibly see some preimage of S_2 in γ_1). Let V_2, W_2 be respectively the min and max of $f^{-M}(S_2) \cap \{x > W_1\}$, and define $\gamma_2 = [V_2, W_2]$. Recursively define V_i, W_i, γ_i in the same way, and define subpaths $\xi_0 = \{x \leq V_1\}$, $\xi_i = [W_i, V_{i+1}]$ for $i < h$, and $\xi_h = \{x \geq W_h\}$ (so ξ_i is a preimage of σ_i in a broad sense).

Since any γ_i gets cancelled in $[f^M(L)]$, then it contains at least one illegal turn (γ_i may be a single point at an illegal turn of L , in this case we abuse notation and still say that γ_i is a path containing one illegal turn). Suppose that γ_i contains only one illegal turn. For any $x_i \in \xi_{i-1}$, $y_i \in \xi_i$, Proposition 6.4 applies to the path $[x_i, y_i]$, and taking those points sufficiently close to γ_i , we may assure that we are in the situation (3'iii), for both x_i, y_i (so the images of endpoints of the ppNp provided by Proposition 6.4 are in σ_{i-1}, σ_i respectively). Proposition 6.4 in particular implies that we find $z_i \in \xi_{i-1}, t_i \in \xi_i$ so that $[z_i, t_i] \subseteq \xi_{i-1} \gamma_i \xi_i$ is ppNp, with periodicity starting before N iterates,⁽⁶⁾ and with period less than N , where N is as in Proposition 6.4. Note that by periodicity, and since f is expanding, then either t_i coincides or it is on the left of z_{i+1} . (Indeed, since f is expanding, the f^n -images of $[z_i, t_i]$ are longer and longer paths, in which some central

⁽⁵⁾ Because if $x > 3$, then $\frac{2}{3}(x+1) < \frac{8}{9}x$.

⁽⁶⁾ See Definition 2.39.

portion cancels because of periodicity. Any point x in the interior of $[z_i, t_i]$ eventually cancels in iterates because otherwise the distance of $f^n(x)$ from $f^n(z_i)$ would exceed the distance between $f^n(z_i)$ and $f^n(t_i)$. By periodicity, z_{i+1} does not cancel, and so it must stay on the right of t_i .)

We say that γ_i is:

- *periodic* if γ_i contains only one illegal turn, and such turn will be called periodic;
- *non-periodic* if γ_i contains at least two illegal turns.

Since $[f^n(L)]$ becomes eventually legal, all illegal turns must disappear, and they can disappear in two ways: either they become legal after some iteration-cancellation of the turn itself, or they are cancelled by overlapping the image of some other illegal turn (since f is train track, no new illegal turns are created).

Since illegal turns in the periodic γ_i 's remain illegal forever (because they are part of a pNp), then they must cancel by overlapping or just because they are at extremities of L and the ppNp provided by Proposition 6.4 strictly contains that extremity of L . There are at most two illegal turns of the last kind, at the extremities of L .

CLAIM 6.11. — *The illegal turns of iterates of two different periodic γ_i 's never overlap.*

Proof. — Suppose the contrary. Let γ_i and γ_j be two periodic paths whose illegal turns eventually overlap. Let τ be the illegal turn of γ_i and ω that of γ_j . Let $\tau_n = f^n(\tau)$ and $\omega_n = f^n(\omega)$. Let $s < N$ be such that both γ_i and γ_j become periodic from step s on, and let p_i, p_j their periods, whose product $p = p_i p_j$ is less than N^2 (by Proposition 6.4). Let n_0 be the first iterate when $\tau_{n_0} = \omega_{n_0}$. Let $N \leq q < N + p$ such that $q \equiv n_0 \pmod{p}$ (note that $N + p \leq N + N^2 < M$). By periodicity τ_q is in the same orbit as τ_{n_0} and ω_q in the same orbit as ω_{n_0} . Thus there is $g \in G$ such that $\omega_q = g\tau_q$. Note that $g \neq \text{id}$ because γ_i and γ_j remains disjoint till iterate M . But now $f^{n_0-q}(\tau_q) = \tau_{n_0} = \omega_{n_0} = f^{n_0-q}(\omega_q) = f^{n_0-q}(g\tau_q) = g\tau_{n_0}$ forces $g = \text{id}$, a contradiction. \square

Therefore, except possibly for the two extremal γ_i 's, we can associate to any periodic γ_j the illegal turn $\tau_{\text{canc}}(\gamma_j)$ of some other γ_k which eventually cancels the unique illegal turn of γ_j . By the claim above, the turn $\tau_{\text{canc}}(\gamma_j)$ is one of the illegal turns of $[f^M(L)]$ that come from a non-periodic γ_k . Again by the claim, different periodic γ_i 's have associated different non-periodic turns, that is $\tau_{\text{canc}}(\gamma_i) \neq \tau_{\text{canc}}(\gamma_j)$ for $i \neq j$. That is, τ_{canc} is injective.

Let A, B be respectively the number of periodic and non-periodic illegal turns we see in $[f^M(L)]$. Since τ_{canc} is injective and defined possibly except for at most two γ_i 's, we have $A - 2 \leq B$. It follows that

$$\frac{(A - 2) + B}{(A - 2) + 2B} \leq \frac{2}{3}$$

hence

$$A + B \leq \frac{2}{3}(A - 2 + 2B) + 2 = \frac{2}{3}(A - 2 + 2B + 3) = \frac{2}{3}(A + 2B + 1).$$

The number of illegal turns in $[f^M(L)]$ is $A + B$ by definition. Any non-periodic illegal turn in $[f^M(L)]$ contributes to at least two⁽⁷⁾ illegal turns in L , so the number of illegal turns of L is at least $A + 2B$ and the lemma is proved. \square

Remark 6.12. — The statement is sharp, as you can build a path with two illegal turns that survives till M but then disappears, just by a concatenation of two pNp's to which we cut a suitable portion near the ends.

DEFINITION 6.13. — For a simplicial G -tree X denote by a_X the length of the shortest edge of X .

LEMMA 6.14. — Let $[\phi] \in \text{Out}(\mathcal{G})$. Let $Y \in \mathcal{O}(\mathcal{G})$ such that there exists a simplicial train track map $f_Y: Y \rightarrow Y$ representing $[\phi^{-1}]$ (for example if $[\phi]$ is \mathcal{G} -irreducible and $Y \in \text{Min}(\phi^{-1})$ admits a simplicial train track). For any constant $C_1 > 0$, and any $X \in \mathcal{O}(\mathcal{G})$, set

$$D = \frac{C_1}{a_Y} + 1 \quad D' = D\Lambda(X, Y)\Lambda(Y, X).$$

Then, with these constants, the following holds true for any $g \in \text{Hyp}(\mathcal{G})$: if $\frac{\ell_X(\phi^n(g))}{\ell_X(g)} < 1/D'$, then the axis of g in Y contains an f_Y -legal subpath of length at least C_1 .

Proof. — First, observe that if $\frac{\ell_X(\phi^n(g))}{\ell_X(g)} < 1/D'$, then

$$\begin{aligned} \frac{\ell_Y(\phi^{-n}\phi^n(g))}{\ell_Y(\phi^n(g))} &= \frac{\ell_Y(g)}{\ell_Y(\phi^n(g))} \\ &= \left(\frac{\ell_X(g)}{\ell_X(\phi^n(g))} \right) \left(\frac{\ell_Y(g)}{\ell_X(g)} \right) \left(\frac{\ell_X(\phi^n(g))}{\ell_Y(\phi^n(g))} \right) \\ &\geq \frac{\ell_X(g)}{\ell_X(\phi^n(g))} \frac{1}{\Lambda(Y, X)} \frac{1}{\Lambda(X, Y)} \\ &> D. \end{aligned}$$

⁽⁷⁾ Because any illegal turn of $[f^M(L)]$ which is not in a periodic γ_i comes from a non-periodic γ_i , each of which contains at least two illegal turns.

Let L be the axis of $\phi^n(g)$ in Y . Let n_g be the number of f_Y -illegal turns in a fundamental domain of L . Then if $[f_Y^{-n}(L)]$ would not contain any legal subpath of length at least C_1 , since f_Y is train track, we would get

$$\ell_Y(g) \leq C_1 n_g.$$

But also (since f_Y is simplicial),

$$\ell_Y(\phi^n(g)) \geq n_g a_Y.$$

Hence, $\frac{\ell_Y(g)}{\ell_Y(\phi^n(g))} \leq \frac{C_1}{a_Y} < D$, contradicting the above inequality. \square

LEMMA 6.15. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ with $\lambda(\phi) = \lambda > 1$. Let $X, Y \in \mathcal{O}(\mathcal{G})$ such that there exist simplicial train track maps $f_X: X \rightarrow X$, $f_Y: Y \rightarrow Y$ representing $[\phi], [\phi^{-1}]$ respectively. Let $h: X \rightarrow Y$ be a straight G -equivariant map.*

Then, for any $C_2 > 0$, there exist $b \in \mathbb{N}$ and $0 \leq L_0 \in \mathbb{R}$ such that, for any path β in X , if β is pre-legal (for f_X) and $l_X(\beta) \geq L_0$, then either

- (i) $[f_X^b(\beta)]$ contains an f_X -legal subpath of length at least C_2 , or
- (ii) $[h(\beta)]$ contains an f_Y -legal subpath of length at least C_2 .⁽⁸⁾

Proof. — We start by setting some constants.

The constant K . There exists a uniform constant K (depending only on X) such that for any path $\beta = [u, v]$ in X there exists $g \in G$ such that:

- (i) u is in the axis of g ;
- (ii) $v \in [u, gu]$;
- (iii) the distance from v to gu is bounded above by K .

That is, if we look at the quotient graph of groups, we can complete the image of any β to a cyclically reduced loop by adding a path of length at most K .

The constant C_1 . Let $C = \text{cc}(f_X) + 1$ and

$$C_1 = \max\{C, 2C_2 + 2\text{Lip}(h)K\}.$$

The constant M . Our constant M is the same one that appears in Proposition 6.4, Corollary 6.8, and Lemma 6.10.

We then set, bearing in mind Lemma 6.14, the following constants.

⁽⁸⁾ We note that this result is an analogue of [2, Lemma 2.10]. However, that lemma requires the automorphism to be nongeometric, as it is otherwise false. We do not make the nongeometric assumption, and it is this that requires us to restrict to pre-legal paths.

The constants D, D', D'' .

$$D = \frac{C_1}{a_Y} + 1, \quad D' = D\Lambda(X, Y)\Lambda(Y, X), \quad D'' = 2D'.$$

The constants $b, b_1, b_2 \in \mathbb{N}$. They are so that

$$\left(\frac{8}{9}\right)^{b_1} C_1 \leq \frac{a_X}{D''}, \quad \lambda^{b_2} \geq C_2, \quad b = b_1 M + b_2.$$

The constants $L_1, L_0 > 0$. They are defined by:

$$\frac{1}{D''} + \frac{\lambda^b K}{L_1} = \frac{1}{D'}, \quad L_0 = \max\{L_1, 4C_2 D''\}.$$

We note that for any $L \geq L_1$ (in particular if $L \geq L_0$) we have,

$$\frac{1}{D''} + \frac{\lambda^b K}{L} \leq \frac{1}{D'}.$$

We now argue as follows. If $[f_X^b(\beta)]$ contains a f_X -legal subpath of length at least C_2 , we are done. Otherwise, consider paths $\beta, [f_X^M(\beta)], \dots, [f_X^{Mb_1}(\beta)]$. If any of these paths contain an f_X -legal subpath of length at least $C = \text{cc}(f_X) + 1$, then by Lemma 2.38, $[f_X^{Mb_1+b_2}(\beta)] = [f_X^b(\beta)]$ contains an f_X -legal subpath of length at least $\lambda^{b_2} \geq C_2$ and we are done.

Hence we may assume that maximal legal subpaths in each of these paths have length less than C . In particular $l_X([f_X^b(\beta)])$ is bounded by the number of its illegal turns times $\min\{C, C_2\}$. Let η_β be the number of illegal turns of β .

By Lemma 6.10, either the number of illegal turns in $[f_X^b(\beta)]$ is at most $\left(\frac{8}{9}\right)^{b_1} \eta_\beta$, or the number of illegal turns in some $[f_X^{kb_1}(\beta)]$, and hence also in $[f_X^b(\beta)]$, is at most 3.

In either cases, our choice of constants guarantees that

$$\frac{l_X([f_X^b(\beta)])}{l_X(\beta)} \leq \frac{1}{D''}.$$

Indeed, in the latter case we easily get $\frac{l_X([f_X^b(\beta)])}{l_X(\beta)} \leq 4C_2/L_0 \leq \frac{1}{D''}$. In the former case we have $l_X(\beta) \geq a_X \eta_\beta$, whence

$$\frac{l_X([f_X^b(\beta)])}{l_X(\beta)} \leq \left(\frac{8}{9}\right)^{b_1} C/a_X \leq C/C_1 D'' \leq 1/D''.$$

Next we complete β to a path $\beta\gamma$, as in the definition of our constant K . That is, $\beta\gamma$ is cyclically reduced, it is the fundamental domain of a hyperbolic element $g \in G$, and γ has length at most K . Note that

$$\ell_X(\phi^b(g)) \leq l_X([f_X^b(\beta\gamma)]) \leq l_X([f_X^b(\beta)]) + \lambda^b K,$$

and therefore (since $\ell_X(g) = \ell_X(\beta\gamma)$),

$$\frac{\ell_X(\phi^b(g))}{\ell_X(g)} \leq \frac{\ell_X([f^b(\beta\gamma)])}{\ell_X(\beta\gamma)} \leq \frac{1}{D''} + \frac{\lambda^b K}{\ell_X(\beta\gamma)} \leq \frac{1}{D'}.$$

Hence, by Lemma 6.14, the axis of g in Y contains an f_Y -legal subpath of length at least C_1 . This means that $[h(\beta\gamma)]$ contains an f_Y -legal subpath of length at least $C_1/2$ (the number $C_1/2$ arises from the fact that the legal subpath is really a subpath of the axis, and not necessarily of $[h(\beta\gamma)]$). We then get that $[h(\beta)]$ contains an f_Y -legal subpath of length at least $C_1/2 - \text{Lip}(h)K \geq C_2$, because the length of $[h(\gamma)]$ is at most $\text{Lip}(h)K$. \square

LEMMA 6.16. — *Let $f: X \rightarrow X$ be a simplicial train track map representing some $[\phi] \in \text{Out}(\mathcal{G})$. If p is any periodic line in X which is a concatenation of pNp 's, then any group element acting periodically on p is elliptic in $X_{+\infty}$.*

Proof. — This follows directly from the construction of $X_{+\infty}$. \square

THEOREM 6.17. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) > 1$ (that is, a relatively irreducible automorphism with exponential growth). Let $X \in \text{Min}(\phi)$ and $Y \in \text{Min}(\phi^{-1})$, and denote by $X_{+\infty}$ and $Y_{-\infty}$ the corresponding attracting tree and repelling tree for ϕ , respectively (Definition 2.63). Then there exists an $\epsilon > 0$ such that for all $g \in G$, either*

- $\ell_{X_{+\infty}}(g) = \ell_{Y_{-\infty}}(g) = 0$, or
- $\max\{\ell_{X_{+\infty}}(g), \ell_{Y_{-\infty}}(g)\} \geq \epsilon$.

Proof. — Without loss of generality (Proposition 2.65 and Remark 2.32), we may assume that each of X, Y supports a simplicial train track map — let's call these f_X and f_Y — representing $[\phi]$ and $[\phi^{-1}]$ respectively, so that we may apply Corollary 6.8. Let C_1 be larger than the critical constants for f_X, f_Y plus one, and then apply Corollary 6.8 to get constants M_X, M_Y of which we take the larger, and call this M_1 . By Lemma 2.38, for any path ρ in X , containing a legal segment of length C_1 , $[f_X^n(\rho)]$ contains a legal segment of length at least $(\lambda(\phi))^n$. Similarly for Y .

We first show that $X_{+\infty}$ and $Y_{-\infty}$ have the same elliptic (and hence hyperbolic) elements. To this end, suppose that g is $X_{+\infty}$ -elliptic. This implies that $\ell_X(\phi^n(g))$ is bounded for all $n \geq 1$, by some constant A (depending on g). (No legal segment in the axis for $\phi^n(g)$ can be longer than the critical constant, and the number of illegal turns is bounded.) Hence $\ell_Y(\phi^n(g)) \leq A\Lambda(X, Y)$ is also bounded for all $n \geq 1$. In particular, if we realise $\phi^n(g)$ as a periodic line in Y , then the number of illegal turns in its period is uniformly bounded for all $n \geq 1$ (just because the length

of paths is bounded below by a_Y times the number of turns). Let A_1 be greater than this number of illegal turns.

Suppose n is large and apply A_1 times $f_Y^{M_1}$ to the axis L of $\phi^n(g)$ in Y . By Corollary 6.8, we have either reduced the number of illegal turns by A_1 , which is impossible, or we have a legal segment of length C_1 in $[f_Y^{A_1 M_1}(L)]$, or L is a product of ppNp's in Y and hence $Y_{-\infty}$ elliptic (Lemma 6.16). We argue that only the last can occur, since otherwise the length of $[f_Y^{A_1 M_1 j}(L)]$ in Y would be at least $(\lambda(\phi^{-1}))^j$. But we know that $\ell_Y(g) \leq A\Lambda(X, Y)$, and we can take j so that $(\lambda(\phi^{-1}))^j > A\Lambda(X, Y)$ and $n = A_1 M_1 j$, getting to a contradiction. (Note that if L is the axis for $\phi^n(g)$ in Y , then $[f_Y^n(L)]$ is the axis for g in Y .)

Therefore $\ell_{X_{+\infty}}(g) = 0$ implies that $\ell_{Y_{-\infty}}(g) = 0$, and vice versa, by symmetry.

Next we set our notation and constants:

- (i) $h: X \rightarrow Y$ is a straight G -equivariant map;
- (ii) $C_2 = \max\{\text{cc}(f_X) + 1, \text{cc}(f_Y) + 1 + 2\text{BCC}(h)\}$;
- (iii) L_0, b are the constants from Lemma 6.15 with the previous value of C_2 ;
- (iv) J is any integer greater than $7L_0/a_X$;
- (v) M is the constant from Proposition 6.4;
- (vi) $\epsilon = \min\{1, 1/\lambda(\phi)^b, 1/\lambda(\phi)^{JM}\}$.

We are going to show that for any $X_{+\infty}$ -hyperbolic g

$$\max\{\ell_{X_{+\infty}}(g), \ell_{Y_{-\infty}}(g)\} \geq \epsilon.$$

Consider an $X_{+\infty}$ -hyperbolic g , which we represent as a periodic line L in X . First note that the statement here is really one about long group elements. More precisely, if $\ell_X(g)$ is bounded above by $7L_0$, then we may apply Proposition 6.4 J times. Since J is greater than the number of edges — whence that of illegal turns — in a period of L we cannot reduce the number of illegal turns J times. We also cannot write L (or $[f_X^{JM}(L)]$) as a product of ppNp's, as that would imply that g was $X_{+\infty}$ -elliptic. Hence, $[f_X^{JM}(L)]$ must contain an f_X -legal segment of length at least $\text{cc}(f_X) + 1$. Then, by Lemma 2.38, $[f_X^{JM+j}(L)]$ must contain an f_X -legal segment of length at least $\lambda(\phi)^j$, and so $\ell_{X_{+\infty}}(g) \geq 1/\lambda(\phi)^{JM} \geq \epsilon$.

Next, since g is $X_{+\infty}$ -hyperbolic, there exists some $n > JM$ such that $[f_X^n(L)]$ contains a legal subpath longer than the critical constant for f_X . Hence, by increasing n , we can assume that $[f_X^n(L)]$ contains arbitrarily long legal segments. In particular we shall assume that $[f_X^n(L)]$ contains a legal subpath of length at least $2\text{cc}(f_X) + 1$.

Let β be a subpath of L such that $[f_X^n(\beta)]$ is a legal subpath of $[f_X^n(L)]$. (Constructively, take two points p, q in $[f_X^n(L)]$ that bounds a legal subpath, and consider any two preimages in L . These are the endpoints of β . Thus, the endpoints of β lie in $f_X^{-n}([f_X^n(L)])$.)

If it were the case that $l_X(\beta) \geq L_0$, we could apply Lemma 6.15 to conclude that either $[f_X^b(\beta)]$ contains an f_X -legal subpath of length at least C_2 , or $[h(\beta)]$ contains an f_Y -legal subpath of length at least C_2 . In the former case, $[f_X^b(L)]$ contains $[f_X^b(\beta)]$ as a subpath, and hence an f_X -legal subpath of length at least $C_2 \geq \text{cc}(f_X) + 1$. Hence, by Lemma 2.38 the length of $[f_X^{b+j}(L)]$ is at least $\lambda(\phi)^j$ and thus, $\ell_{X+\infty}(g) \geq 1/\lambda(\phi)^b \geq \epsilon$. Similarly, in the latter case, $[h(L)]$ contains an f_Y -legal subpath of length at least $C_2 - 2\text{BCC}(h)$ (cancellation is possible on applying h). Since $C_2 - 2\text{BCC}(h) \geq \text{cc}(f_Y) + 1$, we conclude as before that $\ell_{Y-\infty}(g) \geq 1 \geq \epsilon$. Hence, we may assume that all such β have length less than L_0 .

Finally, we conclude as follows. Choose four points, $p_0, p_1, p_2, p_3 \in [f_X^n(L)]$, such that $[p_0, p_3]$ splits as $[p_0, p_1][p_1, p_2][p_2, p_3]$, where $[p_1, p_2]$ is a maximal legal subpath of $[f_X^n(L)]$ of length at least $2\text{cc}(f_X) + 1$, and $[p_0, p_1]$ and $[p_2, p_3]$ each consists of three maximal legal subpaths of $[f_X^n(L)]$. Then, for any $0 \leq i < n$ we backward-recursively choose preimages $f_X^{i-n}(p_s)$ of these points in each $[f_X^i(L)]$. Thus, for all i the path $[f_X^{i-n}(p_0), f_X^{i-n}(p_3)]$ is a subpath $[f_X^i(L)]$ which splits as

$$[f_X^{i-n}(p_0), f_X^{i-n}(p_1)] [f_X^{i-n}(p_1), f_X^{i-n}(p_2)] [f_X^{i-n}(p_2), f_X^{i-n}(p_3)].$$

Moreover, for all j we have $f_X^j(f_X^{-n}(p_s)) = f_X^{j-n}(p_s)$. Let

$$\begin{aligned} \gamma_0 &= [f_X^{-n}(p_0), f_X^{-n}(p_1)], \\ \beta &= [f_X^{-n}(p_1), f_X^{-n}(p_2)], \\ \gamma_1 &= [f_X^{-n}(p_2), f_X^{-n}(p_3)] \end{aligned}$$

and set $\gamma = \gamma_0\beta\gamma_1$. Since we are assuming that preimages of legal subpaths of $f_X^n(L)$ have length at most L_0 (as otherwise we are done), we have $l_X(\beta) < L_0$, $l_X(\gamma_0), l_X(\gamma_1) < 3L_0$, and $l_X(\gamma) < 7L_0$.

As before, we apply J times f_X^M to γ , where J (defined above) is a bound on the number of illegal turns in γ . We analyse the behaviour of this path using Proposition 6.4. We know that if we get a long legal segment for f_X , this would bound the length $\ell_{X+\infty}(g)$ from below (by $1/\lambda(\phi)^{JM} \geq \epsilon$). Also, we cannot reduce the number of illegal turns J times. Therefore, we are left with the case where for some $j \leq J$ we have that $[f_X^{jM}(\gamma)]$ splits as

$$[f_X^{jM}\gamma] = \delta_0\rho_1 \cdots \rho_\kappa\delta_1,$$

where each subpath has at most one illegal turn and the ρ_i 's are ppNps. Moreover, $[f_X^{jM}(\gamma)]$ has subpaths $[f_X^{jM}(\gamma_0)]$, $[f_X^{jM}(\beta)]$ and $[f_X^{jM}(\gamma_1)]$. We also know that the first and last of these cross at least two illegal turns, because $[p_0, p_1]$ and $[p_2, p_3]$ cross two illegal turns. In particular, $[f_X^{jM}(\beta)]$ is a subpath of $\rho_1 \cdots \rho_\kappa$. However, this is a splitting and therefore $[f_X^n(\beta)]$ is a subpath of $[f_X^{n-jM}(\rho_1 \cdots \rho_\kappa)]$, which is impossible since $[f_X^n(\beta)]$ is a legal path of length at least $2\text{cc}(f_X) + 1$, and no ppNp can contain a legal subpath of length greater than $\text{cc}(f_X)$. \square

7. Co-compactness of the Min-Set

We will prove our main theorem, first under the extra hypothesis that our automorphism is primitive. More specifically, we will prove that the Min-Set of a primitive irreducible automorphism is co-bounded under the action of $\langle \phi \rangle$, which implies that it is co-compact (see Section 4). The general result is proved in the next section, where we drop the primitivity hypothesis.

Let us fix a group G and a free factor system $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ of G . We are going to use north-south dynamics stated in Section 5, where we had the additional assumption of $\text{rank}(\mathcal{G}) = k + r \geq 3$. We remark that in case of lower rank, either we are in the classical CV_2 case, and the co-compactness result is known (a proof can be found for example in [15]) or the result is trivial.

7.1. Ultralimits

At this stage, our strategy is as follows: we will argue by contradiction so that if $\text{Min}_1(\phi)$ is not co-compact, then Theorem 4.26 provides us with a sequence of minimally displaced points Z_i which stay at constant Lipschitz distance from the attracting tree $X_{+\infty}$, but which are at unbounded distance from some basepoint. We can find scaling constants μ_i so that Z_i/μ_i is bounded, and we would like to take the limit of a sub-sequence of Z_i/μ_i . The problem with this is that a priori we do not have sequential compactness unless G is countable (Remark 4.18).

This is a minor issue, as what we really use is the existence of some adherence point of the above sequence. The easiest way to deal with this is to turn to ω -limits (or ultralimits). Of course, the main point of interest is exactly when G is countable, and in this case ω -limits are not needed.

So the reader may wish to simply read all the ω -limits as usual limits, and the arguments stay essentially the same (up to taking subsequences appropriately, and consider suitable \liminf or \limsup in some inequalities). We refer to Section 2 definitions and some basic properties for ω -limits, especially Definitions 2.53, 2.55 and Proposition 2.56. In our situation, we make the following definition of ω -limit of elements in $\overline{\mathcal{O}(\mathcal{G})}$.

DEFINITION 7.1. — *Let ω be a non-principal ultrafilter on \mathbb{N} . Let $(Y_i) \subset \overline{\mathcal{O}(\mathcal{G})}$ be a sequence and let $Y_\infty \in \overline{\mathcal{O}(\mathcal{G})}$. We say that Y_∞ is the ω -limit of Y_i , and write $Y_\infty = \lim_\omega Y_i$ if for any $g \in G$ we have*

$$\ell_{Y_\infty}(g) = \lim_\omega \ell_{Y_i}(g).$$

Remark 7.2. — Suppose that for any $g \in G$, the ω -limit of $\ell_{Y_i}(g)$ exists. Then the corresponding ω -limit length function is indeed the length function of an element in $Y_\infty \in \overline{\mathcal{O}(\mathcal{G})}$. This is because the conditions defining length functions of trees are closed under ω -limits (by Proposition 2.56 and [29]).

PROPOSITION 7.3. — *Let ω be a non-principal ultrafilter on \mathbb{N} . Consider a sequence of points $Z_i \in \mathcal{O}(\mathcal{G})$, and let $X \in \mathcal{O}(\mathcal{G})$. Set $\mu_i = \Lambda(X, Z_i)$. Then $\lim_\omega \frac{Z_i}{\mu_i}$ exists, is unique (depends on ω) and non-trivial.*

Proof. — Notice that for any $g \in G$, $\ell_{Z_i}(g) \leq \mu_i \ell_X(g)$. So each sequence $\frac{\ell_{Z_i}(g)}{\mu_i}$ is a bounded sequence and therefore has a unique ω -limit $l_T(g)$ for some T . Furthermore, Lemma 2.41 provides a finite set $H \subseteq G$ such that, for any i ,

$$\Lambda(X, Z_i) = \max_{h \in H} \frac{\ell_{Z_i}(h)}{\ell_X(h)}.$$

Hence,

$$\begin{aligned} 1 &= \lim_\omega \Lambda(X, Z_i / \mu_i) \\ &= \lim_\omega \max_{h \in H} \frac{\ell_{Z_i}(h) / \mu_i}{\ell_X(h)} \\ &= \max_{h \in H} \frac{\lim_\omega \ell_{Z_i}(h) / \mu_i}{\ell_X(h)} \\ &= \max_{h \in H} \frac{\ell_T(h)}{\ell_X(h)}, \end{aligned}$$

which shows that the limiting tree is non-trivial. Here we have used that the ω -limit commutes with a finite maximum (Proposition 2.56). \square

7.2. Co-compactness of the Min-Set of primitive irreducible automorphisms

For this, and the next section, we fix once and for all a non-principal ultrafilter ω on \mathbb{N} . In the subsequent results we refer to $\Lambda(T, -)$, where T may be a tree in $\overline{\mathcal{O}(\mathcal{G})}$, rather than just in $\mathcal{O}(\mathcal{G})$. We intend the following.

DEFINITION 7.4. — Let $T, W \in \overline{\mathcal{O}(\mathcal{G})}$. We set $\Lambda(T, W)$ to be the supremum of the ratios, $\frac{\ell_W(g)}{\ell_T(g)}$, over all elements which are hyperbolic in T . This is possibly infinite. We also set $\Lambda(T, W) = +\infty$ if there is a T -elliptic group element which is hyperbolic in W .

LEMMA 7.5. — Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) = \lambda > 1$. Let $X \in \text{Min}(\phi)$ and let $X_{+\infty}$ be the corresponding attracting tree. Suppose we have a sequence $(Z_i) \subset \text{Min}_1(\phi)$ for which there is μ_i such that $\lim_{\omega} \frac{Z_i}{\mu_i}$ exists and is a non-trivial tree T . Then, for any positive integer n we have:

- $\Lambda(T, T\phi^n) \leq \lambda^n$;
- if $\Lambda(T, X_{+\infty})$ is finite, then $\Lambda(T, T\phi^n) = \lambda^n$.

Proof. — The first claim is straightforward, since for any T -hyperbolic g

$$\frac{\ell_{T\phi^n}(g)}{\ell_T(g)} = \lim_{\omega} \frac{\ell_{Z_i\phi^n}(g)/\mu_i}{\ell_{Z_i}(g)/\mu_i} = \lim_{\omega} \frac{\ell_{Z_i\phi^n}(g)}{\ell_{Z_i}(g)} \leq \lambda^n.$$

For the second claim note that, for any T -hyperbolic g and any positive integer m ,

$$\begin{aligned} \Lambda(T, X_{+\infty}) &= \Lambda(T\phi^{mn}, X_{+\infty}\phi^{mn}) \\ &\geq \frac{\ell_{X_{+\infty}}(\phi^{mn}(g))}{\ell_T(\phi^{mn}(g))} \\ &= \frac{\ell_{X_{+\infty}}(g)\lambda^{mn}}{\ell_T(\phi^{mn}(g))} \\ &\geq \frac{\ell_{X_{+\infty}}(g)\lambda^{mn}}{\ell_T(g)(\Lambda(T, T\phi^n))^m} \\ &= \frac{\ell_{X_{+\infty}}(g)}{\ell_T(g)} \left(\frac{\lambda^n}{\Lambda(T, T\phi^n)} \right)^m. \end{aligned}$$

We note that at no stage are we dividing by zero here. Indeed, if g is T -hyperbolic, then it is also \mathcal{G} -hyperbolic and:

$$\frac{\ell_T(\phi(g))}{\ell_T(g)} = \lim_{\omega} \frac{\ell_{Z_i}(\phi(g))}{\ell_{Z_i}(g)} \geq \lim_{\omega} \frac{1}{\Lambda(Z_i\phi, Z_i)} \geq \lim_{\omega} \frac{1}{\Lambda(Z_i, Z_i\phi)^C} = \frac{1}{(\lambda(\phi))^C},$$

where the last inequality follows from quasi-symmetry, Lemma 2.19 and the fact that $\text{Min}_1(\phi)$ is uniformly thick (Theorem 2.47). In particular, this says that if we start with a T -hyperbolic group element g , then $\phi(g)$ is also T -hyperbolic (hence $\phi^{mn}(g)$ is also T -hyperbolic). \square

LEMMA 7.6. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) = \lambda > 1$. Let $X \in \text{Min}(\phi)$ and let $X_{+\infty}$ be the corresponding attracting tree. Let $Y \in \text{Min}(\phi)$ such that it admits a simplicial train track representative for $[\phi]$, and let $f_Y: Y \rightarrow X_{+\infty}$ be an optimal map. For any $g \in \text{Hyp}(\mathcal{G})$, and any integer n , we have*

$$\ell_{X_{+\infty}}(\phi^n(g)) \geq \Lambda(Y, X_{+\infty}) \ell_Y(\phi^n(g)) - \ell_Y(g) \frac{B}{a_Y},$$

where $B = \Lambda(Y, X_{+\infty}) \text{vol}(Y)$ is the BCC of f_Y and a_Y is the length of the shortest edge in Y .

Proof. — For $g \in \text{Hyp}(\mathcal{G})$, let η_g be the number of edges in a reduced loop representing g in $G \setminus Y$. Clearly $\ell_Y(g) \geq \eta_g a_Y$. By Proposition 2.64, the axis of $\phi^n(g)$ in Y can be written as a concatenation of at most η_g f_Y -legal paths. By Corollary 2.35,

$$\begin{aligned} \ell_{X_{+\infty}}(\phi^n \gamma) &\geq \Lambda(Y, X_{+\infty}) \ell_Y(\phi^n \gamma) - \eta_g B \\ &\geq \Lambda(Y, X_{+\infty}) \ell_Y(\phi^n(g)) - \ell_Y(g) \frac{B}{a_Y}. \end{aligned} \quad \square$$

LEMMA 7.7. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible, with $\lambda(\phi) > 1$. Then there is a constant $\epsilon > 0$ such that for all $W \in \text{Min}_1(\phi)$, admitting a simplicial train track representative of $[\phi]$, we have*

$$a_W > \epsilon$$

where a_W denotes the length of the shortest edge in W .

Proof. — Note that uniform thickness of minimally displaced points is not enough in order to have a lower bound on the lengths of the edges, but for points supporting simplicial train track representatives there is such a bound, as there are finitely many transition matrices of simplicial train tracks representing $[\phi]$,⁽⁹⁾ and the lengths for edges are given by eigenvectors of the Perron–Frobenius eigenvalue of these matrices. \square

⁽⁹⁾Such transition matrices have integer non-negative entries, and spectral radius $\lambda(\phi)$. Hence each entry of the matrix is bounded by λ_ϕ . So for any given spectral radius λ there are only finitely many matrices with non-negative integer coefficients and spectral radius not exceeding λ .

PROPOSITION 7.8. — Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) = \lambda > 1$. Let $X \in \text{Min}(\phi)$ and let $X_{+\infty}$ be the corresponding attracting tree. Let $T \in \overline{\mathcal{O}(\mathcal{G})}$ be a non-trivial tree which is the ω -limit of a sequence Z_i/μ_i with the following properties:

- (1) the Z_i are uniformly thick and with co-volume 1; that is, $\exists \epsilon_0 > 0 \forall i, Z_i \in \mathcal{O}_1(\epsilon_0)$;
- (2) $\mu_i \rightarrow \infty$;
- (3) there is $\delta > 0$ such that $\Lambda(Z_i, X_{+\infty}) \geq \delta$, for all i ;
- (4) there is a sequence $W_i \in \text{Min}_1(\phi)$ and $K > 0$ such that $\Lambda(Z_i, W_i) \leq K$, for all i (for example if the Z_i themselves belong to $\text{Min}_1(\phi)$).

Then $\Lambda(T, X_{+\infty}) = \infty$.

Proof. — Note that by assumption T is non-trivial. Without loss of generality, we may assume that each W_i supports a simplicial train track representing $[\phi]$ (Remark 2.32).

Note also that the points W_i are uniformly thick because they are minimally displaced (Theorem 2.47) and the same is true for Z_i , by assumption. Therefore, all the points W_i, Z_i belong to some uniform thick part and since the stretching factor Λ is multiplicatively quasi-symmetric when restricted on any thick part $\mathcal{O}_1(\epsilon)$ (Theorem 2.19) it follows that there is some uniform constant C such that

$$\Lambda(Z_i, W_i)^{1/C} \leq \Lambda(W_i, Z_i) \leq \Lambda(Z_i, W_i)^C.$$

In particular, it follows that for the constant $K_1 = K^C$, we get that for any i :

$$\Lambda(W_i, Z_i) \leq \Lambda(Z_i, W_i)^C \leq K^C = K_1.$$

We will prove now that there is a non-trivial tree $S \in \overline{\mathcal{O}(\mathcal{G})}$ so that W_i/μ_i ω -converges to S . We first observe that for every hyperbolic element $g \in G$ and positive integer i , we have that:

$$0 \leq \ell_{\frac{W_i}{\mu_i}}(g) = \frac{\ell_{W_i}(g)}{\mu_i} \leq K \frac{\ell_{Z_i}(g)}{\mu_i}.$$

It follows that the sequence W_i/μ_i is bounded, and so $S = \lim_{\omega} W_i/\mu_i$ exists. Moreover, T, S have finite distances to each other (in particular, S is non-trivial since T is non-trivial, and they admit the same hyperbolic elements) because of the inequalities:

$$\frac{\ell_S(g)}{\ell_T(g)} = \lim_{\omega} \frac{\ell_{W_i}(g)/\mu_i}{\ell_{Z_i}(g)/\mu_i} \leq K$$

and similarly

$$\frac{\ell_T(g)}{\ell_S(g)} = \lim_{\omega} \frac{\ell_{Z_i}(g)/\mu_i}{\ell_{W_i}(g)/\mu_i} \leq K_1.$$

Therefore, it is enough to prove that $\Lambda(S, X_{+\infty})$ is infinite. We argue by contradiction, assuming that

$$\Lambda(S, X_{+\infty}) < \infty.$$

Then, by Lemma 7.5, $\Lambda(S, S\phi^m) = \lambda^m$ for any positive integer m . For all such m , we then choose a S -hyperbolic element g_m such that

$$\ell_S(\phi^m(g_m))/\ell_S(g_m) \geq \lambda^m/2.$$

Now we apply Lemma 7.6 to W_i and get constants $B_i = \Lambda(W_i, X_{+\infty})$ (because $\text{vol}(W_i) = 1$), and $\epsilon_i := a_{W_i}$. By Lemma 7.7, there is a uniform $\epsilon > 0$ so that $\epsilon_i > \epsilon$ for all i . On the other hand, by the properties of Z_i and the triangle inequality, we get that all the distances $\Lambda(W_i, X_{+\infty})$ are uniformly bounded from below by $\frac{\delta}{K}$. Thus

$$\begin{aligned} \ell_{X_{+\infty}}(\phi^m(g_m)) &\geq \Lambda(W_i, X_{+\infty}) \left(\ell_{W_i}(\phi^m(g_m)) - \ell_{W_i}(g_m) \frac{1}{\epsilon} \right) \\ &\geq \frac{\delta}{K} \left(\ell_{W_i}(\phi^m(g_m)) - \ell_{W_i}(g_m) \frac{1}{\epsilon} \right). \end{aligned}$$

Therefore,

$$\begin{aligned} \frac{\ell_{X_{+\infty}}(\phi^m(g_m))}{\ell_{W_i}(\phi^m(g_m))} &\geq \frac{\delta}{K} \left(1 - \frac{\ell_{W_i}(g_m)}{\epsilon \ell_{W_i}(\phi^m(g_m))} \right) \\ &\xrightarrow{\lim_{\omega}} \frac{\delta}{K} \left(1 - \frac{\ell_S(g_m)}{\epsilon \ell_S(\phi^m(g_m))} \right) \geq \frac{\delta}{K} \left(1 - \frac{2}{\epsilon \lambda^m} \right). \end{aligned}$$

For any $0 < \delta_0 < 1$, choose m such that $1 - \frac{2}{\epsilon \lambda^m} \geq 1 - \delta_0$. For this choice of m , let $c_i = \frac{\ell_{X_{+\infty}}(\phi^m(g_m))}{\ell_{W_i}(\phi^m(g_m))}$. Then the calculation above shows that $\lim_{\omega} c_i \geq \frac{\delta}{K}(1 - \delta_0) > 0$. On the other hand, $\Lambda(S, X_{+\infty}) \geq \lim_{\omega} c_i \mu_i$. Hence,

$$\Lambda(S, X_{+\infty}) \geq \lim_{\omega} \frac{\mu_i \delta(1 - \delta_0)}{K} = \infty,$$

contradicting the assumption $\Lambda(S, X_{+\infty}) < +\infty$. \square

Remember that if $T, S \in \text{Min}_1(\phi)$ we have (Proposition 2.65 and basic properties of stable trees):

$$\begin{aligned} T_{+\infty} &= \Lambda(S, T_{+\infty}) S_{+\infty}, \\ \Lambda(S_{+\infty}, T_{+\infty}) &= \Lambda(S, T_{+\infty}), \\ T_{\infty} \phi^{\pm n} &= \lambda(\phi)^{\pm n} T_{\infty}. \end{aligned}$$

PROPOSITION 7.9. — Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) = \lambda > 1$. Let $X \in \text{Min}_1(\phi)$, $Y \in \text{Min}_1(\phi^{-1})$, and let $X_{+\infty}, Y_{-\infty}$ be the corresponding attracting and repelling trees.

Then, for every constant $\nu > 0$, there is a constant $\delta = \delta(\nu, \lambda, X, Y) > 0$ so that for every $Z \in \text{Min}_1(\phi)$ we have

$$\Lambda(Z, X_{+\infty}) \leq \nu \implies \Lambda(Z, Y_{-\infty}) \geq \delta.$$

Proof. — Let $\epsilon > 0$ be the constant given by Theorem 6.17. For any positive number ν , fix a (for instance the smallest) positive integer n_0 for which

$$\lambda^{-n_0} \leq \frac{\epsilon}{4\nu}.$$

Let now $Z \in \text{Min}_1(\phi)$ be such that $\Lambda(Z, X_{+\infty}) \leq \nu$, and let $Z_{+\infty}$ be the corresponding attracting tree. Let g be a candidate that realises $\Lambda(Z, X_{+\infty})$. In particular g and any of its power are not $X_{+\infty}$ -elliptic. Moreover, the length of g with respect to both Z and Z_{∞} is bounded above by 2 (as the volume of Z is 1). If we set now $h = \phi^{-n_0}(g)$, we get

$$\begin{aligned} \ell_{X_{+\infty}}(h) &= \ell_{\Lambda(Z, X_{+\infty})Z_{+\infty}}(h) \\ &= \Lambda(Z, X_{+\infty})\ell_{Z_{+\infty}}(\phi^{-n_0}(g)) \\ &\leq \nu\ell_{Z_{+\infty}}(\phi^{-n_0}(g)) \\ &= \nu\ell_{(Z_{+\infty}\phi^{-n_0})}(g) \\ &= \nu\ell_{(\lambda^{-n_0}Z_{+\infty})}(g) \\ &= \nu\lambda^{-n_0}\ell_{Z_{+\infty}}(g) \\ &\leq 2\nu\lambda^{-n_0} \\ &\leq \frac{\epsilon}{2} \\ &< \epsilon. \end{aligned}$$

Therefore, by Theorem 6.17, it follows that $\ell_{Y_{-\infty}}(h) \geq \epsilon$.

By multiplicative quasi-symmetry of Λ restricted on the thick parts of $\mathcal{O}_1(\mathcal{G})$ (Theorem 2.19), there exists a constant C such that $\Lambda(T, S) \leq \Lambda(S, T)^C$, for any $T, S \in \text{Min}_1(\phi) \cup \text{Min}_1(\phi^{-1})$ (note that C depends only on $[\phi]$ because elements in Min-Sets are uniformly thick because $[\phi]$ is irreducible). In particular,

$$\Lambda(Z, Z\phi^{-n_0}) = \Lambda(Z\phi^{n_0}, Z) \leq \Lambda(Z, Z\phi^{n_0})^C = \lambda^{Cn_0}.$$

Therefore, as the length of h with respect to Z is at most 2, we get that

$$\ell_Z(h) = \ell_Z(\phi^{-n_0}(g)) \leq \ell_Z(g)\Lambda(Z, Z\phi^{-n_0}) \leq 2\lambda^{Cn_0},$$

which implies

$$\Lambda(Z, Y_{-\infty}) \geq \frac{\ell_{Y_{-\infty}}(h)}{\ell_Z(h)} \geq \frac{\epsilon}{2\lambda^{Cn_0}} = \delta,$$

where the quantity δ does not depend on Z . \square

COROLLARY 7.10. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) = \lambda > 1$. Let $X \in \text{Min}_1(\phi)$, $Y \in \text{Min}_1(\phi^{-1})$, and let $X_{+\infty}, Y_{-\infty}$ be the corresponding attracting and repelling trees.*

Let $T \in \overline{\mathcal{O}(\mathcal{G})}$ which is the ω -limit of a sequence Z_i/μ_i with the following properties:

- (1) $Z_i \in \text{Min}_1(\phi)$;
- (2) $\mu_i \rightarrow \infty$;
- (3) *there is $\nu > 1$ such that $1 \leq \Lambda(Z_i, X_{+\infty}) \leq \nu$.*

Then

$$\Lambda(T, X_{+\infty}) = \infty = \Lambda(T, Y_{-\infty}).$$

Proof. — The equality for $X_{+\infty}$, follows by just applying Proposition 7.8 directly on the sequence $Z_i = W_i$, with $\delta = 1$.

For $Y_{-\infty}$, we first apply Theorem 3.5, which provides us with a sequence of minimally displaced points $W_i \in \text{Min}_1(\phi^{-1})$, with the property that for some uniform constant M ,

$$\max\{\Lambda(W_i, Z_i), \Lambda(Z_i, W_i)\} \leq M.$$

Next, we want to apply Proposition 7.8, for ϕ^{-1} . Conditions (1) and (2) are satisfied by our assumptions. By the choice of W_i 's, condition (4) is satisfied, too.

For property (3), we apply Proposition 7.9 for every i , to the point Z_i . By hypothesis $\Lambda(Z_i, X_{+\infty}) < \nu$, and Proposition 7.9 provides the $\delta > 0$ such that $\Lambda(Z_i, Y_{-\infty}) > \delta$, as required. \square

THEOREM 7.11. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -primitive irreducible (that is, a relatively irreducible automorphism with a train track with primitive transition matrix — and hence exponential growth). Then $\text{Min}_1(\phi) = \text{Min}(\phi) \cap \mathcal{O}_1$ is co-compact, under the action of $\langle \phi \rangle$.*

Proof. — Let $X \in \text{Min}_1(\phi)$, $Y \in \text{Min}_1(\phi^{-1})$, and $X_{+\infty}, Y_{-\infty}$ be the corresponding attracting/repelling trees. By Theorem 5.2, $[X_{+\infty}] = [T_\phi^+]$ is the unique attracting class of trees for ϕ , and $[Y_{-\infty}] = [T_\phi^-]$ is the unique repelling class.

We now argue by contradiction and suppose that $\text{Min}_1(\phi)/\langle \phi \rangle$ is not compact. By Theorem 4.26(5), there is a sequence of points Z_1, Z_2, \dots ,

Z_i, \dots of $\text{Min}_1(\phi)$ for which $\Lambda(X, Z_i \phi^m) \geq i$, for every $m \in \mathbb{Z}$. Note that for all $m \in \mathbb{Z}$ and $W \in \mathcal{O}(\mathcal{G})$ we have

$$\begin{aligned}\Lambda(W\phi^m, X_{+\infty}) &= \Lambda(W, X_{+\infty}\phi^{-m}) \\ &= \Lambda(W, \lambda(\phi)^{-m}X_{+\infty}) \\ &= \lambda(\phi)^{-m}\Lambda(W, X_{+\infty}).\end{aligned}$$

Therefore, we can replace points Z_i with some points of their $\langle \phi \rangle$ -orbits (which we will still denote by Z_i) with the extra property

$$1 \leq \Lambda(Z_i, X_{+\infty}) \leq \lambda(\phi).$$

We set $\mu_i = \Lambda(X, Z_i)$. Now $\lim_{\omega} Z_i / \mu_i = T$ for some (non-trivial) tree T on the boundary of $\mathcal{O}(\mathcal{G})$, by Proposition 7.3. From Corollary 7.10, we know

$$(7.1) \quad \Lambda(T, Y_{-\infty}) = \infty = \Lambda(T, X_{+\infty}).$$

In particular, T does not belong to $[T_{\phi}^-]$. On the other hand, by Lemma 7.5, it follows that for every positive integer j

$$\Lambda(T, T\phi^j) \leq \lambda(\phi)^j,$$

or equivalently

$$\Lambda\left(T, \frac{T\phi^j}{\lambda(\phi)^j}\right) \leq 1.$$

By applying the north-south dynamics Theorem 5.2, on $T \notin [T_{\phi}^-]$, we get that $\frac{T\phi^j}{\lambda(\phi)^j}$ projectively converges to T_{ϕ}^+ , which is in the same projective class as $X_{+\infty}$, so there is $c > 0$ such that $\frac{T\phi^j}{\lambda(\phi)^j}$ converges to $cX_{+\infty}$. But in that case, $\Lambda(T, X_{+\infty})$ would be finite, contradicting (7.1). \square

Remark 7.12. — The primitivity assumption is used only in applying north-south dynamics in the last theorem, and not in previous results of this section. In Proposition 7.9 and Corollary 7.10, if one is allowed to use north-south dynamics (for instance for primitive automorphisms), then one can replace any instance of $Y_{-\infty}$ with $X_{-\infty}$.

7.3. Co-compactness of the Min-Set of general irreducible automorphisms

In this subsection, we will prove the co-compactness of the Min-Set for irreducible automorphisms of exponential growth. For this section we fix: a free factor system $\mathcal{G} = (\{G_1, \dots, G_k\}, r)$ of a group G ; an element $[\phi] \in \text{Out}(\mathcal{G})$ which is \mathcal{G} -irreducible, with $\lambda(\phi) = \lambda > 1$; an element

$X \in \text{Min}_1(\phi)$ supporting a simplicial train track map $f: X \rightarrow X$ representing $[\phi]$. In particular there is a genuine automorphism $\psi \in [\phi]$ represented by f (Definition 2.30) and, up possibly to replace ϕ with ψ , we may assume that f represents ϕ (that is, $f(gx) = \phi(g)f(x)$).

We denote by M_f the transition matrix of f (see Section 2.12). If M_f fails to be primitive, then we can partition the edge orbits into blocks so that, for some positive integer s , $M_{f^s} = M_f^s$ is a block diagonal matrix, which is a strictly positive matrix when restricted to a block. Correspondingly, we can define sub-forests, X_1, \dots, X_l , of X consisting of edges, and their incident vertices, belonging to a single block. The following two lemmas are straightforward.

LEMMA 7.13. — *Let f , X and X_i be defined as above. Then:*

- (i) *f permutes the X_i 's;*
- (ii) *each X_i is a G -forest (i.e. a forest which is G -invariant);*
- (iii) *the union of the X_i is X .*

We can then define cylinders.

DEFINITION 7.14. — *A cylinder is a connected component of some X_i .*

Remark 7.15. — We note that it is possible for two cylinders to intersect at a vertex, as long as the cylinders belong to different sub-forests $X_i \neq X_j$.

LEMMA 7.16. — *If C is a cylinder, then $f(C)$ is also a cylinder. Moreover, for any $g \in G$, also $g(C)$ is a cylinder, belonging to the same X_i as C .*

We also have the following.

LEMMA 7.17. — *For any cylinder C , and any vertex $v \in C$:*

- (i) *$\text{Stab}_G(C)$ contains a G -hyperbolic element;*
- (ii) *$\text{Stab}_G(v) \leq \text{Stab}_G(C)$.*

Proof. — Without loss of generality, the map f^s has a block diagonal transition matrix where the positive entry in every block is at least 3. Choose an edge e in C . Since f^s is train track, $f^s(e)$ is a legal path. The condition on f^s means that $f^s(e)$ crosses the orbit of e at least three times, and it is contained in the same X_i as C . This means that C contains a path crossing e , ge and he for some $g, h \in G$; where these three edges are distinct. Clearly, $g, h \in \text{Stab}_G(C)$. Since the action of G on X is edge-free, this implies that if both g and h are elliptic, then gh is hyperbolic (as e, ge, he are all in the legal path $f^s(e)$). Hence $\text{Stab}_G(C)$ contains a hyperbolic element. Finally, an element of $\text{Stab}_G(v)$ must send C to another

cylinder containing v but belonging to the same subforest as C . This means that it must preserve C . \square

We now define a new tree \mathcal{T} from this information, which remembers the construction of the dual tree of the partition of X in cylinders. We note that this is not a \mathcal{G} -tree because Lemma 7.17 tells us that vertex stabilisers are too big and in general edge stabilisers are not trivial. More precisely:

DEFINITION 7.18. — We define a G -tree \mathcal{T} as follows. This is a bi-partite tree:

- type I vertices are the cylinders of X ;
- type II vertices are the vertices of X which belong to at least two distinct cylinders.

The edges of \mathcal{T} are the pairs (C, v) where C is a type I vertex, and v is a type II vertex contained in C .

It is an easy exercise to see that \mathcal{T} is a G -tree.

PROPOSITION 7.19. — We get the following:

- (i) the stabiliser of an edge (C, v) of \mathcal{T} is equal to $\text{Stab}_G(v)$;
- (ii) f induces a map, $F: \mathcal{T} \rightarrow \mathcal{T}$ representing ϕ (that is $F(gx) = \phi(g)F(x)$), which sends vertices to vertices — preserving type — and edges to edges;
- (iii) the irreducibility of $[\phi]$ implies that all the edge stabilisers of \mathcal{T} are non-trivial.

Proof. — The first point follows from the second part of Lemma 7.17. The second point follows from Lemma 7.16, and the fact that f maps vertices to vertices.

For the final point, note that if \mathcal{T} had an edge with trivial stabiliser, we could collapse all the edges with non-trivial stabiliser, and get a new G -tree $\overline{\mathcal{T}}$ and a new map \overline{F} on this tree representing ϕ . Since the action of this tree is edge-free and non-trivial, this would correspond to a proper free factor system for G , which would be ϕ -invariant. However, Lemma 7.17 implies that this free factor system properly contains $[\mathcal{G}]$. Therefore, we would obtain a ϕ -invariant proper free factor system properly containing $[\mathcal{G}]$, a contradiction to the irreducibility of $[\phi]$. \square

THEOREM 7.20. — Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) > 1$ (that is, a relatively irreducible automorphism with exponential growth). Then the action of $\langle \phi \rangle$ on $\text{Min}_1(\phi) = \text{Min}(\phi) \cap \mathcal{O}_1$ is co-compact.

Proof. — We shall deduce this theorem from the primitive case. We have our base-point $X \in \text{Min}(\phi)$ which supports our simplicial train track map f ,

representing ϕ , but with (potentially) imprimitive transition matrix. Let $X^{+\infty}$ be the attracting tree corresponding to X . We use here the notation $X^{+\infty}$ instead of $X_{+\infty}$, as we did in the rest of the paper, for notational reasons of this proof.

We argue by contradiction and suppose that the action is not co-compact. Then, by Theorem 4.26(8), we may find a sequence of points $Y_i \in \text{Min}(\phi)$ such that:

- (i) $\text{vol}(Y_i)$ are uniformly bounded;
- (ii) $\Lambda(Y_i, X^{+\infty}) = 1$;
- (iii) $\mu_i := \Lambda(X, Y_i)$ is unbounded.

We then define $T = \lim_{\omega} Y_i / \mu_i$ which exists and is non-trivial by Proposition 7.3. In order to reach the desired contradiction, we will show that such non-trivial T is trivial.

Note that by Proposition 2.65, the first and second points imply that

$$Y_i^{+\infty} = \lim_{m \rightarrow \infty} \frac{Y_i \phi^m}{\lambda^m} = X^{+\infty}.$$

Consider a cylinder C in X , with stabiliser $H = \text{Stab}_G(C)$. We note that H is a free factor of G and $[\phi^s]$ induces by restriction an automorphism class $[\phi_H^s]$ of H . Also, \mathcal{G} induces a free factor system \mathcal{H} of H . The restriction of f^s induces a train track representative of $[\phi_H^s]$ with primitive transition matrix. Moreover, $[\phi_H^s]$ is \mathcal{H} -irreducible: f permutes the forests X_i , and if f^s shows an invariant free factor system on H , this can be translated to others $\text{Stab}_G(C')$ by iterating f , and so producing a global free factor system which is $[\phi]$ -invariant, which cannot exist because $[\phi]$ is irreducible. In particular, Theorem 7.11 applies to $[\phi_H^s]$.

Next, for each of the G -trees above, we may form the minimal invariant H subtree. We denote this invariant subtree with a subscript H , for instance $Y_{i,H}^{+\infty}$. The fact that $Y_i^{+\infty} = X^{+\infty}$ implies that $Y_{i,H}^{+\infty} = X_H^{+\infty}$ and hence $\Lambda(Y_{i,H}, X_H^{+\infty}) = 1$. We still get that $X_H, Y_{i,H}$ are minimally displaced points for ϕ_H^s , whose volumes are uniformly bounded. By Theorem 7.11, ϕ_H^s acts co-compactly on its minimally displaced set and this, by Theorem 4.26(7), means that

$$\Lambda(X_H, Y_{i,H}) = \sup_{\substack{h \in H \\ \ell_X(h) \neq 0}} \frac{\ell_{Y_i}(h)}{\ell_X(h)}$$

is bounded. But since

$$\ell_T(h) = \frac{\lim_{\omega} \ell_{Y_i}(h)}{\mu_i},$$

and μ_i is unbounded, we deduce that $\ell_T(h) = 0$ for all $h \in H$. By Lemma 2.59, this implies that H fixes a unique point of T , and that this is the same point fixed by any $\text{Stab}_G(v)$, for $v \in C$ with nontrivial stabiliser.

In particular, we may define a G -equivariant map from \mathcal{T} to T by mapping each vertex to the unique point of T which is fixed by the corresponding (and non-trivial) stabiliser. By Lemmas 2.59 and 7.19, each edge is actually mapped to a point. This means that the whole G -tree \mathcal{T} is mapped to a point. In this case, as the map from \mathcal{T} to T is G -equivariant, there would be a fixed point for the whole group G . But this would imply that T is trivial in the sense of translation length functions, contradicting the non-triviality of T . \square

COROLLARY 7.21. — *Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible and with $\lambda(\phi) > 1$. Then $\text{Min}_1(\phi)$, equipped with the symmetric Lipschitz metric, is quasi-isometric to a line.*

Proof. — The idea is to simply apply the Svarc–Milnor Lemma. The action of $\langle \phi \rangle$ is clearly properly discontinuous and Theorem 7.20 gives us cocompactness. The only obstacle is that the symmetric Lipschitz metric d_{sym} is not geodesic (or, even, a length metric). Define the intrinsic metric d_I to be the infimum of lengths of paths between any two points. Notice that since $\text{Min}_1(\phi)$ is thick, quasi-symmetry implies that the asymmetric Lipschitz metric d_{out} and d_{sym} are bi-Lipschitz equivalent functions. Since d_{out} is a geodesic asymmetric metric, we deduce that d_I and d_{sym} are also bi-Lipschitz equivalent, and we are done. \square

8. Applications

8.1. Relative centralisers

In this section we give an application of our main result, regarding relative centralisers of relatively irreducible automorphisms with exponential growth.

THEOREM 8.1. — *Let G be a group, \mathcal{G} a non-trivial free factor system for G , and $\mathcal{O}_1(\mathcal{G})$ be the corresponding co-volume 1 section of relative outer space. Let $[\phi] \in \text{Out}(\mathcal{G})$ be \mathcal{G} -irreducible with exponential growth, and let $X \in \text{Min}_1(\phi)$. Let $C(\phi)$ be the relative centraliser of $[\phi]$ in $\text{Out}(G, \mathcal{G})$. Then there is a finite index subgroup $C_0(\phi)$ of $C(\phi)$ such that $C_0(\phi)$ is the (internal) direct product*

$$C_0(\phi) = C_X(\phi) \times \langle [\phi] \rangle,$$

where $C_X(\phi) = \{[\psi] \in C_0(\phi) : X\psi = X\} = \text{Stab}(X) \cap C_0(\phi)$.

Proof. — First, note that $C(\phi)$ preserves $\text{Min}(\phi)$. By Theorem 7.20, there is a fundamental domain K for the action of $\langle [\phi] \rangle$ on $\text{Min}_1(\phi)$, which consists of finitely many simplices. Without loss of generality we assume that $X \in K$. For any $[\psi] \in C(\phi)$ we have that $X\psi \in K\langle [\phi] \rangle$, since K is a fundamental domain.

Define $C_0(\phi) := \{[\psi] \in C(\phi) : X\psi \in X\langle [\phi] \rangle\}$. Since K is finite, $C_0(\phi)$ is a finite index subgroup of $C(\phi)$ containing $[\phi]$. By definition, for every $[\psi] \in C_0(\phi)$, there is an $n \in \mathbb{Z}$ such that $X\psi = X\phi^n$. Hence $[\psi] = [\alpha\phi^n]$ for some $[\alpha] \in \text{Stab}(X) \cap C_0(\phi) = C_X(\phi)$. Therefore $C_0(\phi) = C_X(\phi)\langle [\phi] \rangle$.

(Morally, one could take $C_0(\phi)$ to be the subgroup which acts trivially on $\text{Min}_1(\phi)/\langle [\phi] \rangle$. This is a finite index subgroup since automorphisms preserve simplices and metric structures, so therefore each point in the quotient space has a finite orbit under the action of $C(\phi)$.)

Moreover, since $[\phi]$ acts without periodic points in $\mathcal{O}(\mathcal{G})$, $\langle [\phi] \rangle \cap C_X(\phi) = \{1\}$. As $\langle [\phi] \rangle$ commutes with $C_X(\phi)$ we get that $C_0(\phi)$ is the direct product, $C_0(\phi) = C_X(\phi) \times \langle [\phi] \rangle$. \square

Note that the previous result generalises a well-known result for free groups, that centralisers of irreducible automorphisms with irreducible powers are virtually cyclic (see [2]). It also generalises a result of the third author who proved a similar result for relative centralisers of relatively irreducible automorphisms, with the extra hypothesis that all the powers of the automorphism are irreducible (see [32]).

8.2. Centralisers in $\text{Out}(F_3)$

In this section we study centralisers of automorphisms in $\text{Out}(F_3)$. The main result of this section is the following.

THEOREM 8.2. — *Centralisers of elements in $\text{Out}(F_3)$ are finitely generated.*

Before going into the proof, we need to quote some preliminary facts. Our proof is based on Remark 2.46: any automorphism $[\phi] \in \text{Out}(F_3)$ is irreducible with respect to some relative outer space $\mathcal{O}(\mathcal{G})$, for some free factor system \mathcal{G} of the free group F_3 . Equivalently, in the language of free factor systems, \mathcal{G} is a maximal $[\phi]$ -invariant free factor system. However, a maximal free factor system for $[\phi]$ is not necessarily unique. In fact, there are automorphisms with infinitely many different maximal invariant free factor systems. The following theorem shows that under the extra assumption that $[\phi]$ does not act periodically on any free splitting (i.e. point

of some relative outer space), there are finitely many maximal invariant free factor systems. This is proved by Guirardel and Horbez in [19].

PROPOSITION 8.3 ([19]). — *Let $[\phi] \in \text{Out}(F_n)$. Suppose that there is no free splitting of F_n which is preserved by some power of $[\phi]$. Then there are finitely many maximal $[\phi]$ -invariant free factor systems $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_K$. As a consequence, the relative centraliser $C_{\mathcal{G}_i}(\phi)$ has finite index in $C(\phi)$, for $i = 1, \dots, K$.*

Proof. — The first part is a special case of [19, Corollary 1.14]. For the second part, we note that $C(\phi)$ preserves the finite set of maximal $[\phi]$ -invariant free factor systems $\{\mathcal{G}_1, \dots, \mathcal{G}_K\}$. As the relative centraliser with respect to the free factor system \mathcal{G}_i is simply $C_{\mathcal{G}_i}(\phi) = \text{Out}(F_n, \mathcal{G}_i) \cap C(\phi)$, the result follows. \square

On the other hand, we need to understand the complementary case of an automorphism that acts periodically on a free splitting. This case has also been studied in [19]. If T, S are two F_n -trees with trivial edge stabilisers (i.e. free splittings), then we say that T dominates S if point stabilisers in T are elliptic in S . In other words, if $T \in \mathcal{O}(\mathcal{G}_1)$, $S \in \mathcal{O}(\mathcal{G}_2)$, then T dominates S if and only if $\mathcal{G}_1 \leq \mathcal{G}_2$. Alternatively, T dominates S if $\Lambda(T, S) < \infty$. From [19] we can also extract the following proposition.

PROPOSITION 8.4. — *Let $[\phi] \in \text{Out}(F_n)$. Let's assume that there is a power of $[\phi]$ fixing a free splitting. Then there is a maximal (with respect to domination) $\langle [\phi] \rangle$ -periodic free splitting $T \in \mathcal{O}(\mathcal{G})$, for some free factor system \mathcal{G} . All such maximal free splittings belong to the same relative outer space $\mathcal{O}(\mathcal{G})$. Moreover, if $[\phi]$ has infinite order, then the centraliser $C(\phi)$ preserves the free factor system \mathcal{G} .*

Proof. — The first part is [19, Proposition 6.2] for the cyclic subgroup $H = \langle [\phi] \rangle$. The second part follows by [19, Theorem 8.32]. \square

Remark 8.5. — We recall that *maximal*, invariant, free factor systems are defined to be maximal with respect to the natural ordering \leq on free factor systems of F_n . It is important to mention here a *maximal* free splitting means that it belongs to the *minimal*, in terms of the ordering, relative outer space!

The linear growth case cannot be really studied using the methods that are presented in this paper, so we need the following result.

THEOREM 8.6 ([1]). — *Centralisers of linearly growing automorphisms in $\text{Out}(F_n)$ are finitely generated.*

We need also the following well-known result for $\text{Out}(F_2)$.

THEOREM 8.7. — *Centralisers of infinite order elements in $\text{Out}(F_2)$ are virtually cyclic.*

Proof. — This is clear as $\text{Out}(F_2)$ is virtually F_2 and centralisers of non-trivial elements in F_2 are cyclic. \square

Now we are in position to start the proof of the main result of this section.

Proof of Theorem 8.2. — The possible free factor systems of F_3 have one of the following types (for some free basis $\{a, b, c\}$):

- (1) $\mathcal{G} = \emptyset$ (note that in this case $\mathcal{O}(\mathcal{G}) = \text{CV}_3$);
- (2) $\mathcal{G} = \{\langle a \rangle\}$;
- (3) $\mathcal{G} = \{\langle a \rangle, \langle b \rangle\}$;
- (4) $\mathcal{G} = \{\langle a \rangle, \langle b \rangle, \langle c \rangle\}$;
- (5) $\mathcal{G} = \{\langle a, b \rangle\}$;
- (6) $\mathcal{G} = \{\langle a, b \rangle, \langle c \rangle\}$.

Remark 8.8. — The stabilisers of points of a relative outer space of a free product are described in [20], in terms of the elliptic free factors G_i of \mathcal{G} and the automorphisms groups $\text{Aut}(G_i)$. In cases (1)–(4), stabilisers of points are virtually \mathbb{Z}^k for some uniformly bounded k . In particular, any subgroup of the stabiliser in these cases is finitely presented.

Let $[\phi] \in \text{Out}(F_3)$. Let's first assume that our automorphism and all of its powers do not fix a point of some relative outer space of F_3 . As noticed (Remark 2.46), there is some relative outer space $\mathcal{O}(\mathcal{G})$ for which $[\phi]$ is irreducible. Note that under our assumption that no power of $[\phi]$ fixes a free splitting, we get that $\lambda_{\mathcal{O}}(\phi) > 1$. Therefore, cases (5) and (6) of the above list cannot appear under our assumptions, as the corresponding relative outer spaces are consisted by a single point (assuming volume equal to 1) and so there are no automorphisms of $\text{Out}(\mathcal{G})$ with $\lambda_{\mathcal{G}}(\phi) > 1$ (all such automorphisms fix a point of $\mathcal{O}(\mathcal{G})$).

In any other case, by Theorem 8.1, $C_{\mathcal{G}}(\phi)$ has a finite index subgroup which is a \mathbb{Z} -extension of $C_X(\phi)$, where $C_X(\phi)$ is the subgroup of $C_{\mathcal{G}}(\phi)$, acting trivially on some $X \in \mathcal{O}(\mathcal{G})$. By Remark 8.8, $C_X(\phi)$ is finitely presented. Therefore, $C_{\mathcal{G}}(\phi)$ is finitely presented, as a \mathbb{Z} -extension of a finitely presented group. By Proposition 8.3, the centraliser $C(\phi)$ of $[\phi]$ in $\text{Out}(F_3)$ has a finite index subgroup which is finitely presented (the group $C_{\mathcal{G}}(\phi)$), and therefore $C(\phi)$ is finitely presented itself. In particular, $C(\phi)$ is finitely generated.

We now assume that our automorphism has a power that fixes a point of some relative outer space of F_3 . There is a maximal such free splitting with

respect to domination, by Proposition 8.4, and there is a free factor system \mathcal{G} such that all such maximal free splittings belong to the same relative outer space, $\mathcal{O}(\mathcal{G})$. We deal with the possible cases for \mathcal{G} as enumerated above.

In case (1), $[\phi^k]$ fixes a point of CV_3 , then $[\phi]$ has finite order, and by [25], $C(\phi)$ is finitely presented.

In cases (2)–(4), $[\phi^k]$ fixes a point T of the corresponding relative outer space. By the description of stabilisers of points in [20], it is easy to see that $[\phi^k]$ (and so $[\phi]$) has linear growth as an automorphism of $\text{Out}(F_3)$ and so the result follows by Theorem 8.6.

For case (5), note that the corresponding relative outer space is a single point, hence preserving this kind of free factor system is equivalent to preserving the corresponding splitting. Hence, by Proposition 8.4, since all periodic maximal free splittings belong to the same relative outer space, there is a unique one. Since the $[\phi]$ image of a periodic splitting is again periodic, this means that the splitting is actually fixed by $[\phi]$.

We switch now to elements of $\text{Aut}(F_3)$. Let $\Phi \in [\phi]$ which actually fixes H (not just up to conjugacy). In other words, $\Phi(H) = H$. Consider the restriction Φ_H of Φ on H , which induces an element of $\text{Out}(H)$. If Φ_H has finite order as an outer automorphism, then it is easy to see that Φ has linear growth and so, as before, $C(\phi)$ is finitely generated by Theorem 8.6.

So, let's assume now that Φ_H has infinite order as an outer automorphism. The subgroup of $\text{Aut}(F_3)$ projecting to the centraliser $C(\phi)$ in $\text{Out}(F_3)$ is $C = \{\Theta \in \text{Aut}(F_3) : [\Theta, \Phi] \in \text{Inn}(F_3)\}$. We will show that C is finitely generated, which will imply that $C(\phi)$ is finitely generated.

By Proposition 8.4, if $\Theta \in C$, then $[\Theta] \in \text{Out}(F_3)$ fixes the conjugacy class of H , so we have a well-defined homomorphism $\pi: C \rightarrow \text{Out}(H)$. It is easy to see that the image of π is in fact contained in the centraliser of $[\Phi_H]$ in $\text{Out}(H)$, which, by Theorem 8.7, is virtually cyclic. Therefore $C^0 = \pi^{-1}([\Phi_H])$ is a finite index subgroup of C . Hence, it is enough to show that C^0 is finitely generated.

Let $\Theta \in C^0$. We assume without loss, up to composing with an inner automorphism of F_3 , that $\Theta(H) = H$. As $\Theta \in C^0$, the restriction of Θ on H , which we denote by Θ_H , is of the form $\Theta_H = \Phi_H^k \text{ad}(h)$, where $\text{ad}(h) \in \text{Inn}(H)$, for some $k \in \mathbb{Z}$ and $h \in H$. Therefore, if we denote by C^1 the subgroup of C^0 of those automorphisms acting as the identity on H , we get that C^0 is generated by the generators of C^1 , Φ_H , and the generators of $\text{Inn}(H)$ (which is clearly finitely generated). In particular, it is enough to show that C^1 is finitely generated.

Recall that we are working with a free basis $\{a, b, c\}$, with $H = \langle a, b \rangle$. Since $\Phi(H) = H$, we must have $\Phi(c) = h_1 c^\epsilon h_2$, where $\epsilon \in \{-1, 1\}$ and $h_1, h_2 \in H$, and a similar equation holds for elements of C^1 . Up to passing to a finite index subgroup C^2 of C^1 , we can assume that $\Theta(c) = xcy$, $x, y \in H$. As we pass to finite index subgroup, it is clear that it is enough to prove that C^2 is finitely generated.

Since $\Theta \in C^2 < C^1$, hence $\Theta_H = \text{Id}_H$, we get that $\Phi\Theta(a) = \Theta\Phi(a)$ and $\Phi\Theta(b) = \Theta\Phi(b)$. The remaining part of the proof is to write down the equations corresponding to $\Phi\Theta(c) = \Theta\Phi(c)$, which is equivalent to the fact that Φ and Θ commute (under our assumptions that Θ acts as the identity of H and Φ preserves H , it is clear that Φ and Θ commute up to inner automorphism if and only if they genuinely commute).

We have:

$$\begin{aligned}\Phi\Theta(c) &= \Phi(xcy) = \Phi(x)h_1 c^\epsilon h_2 \Phi(y), \\ \Theta\Phi(c) &= \Theta(h_1 c^\epsilon h_2) = h_1(xcy)^\epsilon h_2.\end{aligned}$$

Let's first assume that $\epsilon = 1$. In this case, the automorphisms Φ, Θ commute if and only if

$$\begin{cases} \Phi(x)h_1 = h_1x \\ \Phi(y)h_2 = h_2y \end{cases} \implies \begin{cases} \Phi(x) = h_1xh_1^{-1} \\ \Phi(y) = h_2^{-1}yh_2. \end{cases}$$

Note that it is well-known that the subgroups $S_{\Phi, h} = \{z : \Phi(z) = hzh^{-1}\}$ is finitely generated for every Φ and every $h \in H$ (for example see [4] — since $S_{\Phi, h}$ is just the fixed subgroup of Φ composed with an inner automorphism). In our case, as any Θ with the requested properties is uniquely determined by $x \in S_{\Phi, h_1}, y \in S_{\Phi, h_2^{-1}}$, the above equations identify the subgroup C^2 with the product of S_{Φ, h_1} and $S_{\Phi, h_2^{-1}}$, which means that it is finitely generated. Therefore, the proof concludes in this case.

In case $\epsilon = -1$, the automorphisms commute if and only if

$$\Phi(x)h_1 = h_1y^{-1} \quad \text{and} \quad h_2\Phi(y) = x^{-1}h_2$$

which is equivalent to

$$\begin{cases} \Phi^2(y) = \Phi(h_2)h_1^{-1}yh_1(\Phi(h_2))^{-1} \\ x = h_2\Phi(y^{-1})h_2^{-1} \end{cases} \implies \begin{cases} y \in S_{\Phi^2, \Phi(h_2)h_1^{-1}} \\ x = h_2\Phi(y^{-1})h_2^{-1} \end{cases}$$

and the thesis follows as above, since $S_{\Phi^2, \Phi(h_2)h_1^{-1}}$ is finitely generated.

Case (6) is similar to case (5) and so we skip the details. \square

BIBLIOGRAPHY

- [1] N. ANDREW & A. MARTINO, “Centralisers of linear growth automorphisms of free groups”, 2022, <https://arxiv.org/abs/2205.12865>.
- [2] M. BESTVINA, M. FEIGHN & M. HANDEL, “Laminations, trees, and irreducible automorphisms of free groups”, *Geom. Funct. Anal.* **7** (1997), no. 2, p. 215-244.
- [3] ———, “The Tits alternative for $\text{Out}(F_n)$ I: Dynamics of exponentially-growing automorphisms”, *Ann. Math.* **151** (2000), no. 2, p. 517-623.
- [4] M. BESTVINA & M. HANDEL, “Train tracks and automorphisms of free groups”, *Ann. Math.* **135** (1992), no. 1, p. 1-51.
- [5] M. M. COHEN & M. LUSTIG, “The conjugacy problem for Dehn twist automorphisms of free groups”, *Comment. Math. Helv.* **74** (1999), no. 2, p. 179-200.
- [6] M. CULLER & J. W. MORGAN, “Group actions on \mathbb{R} -trees”, *Proc. Lond. Math. Soc.* **55** (1987), p. 571-604.
- [7] F. DAHMANI & R. LI, “Relative hyperbolicity for automorphisms of free products and free groups”, *J. Topol. Anal.* **14** (2022), no. 1, p. 55-92.
- [8] S. DOWDALL, I. KAPOVICH & C. J. LEININGER, “Dynamics on free-by-cyclic groups”, *Geom. Topol.* **19** (2015), no. 5, p. 2801-2899.
- [9] S. FRANCAVIGLIA & A. MARTINO, “Metric properties of outer space”, *Publ. Mat., Barc.* **55** (2011), no. 2, p. 433-473.
- [10] ———, “The isometry group of outer space”, *Adv. Math.* **231** (2012), no. 3-4, p. 1940-1973.
- [11] ———, “Stretching factors, metrics and train tracks for free products”, *Ill. J. Math.* **59** (2015), no. 4, p. 859-899.
- [12] ———, “Displacements of automorphisms of free groups I: Displacement functions, minpoints and train tracks”, *Trans. Am. Math. Soc.* **374** (2021), p. 3215-3264.
- [13] ———, “Displacements of automorphisms of free groups II: Connectedness of level sets”, *Trans. Am. Math. Soc.* **375** (2022), p. 2511-2551.
- [14] S. FRANCAVIGLIA, A. MARTINO & D. SYRIGOS, “The minimally displaced set of an irreducible automorphism is locally finite”, *Glas. Mat., III. Ser.* **55** (2020), no. 2, p. 301-336.
- [15] ———, “The minimally displaced set of an irreducible automorphism of F_N is co-compact”, *Arch. Math.* **116** (2021), no. 4, p. 369-383.
- [16] M. GROMOV, “Asymptotic invariants of infinite groups”, in *Geometric group theory, Vol. 2 (Sussex, 1991)* (G. A. Niblo & M. A. Roller, eds.), London Mathematical Society Lecture Note Series, vol. 182, Cambridge University Press, 1993, p. 1-295.
- [17] V. GUIRARDEL, “Approximations of stable actions on \mathbb{R} -trees”, *Comment. Math. Helv.* **73** (1998), p. 89-121.
- [18] V. GUIRARDEL & C. HORBEZ, “Algebraic laminations for free products and arational trees”, *Algebr. Geom. Topol.* **19** (2019), no. 5, p. 2283-2400.
- [19] ———, “Measure equivalence rigidity of $\text{Out}(F_N)$ ”, 2021, <https://arxiv.org/abs/2103.03696>.
- [20] V. GUIRARDEL & G. LEVITT, “The outer space of a free product”, *Proc. Lond. Math. Soc. (3)* **94** (2007), no. 3, p. 695-714.
- [21] R. GUPTA, “Loxodromic Elements in the relative free factor complex”, *Geom. Dedicata* **196** (2018), p. 91-121.
- [22] M. HANDEL & L. MOSHER, *Axes in outer space*, vol. 213, Memoirs of the American Mathematical Society, vol. 1004, American Mathematical Society, 2011, vi+104 pages.
- [23] C. HORBEZ, “Hyperbolic graphs for free products, and the Gromov boundary of the graph of cyclic splittings”, *J. Topol.* **9** (2016), p. 401-450.

- [24] ———, “The boundary of the outer space of a free product”, *Isr. J. Math.* **221** (2017), p. 179-234.
- [25] S. KALAJDŽIEVSKI, “Automorphism group of a free group: centralizers and stabilizers”, *J. Algebra* **150** (1992), no. 2, p. 453-502.
- [26] S. KRSTIĆ, M. LUSTIG & K. VOGTMANN, “An equivariant Whitehead algorithm and conjugacy for roots of Dehn twist automorphisms”, *Proc. Edinb. Math. Soc., II. Ser.* **44** (2001), no. 1, p. 117-141.
- [27] G. LEVITT & M. LUSTIG, “Irreducible automorphisms of F_n have north-south dynamics on compactified outer space”, *J. Inst. Math. Jussieu* **2** (2003), no. 1, p. 59-72.
- [28] J. P. MUTANGUHA, “Irreducible nonsurjective endomorphisms of F_N are hyperbolic”, *Bull. Lond. Math. Soc.* **52** (2020), no. 5, p. 960-976.
- [29] W. PARRY, “Axioms for translation length functions”, in *Arboreal group theory (Berkeley, CA, 1988)*, Mathematical Sciences Research Institute Publications, vol. 19, Springer, 1991, p. 295-330.
- [30] F. PAULIN, “The Gromov topology on \mathbb{R} -trees”, *Topology Appl.* **32** (1989), no. 3, p. 197-221.
- [31] J.-P. SERRE, *Trees*, Springer, 1980, ix+142 pages.
- [32] D. SYRIGOS, “Irreducible laminations for IWIP automorphisms of free products and centralisers”, 2014, <https://arxiv.org/abs/1410.4667>.
- [33] ———, “Asymmetry of outer space of a free product”, *Commun. Algebra* **46** (2018), no. 8, p. 3442-3460.

Manuscrit reçu le 27 avril 2022,
révisé le 22 novembre 2023,
accepté le 12 avril 2024.

Stefano FRANCAVIGLIA
University of Bologna (Italy)
stefano.francaviglia@unibo.it
Armando MARTINO
University of Southampton (UK)
a.martino@soton.ac.uk
Dionysios SYRIGOS
University of Southampton (UK)
d.syrigos@soton.ac.uk