DISPLACEMENTS OF AUTOMORPHISMS OF FREE GROUPS II: CONNECTIVITY OF LEVEL SETS AND DECISION PROBLEMS

STEFANO FRANCAVIGLIA

Dipartimento di Matematica of the University of Bologna

ARMANDO MARTINO

Mathematical Sciences, University of Southampton

Abstract. This is the second of two papers in which we investigate the properties of displacement functions of automorphisms of free groups (more generally, free products) on the Culler-Vogtmann Outer space CV_n and its simplicial bordification. We develop a theory for both reducible and irreducible autormorphisms. As we reach the bordification of CV_n we have to deal with general deformation spaces, for this reason we developed the theory in such generality. In first paper [13] we studied general properties of the displacement functions, such as well-orderability of the spectrum and the topological characterization of min-points via partial train tracks (possibly at infinity).

This paper is devoted to proving that for any automorphism (reducible or not) any level set of the displacement function is connected. Here, by the "level set" we intend to indicate the set of points displaced by *at most* some amount, rather than exactly some amount; this is sometimes called a "sub-level set".

As an application, this result provides a stopping procedure for brute force search algorithms in CV_n. We use this to reprove two known algorithmic results: the conjugacy problem for irreducible automorphisms and detecting irreducibility of automorphisms.

Note: the two papers were originally packed together in the preprint [12] We decided to split that paper following the recommendations of a referee.

Contents

E-mail addresses: `stefano.francaviglia@unibo.it, A.Martino@soton.ac.uk`.

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

1. Introduction

We consider F_n the free group of rank n, usually with a basis B (a free generating set). We are interested in the automorphism group, $\mathrm{Aut}(F_n)$ and the Outer automorphism group, which is defined as $\mathrm{Out}(F_n) = \mathrm{Aut}(F_n)/\,\mathrm{Inn}(F_n)$.

That said, as the reader will notice, in this paper all results are about general deformation spaces, and our statements are of the form "let $[\phi] \in \mathrm{Out}(\Gamma)$" or "let $X \in \overline{\mathcal{O}_{\mathrm{gr}}(\Gamma)}^{\infty}$" and so on. Let's briefly explain the notation and why we need to work in such generality. The reason is that classical Culler-Vogtmann space CV_n is perfect for studying irreducible automorphisms, but if one is interested in possibly reducible automorphisms, some more general space is needed. If for instance an automorphism ϕ is represented by a simplicial map f on a finite graph X, it may happen that in X we have a collection of subgraphs $A_1, ..., A_k$ so that $\cup_i A_i$ is preserved by f. In this case it may be necessary to study both the invariant collection and the quotient obtained by collapsing any A_i to a point. So the typical object we have to deal with is a deformation space of finite unions of graphs of groups. Concretely, our proofs boil down to induction proofs where the inductive step needs to deal with both the (disconnected) collection $\cup_i A_i$ and the map(s) that f induces there, as well as the quotient graph of groups obtained by 'collapsing' the A_i in X, but keeping track of the fundamental group of the collapsed part; this leads to a graph of groups with trivial edge groups. So, even though our main focus is CV_n, it turns out to be no more complicated to deal with arbitrary free products and their deformation spaces, and our proofs need in fact to deal with the case of a finite graph of groups, with trivial edge groups, which *may not be connected*. This is what Γ refers to. We direct the reader to Section 3, and in particular Remark 3.8 for more discussion on this.

Nevertheless, since our general theorems specialise to results about classical CV_n and $\mathrm{Out}(F_n)$, in this introduction we will stick as much as possible to that classical setting.

In recent years there has been a great deal of attention given to the Lipschitz metric on CV_n, see [1], [2], [3] for instance. It has been considered even more generally in [25].

In the first part, [13], we proved results concerning the Lipschitz metric on a class of deformation spaces, of which a key example is the Culler-Vogtmann space of a free group, CV_n. We showed that, given an automorphism of a free group, the points of minimal displacement - for a given automorphism, the distance between a point in CV_n and its image - correspond to the points which support partial train track maps, thus generalizing known results about irreducible automorphisms.

In [23] it is shown that, in the irreducible case, these points of minimal displacement (equivalently, the points which support train track maps) form a connected subset of CV_n and this is used to solve the conjugacy problem. Our results here arise out of a desire to generalize those results to the reducible case, and we also employ Peak Reduction as a key tool.

The generalization of this result for arbitrary, possible reducible, automorphisms, requires some care, however. To start with, given an automorphism ϕ, one can define the infimum over all displacements of points in CV_n, to obtain $\lambda(\phi)$. However, in general there might exist no points in CV_n which are displaced by this amount. Our point of view is to pass to the simplicial bordification of CV_n, otherwise known as the free splitting complex, $\mathcal{FS}_n$. One can define displacements for points in $\mathcal{FS}_n$, though in some cases these will be infinite. (A point in CV_n is a marked graph, and a point in $\mathcal{FS}_n$ arises by collapsing a subgraph. These induced points will have finite displacement exactly when the subgraphs are ϕ-invariant[1]). However, the infimum of all displacements of points in $\mathcal{FS}_n$ will, in general, be less than those in CV_n.

Bearing in mind these complications, and the fact that in the whole paper we work with more general deformation spaces, our main Theorem, which is a special case of Theorem 5.3, is the following:

Theorem (Connectivity of Level Sets). Let $[\phi] \in \mathrm{Out}(F_n)$. Let $\lambda(\phi)$ be the infimum of displacements, with respect to the Lipschitz metric, of all points in CV_n. Then the set of points of $\mathcal{FS}_n$ which are displaced by exactly $\lambda(\phi)$, is connected.

Remark. As stated in the abstract we generally intend the "level set" to be the set of points displaced by *at most* some amount; this is sometimes referred to as a "sub-level set". The subsequent Theorem has precisely this kind of statement. Hence the statement above is more properly a statement about the minimally displaced set, although our proofs deal with both at the same time.

However, note that $\lambda(\phi)$ is the infimum of displacements in CV_n; however, it might not be the infimum of displacements of points in $\mathcal{FS}_n$.

Moreover, our techniques allow us to *regenerate* paths from $\mathcal{FS}_n$ to CV_n without disturbing the displacements by very much. Hence, as part of the same Theorem 5.3, we also prove:

Theorem Let $[\phi] \in \mathrm{Out}(F_n)$. Let $\lambda(\phi)$ be the infimum of displacements, with respect to the Lipschitz metric, of all points in CV_n. Then, for any $\varepsilon > 0$ the set of points of CV_n which are displaced by at most $\lambda(\phi) + \varepsilon$, is connected.

In the case where the automorphism is irreducible, there are points in CV_n which are displaced by exactly the minimum, $\lambda(\phi)$. Moreover, every point on the boundary has infinite displacement (Remark 3.19) and hence the connectivity of the level set becomes a statement about CV_n, as in Corollary 5.4:

Corollary Let $[\phi] \in \mathrm{Out}(F_n)$ be irreducible. Let $\lambda(\phi)$ be the infimum of displacements, with respect to the Lipschitz metric, of all points in CV_n. Then the set of points of CV_n which are displaced by $\lambda(\phi)$, is connected.

Remark 1.1. Given an automorphism, ϕ, of the free group, one can construct a relative train track representative for ϕ. The quantity, $\lambda(\phi)$ is then simply the maximum Perron-Frobenius eigenvalue of any stratum.

More generally, if we are given a ϕ-invariant free factor system, then one can build a relative train track representative of ϕ which sees this free factor system as an invariant subgraph. There is a corresponding deformation space where one collapses this

[1] See [13] or Section 3.4 for more details on this point.

subgraph, and the minimum displacement in that deformation space is the maximum Perron-Frobenius eigenvalue of any stratum *above* the invariant subgraph.

We can think of $\mathcal{FS}_n$ as a union of such deformation spaces, with the displacements being infinite when the collapsed object is not invariant. This is why the minimum displacement in $\mathcal{FS}_n$ need not be equal to that in CV_n - they are different if one can collapse an invariant subgraph which carries all the maximum Perron-Frobenius eigenvalues.

A simple example is the following. Consider this automorphism, ϕ, of the free group on a, b, c:

$$\begin{array}{rcl} c & \mapsto & ca \\ b & \mapsto & ba \\ a & \mapsto & aba \end{array}$$

This is then a relative train track map, with two strata, the bottom one given by a, b and the top one by c.

Let λ be the larger eigenvalue of the matrix

$$\begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}.$$

This is the Perron-Frobenius eigenvalue of the bottom stratum, with the top stratum having 1 as its Perron-Frobenius eigenvalue. It is then easy to see that $\lambda(\phi) = \lambda$, but there are points in $\mathcal{FS}_n$ which are fixed by ϕ and so have multiplicative displacement 1; namely, take the point obtained by collapsing a, b. That is, the graph of groups with one edge, one vertex, a trivial edge group and a vertex group generated by a and b.

Naturally, since our results generalize those of [23], we obtain a solution of the conjugacy problem for irreducible automorphisms in the same way. However, it seems that our techniques allow for a more elementary interpretation, and also opens up the possibilty for attempting the algorithm in the reducible case. However, there are further complications that arise in the reducible case, due to the fact that the minimally displaced set enters the thin part, and so we do not easily obtain bounds on the number of points we need to enumerate.

In any case, we can describe this algorithm in the irreducible case, with explicit constants, rather straightforwardly. Moreover, we also provide an algorithm to detect irreducibility; this result was first proved in [21] and improved in [22] (also, see [6] and [7] which give another algorithm for detecting irreducibility).

Finally, it may be worth noting that, thank to the generality of Theorem 5.3, our algorithms easily generalise to a class of groups bigger than just free groups; concretely, groups of the kind $G_1 * \cdots * G_p * F_n$ where the G_i are finite groups (see Theorem 2.11 and Section 9.1).

Acknowledgements: We would like to thank both the Università di Bologna and the Universitat Politécnica de Catalunya, for their hospitality during several visits. We would also like to thank the referee of the earlier paper [12], as well as the current referee for their patience and extraordinary efforts in improving this paper, with their many helpful comments and suggestions.

2. Algorithms

In order to motivate the detailed discussion which follows, we provide here the two algorithms for solving conjugacy in the irreducible case and for detecting irreducibility. We present these algorithms as naively as possible, in order to make them more accessible. That is, one could understand and implement them without any knowledge of the Lipschitz metric, Culler-Vogtmann space or partial train track maps. As such we have made no attempt to streamline the algorithms in any way; they are brute force searches in an exponential space.

However, we would stress that our point of view is fundamentally that these procedures would be better run as path searches in Culler-Vogtmann space, enumerating optimal maps and calculating displacements via candidates. That abundance of terminology would make the algorithms much harder to describe, so we instead translate everything to a more manageable setting; bases of F_n and generating sets for $\mathrm{Out}(F_n)$. However, the technical point of view is more helpful in developing an intuition of the processes and is likely the way to vastly improve the algorithmic complexity.

Let us now describe our algorithms, whose correctness is proved at the end of the paper. First, we recall some terminology. In order to work algorithmically with $\mathrm{Out}(F_n)$ we need a generating set. The best known of these is the set of Nielsen generators, but it is more convenient for us to work with the following:

Definition 2.1 (CMT Automorphisms, [15] and [14])**.** A *CMT* automorphism of F_n is one that is induced by a change of maximal tree. More precisely, fix a basis, B, of F_n. Let $R = R_B$ be the marked rose corresponding to B; that is, R is a graph with one vertex v and n edges called petals, and we have a fixed isomorphism between F_n and $\pi_1(R, v)$ where each element of B corresponds to a petal of R under this isomorphism.

Let X be a graph with fundamental group of rank n, and let T, T' be two maximal trees of X. Collapsing T and T' we obtain two roses R_T and $R_{T'}$. Let $\rho_T, \rho_{T'}$ be the corresponding projections from X to $R_T, R_{T'}$, and let $\alpha_T, \alpha_{T'}$ be homeomorphisms from R to $R_T, R_{T'}$ respectively. Then the (outer) automorphism induced by changing the maximal tree from T to T' is the (homotopy class of the) map $\alpha_{T'}^{-1}\rho_{T'}{\rho_T}^{-1}\alpha_T : R \to R$, where the inverse denotes a homotopy inverse.

Thus the set of CMT automorphisms of F_n, relative to B, is the set of all such change of maximal tree automorphisms.

The set of CMT automorphisms relative to B includes all Whitehead automorphisms, (see [15, Theorem 5.5], and [27]) and is a finite set which generates $\mathrm{Out}(F_n)$. Note also that in case $T = T'$, by varying $\alpha_T, \alpha_{T'}$ we obtain all graph automorphisms of R, including inversions of generators, which therefore are CMT automorphisms. Also, note that the property of being CMT, depends on a fixed chosen basis, B, of F_n (the petals of R).

Remark 2.2. The definition of CMT automorphisms just given is close to that given in [27], but there is an alternative definition via CV_n as follows. We call two marked roses, R_1, R_2 adjacent if there is a simplex, Δ, in CV_n, admitting faces, Δ_1 and Δ_2 such that R_i belongs to Δ_i. This is equivalent to saying there is a marked graph X, admitting two maximal trees, T_1, T_2 whose collapse produces the marked graphs, R_1, R_2, respectively. Then,

$$CMT_R(F_n) = \{[\phi] \in \mathrm{Out}(F_n) \ : \ \phi(R) \text{ is adjacent to } R\}.$$

This is the same as the previous definition by setting $R = R_B$, as above.

Next we need a notion of size of an automorphism, which will provide a termination criterion for our algorithms.

Definition 2.3. Let $[\phi] \in \mathrm{Out}(F_n)$, and let B be a basis of F_n. Define $||\phi||_B$ to be $\sup_{1 \neq g \in F_n} \frac{||\phi g||_B}{||g||_B}$, where $||g||_B$ denotes the cyclic reduced length of g with respect to B. This supremum is a maximum and is realised by an element of cyclic length ≤ 2 (see Lemma 4.1).

Remark 2.4. Note that for any constant, C, there are only finitely many $[\phi] \in \mathrm{Out}(F_n)$ such that $||\phi||_B \leq C$ (see [8], Lemma 4.10).

This also follows since $||\phi||_B$ is really $\Lambda(R, \phi(R))$ in disguise (see section 3.4) where $R = R_B$ is the uniform marked rose corresponding to B with all edges length $1/n$ (so that R has volume 1). The remark then follows from the fact that, for any given C, there are only finitely many (marked, volume 1) roses, R_1 such that $\Lambda(R, R_1) \leq C$, and the stabiliser of any point, and in particular R, is finite.

Our first application is then as follows. (See Section 9 for the proof.)

Theorem 2.5. *The following is an algorithm to determine whether two irreducible automorphisms are conjugate.*

Let $[\phi], [\psi]$ be two irreducible outer automorphisms of F_n, and B a basis of F_n.

- *Choose any $\mu > \max\{||\phi||_B, ||\psi||_B\}$.*
- *Inductively construct a finite set, $S = S_{\phi,\mu}$, as follows (which depends on both ϕ and μ):*
  - *Start with $S_0 = \{\phi\}$.*
  - *Set $K = n(3n-3)\mu^{3n-1}$.*
  - *Inductively put S_{i+1} to be all possible automorphisms $\zeta\phi_i\zeta^{-1}$, where ϕ_i is any element of S_i, ζ is any CMT automorphism, subject to the constraint that $||\zeta\phi_i\zeta^{-1}||_B \leq K$. (Since the identity is a CMT automorphism according Definition 2.1, we have $S_{i-1} \subseteq S_i$).*
  - *End this process when $S_i = S_{i+1}$, and let this final set be S.*
- *Then ψ is conjugate to ϕ if and only if $\psi \in S$.*

Of course, one would like to also be able to decide when an automorphism is irreducible when it is given by images of a basis, for instance. In order to do so, we recall the definition of irreducibility.

Definition 2.6 (See [5])**.** An (outer) automorphism, $[\phi]$ of F_n is called *reducible* if there are free factors, $F_{n_1}, \ldots, F_{n_k}, F_{n_\infty}$ such that $F_n = F_{n_1} * \ldots F_{n_k} * F_{n_\infty}$ and each $\phi(F_{n_i})$ is conjugate to $F_{n_{i+1}}$ (subscripts taken modulo k). If $k = 1$ we further require that $F_{n_\infty} \neq 1$. (In general $\phi(F_{n_\infty})$ is not conjugate to F_{n_∞}). Otherwise $[\phi]$ is called irreducible.

Equivalently, $[\phi]$ is reducible if it is represented by a homotopy equivalence f, on a core graph X, such that X has a proper subgraph X_0, with non-trivial fundamental group, such that $f(X_0) = X_0$. (Being represented by f means that there is an isomorphism, $\tau : F_n \to \pi_1(X)$ such that $\phi = \tau^{-1} f_* \tau$).

We add the following, which constitutes an obvious way that one can detect irreducibility by inspection.

Definition 2.7. Consider F_n with basis B and let $[\phi]$ be an outer automorphism of F_n. We say that $[\phi]$ is visibly reducible with respect to B, or simply visibly reducible, if there exist disjoint subsets $B_1, \ldots, B_k$ of B such that $\phi(\langle B_i \rangle)$ is conjugate to $\langle B_{i+1} \rangle$ (with subscripts taken modulo k). If $k = 1$ we also require that $B_1 \neq B$.

More generally, we say that a homotopy equivalence on the rose is visibly reducible if it is visibly reducible with respect to the basis given by petals.

This is, in fact, easy to check by classical methods due to Stallings, [26].

Lemma 2.8. *If $[\phi]$ is visibly reducible, then it is reducible. Moreover, there is an algorithm to determine if $[\phi]$ is visibly reducible with respect to B.*

Proof. The first statement is clear, since each subset of a basis generates a free factor, and disjoint subsets generate complementary free factors. Since there are only finitely many subsets to check, we simply need to determine if the conditions that $\phi(\langle B_i \rangle)$ is conjugate to $\langle B_{i+1} \rangle$ hold. But this can readily be checked since two subgroups of a free group are conjugate if and only if the core of their Stallings graphs are equal, [26]. □

We can now describe our second algorithm. (See Section 9 for the proof.)

Theorem 2.9. *The following is an algorithm to determine whether or not an outer automorphism of F_n is irreducible.*

Let $[\phi] \in \mathrm{Out}(F_n)$, and B a basis of F_n. Construct $S = S_\phi$ as above. Namely,

- *Choose any $\mu > ||\phi||_B$.*
- *Inductively construct the finite set, $S = S_{\phi,\mu}$:*
  - *Start with $S_0 = \{\phi\}$.*
  - *Set $K = n(3n-3)\mu^{3n-1}$.*
  - *Inductively put S_{i+1} to be all possible automorphisms $\zeta\phi_i\zeta^{-1}$, where ϕ_i is any element of S_i, ζ is any CMT automorphism, subject to the constraint that $||\zeta\phi_i\zeta^{-1}||_B \leq K$. (Since the identity is a CMT automorphism according Definition 2.1, we have $S_{i-1} \subseteq S_i$).*
  - *End this process when $S_i = S_{i+1}$, and let this final set be S.*
- *Let S^+ be the set of all possible automorphisms $\zeta\phi_i\zeta^{-1}$, where ϕ_i is any element of S, ζ is any CMT automorphism, with no other constraint.*
- *If some $\psi \in S^+$ is visibly reducible with respect to B, then ϕ is reducible. Otherwise, ϕ is irreducible.*

Remark 2.10. In both Theorems 2.5, and 2.9, the set S is a subset of the set of automorphism classes with $||\phi||_B \leq K$, which is finite (Remark 2.4). Therefore both algorithms stop in a finite, effectively computable, time.

We also explain, in Section 9.1, how to implement essentially the same algorithms in the case where one has a free product of finite groups with a free group.

Theorem 2.11. *Let $G = G_1 * \ldots * G_p * F_n$ be a free product where the G_i are finite groups and F_n is a free group of rank n. Let $\mathcal{G} = \{\{G_i\}, n\}$ be the free splitting induced from the finite groups G_i. Then the following problems are algorithmically decidable:*

- *Deciding whether a given $[\phi] \in \mathrm{Out}(G)$ is irreducible (relative to $\mathcal{G}$),*
- *Deciding whether two $\mathcal{G}$-irreducible automorphisms, $[\phi], [\psi]$ are conjugate in $\mathrm{Out}(G)$.*

Remark 2.12. Note that any automorphism of G preserves $\mathcal{G}$, so $\mathrm{Out}(G) = \mathrm{Out}(\mathcal{G})$ in this case.

3. Preliminaries and notation (from [13])

Throughout the paper, we use the definitions and notation of [13]. We briefly recall them here, referring the reader to [13] for a detailed discussion.

Before of that, we wish to recall the reasons for giving new definitions and working in a so general setting. Our principal motivation was to study outer automorphisms of free groups that are possibly reducible. This naturally leads to consider simplicial bordifications of Culler-Vogtmann Outer spaces. Namely, if Γ is a marked graph with fundamental group F_n – the rank-n free group – then any automorphism $\phi : F_n \to F_n$ can be represented by a simplicial map $f : \Gamma \to \Gamma$. When ϕ is reducible, it may happen that Γ exhibits a collection $\Gamma_1, \dots, \Gamma_k$ of f-invariant subgraphs. In order to study the properties of ϕ it may help to collapse such sub-graphs to points. So one is naturally lead to study two kind of deformation spaces: that of actions on trees with possibly non-trivial vertex stabilisers (when we collapse the Γ_i's) and product of such spaces (when we consider the restriction of f to the invariant collection).

Summing up, the typical object we need to understand is a disjoint union of metric trees, where a group G acts with possibly non-trivial vertex stabilisers. We therefore work in such a general setting, as developed in [13], but the reader is invited to keep in mind the case of CV_n and its bordification.

3.1. **Splittings, $\mathcal{G}$-trees, outer spaces, and automorphisms.** F_n will always denote the free group of rank n. We will consider groups G equipped with free a splitting $G = G_1 * \cdots * G_p * F_n$. We do not assume G_i is indecomposable, and our main interest is indeed when G itself is a free group.

Definition 3.1. Given a group G, a *free splitting* $\mathcal{G}$ of G is a pair $(\{G_1, \dots, G_p\}, n)$ where $\{G_i\}$ is a collection of subgroups of G (and $n \in \mathbb{N}$) such that $G = G_1 * \cdots * G_p * F_n$. Two splittings $\mathcal{G} = (\{G_i\}, n)$ and $\mathcal{H} = (\{H_i\}, m)$ of G are *of the same type* if $m = n$ and, up to reordering and conjugacy of the G_i, they have the same factor subgroups. That is, we do not require the named (conjugacy class of the) free group factor at the end to be preserved. The Kurosh *rank* of the splitting is $n + p$. We say that $\mathcal{H}$ is a *sub-splitting* of $\mathcal{G}$ if every H_i decomposes as $H_i = G_{j_1} * \cdots * G_{j_{l_i}} * F_{s_i}$ and $n = m + \sum s_i$.

Remark 3.2. We admit the trivial splitting $G = F_n$, $(\emptyset, n)$. That is the splitting with no free factors groups. In this case our discussion will amount to considering the free group F_n and the classical Culler-Vogtmann Outer space CV_n.

Remark 3.3. Free splittings are also referred to as *free factor systems* in the literature (originally introduced in [4], and also used in [18], [19] and [20]). The viewpoint of [13] and the present paper is that of taking a fixed free factor system - a free splitting - and studying its deformation space. We refer to [13] for more details. We just notice here that a "splitting" in general refers to any action on a tree and the induced graph of groups decomposition, but no confusion should arise since all of the splittings we consider are "free", in the sense that the edge stabilisers in the tree are trivial (equivalently, the splitting which arises is a free factor system).

Definition 3.4. Given a group G endowed with a free splitting $\mathcal{G} = (\{G_i\}, n)$, a simplicial *$\mathcal{G}$-tree* is a metric simplicial tree, endowed with a faithful simplicial G-action via isometries, trivial edge-stabilisers, and such that for every G_i there is exactly one orbit of vertices whose stabiliser is conjugate to G_i. Such vertices are called *non-free.* Other vertices (those with trivial stabilisers) are called *free* vertices.

A *$\mathcal{G}$-graph* is a finite connected metric graph of groups X whose *topological* fundamental group is F_n, with trivial edge-groups, and endowed with a G-marking, that is, there is a fixed isomorphism between its fundamental group – *as graph of groups* – and G, such that the splitting given by vertex groups is equivalent to $\mathcal{G}$.

If the splitting $\mathcal{G}$ of a group G is clear from the context, we may use notation G-tree instead of $\mathcal{G}$-tree. Same for graph.

The rank of a $\mathcal{G}$-tree (or graph) is the Kurosh rank of the splitting (as defined in Definition 3.1).

A $\mathcal{G}$-tree is *mimimal* if it has no proper G-invariant sub-tree (in particular, it has no free leaves, and G acts without global fixed point). A graph of groups with trivial edge-groups is a *core graph* if its leaves (if any) have non-trivial vertex group. Given a graph of groups X, with trivial edge groups and non-trivial fundamental group (as graph of groups), we define $\operatorname{core}(X)$ to be the maximal core sub-graph of X. If X has trivial fundamental group (as graph of groups) we define $\operatorname{core}(X)$ to be just a point of X. We say that $\operatorname{core}(X)$ is trivial when it is a point, namely when X is topologically a tree with at most one non-free vertex.

Bass-Serre theory provides a correspondence between minimal $\mathcal{G}$-tree and core $\mathcal{G}$-graphs, so one can equivalently works either with trees or graphs. The equivalence tree-graph is made explicit as follows: Given a minimal $\mathcal{G}$-tree, its quotient by the G-action is a core $\mathcal{G}$-graph.

Two simplicial $\mathcal{G}$-trees are considered *equivalent* if there is a G-equivariant isometry between them, and the corresponding notion of equivalence is given for graphs.

In some setting it will be more convenient using trees, in others, graphs. For this purpose we introduce the following notation.

Notation 3.5 (Tilde-underline notation)**.** Let $\mathcal{G}$ be a free splitting of a group G. If X is a $\mathcal{G}$-graph, then $\widetilde{X}$ denotes its universal covering, which is a $\mathcal{G}$-tree. As usual, if $x \in X$ then $\widetilde{x}$ will denote a lift of x in $\widetilde{X}$. The same for subsets: if $A \subset X$ is connected then $\widetilde{A} \subset \widetilde{X}$ is a connected component of the preimage of A. On the converse situation, if T is a $\mathcal{G}$-tree with finite edge-orbits, we denote by $\underline{T}$ the quotient $\mathcal{G}$-graph. Same notation for points and subsets. So, $\underline{\widetilde{X}} = X$ for both graphs and trees.

Notation 3.6. If X is a connected graph of groups with trivial edge groups, by a X-graph (or tree) we mean a $\pi_1(X)$-graph (resp. tree), that is to say, a $\mathcal{G}$-graph (resp. tree) where G is the fundamental group of X as graph of groups, and $\mathcal{G}$ is the splitting of G given by vertex groups.

If $\Gamma = \sqcup\Gamma_i$ is a disjoint finite union of finite graphs of groups with trivial edge-groups, a Γ-graph is a disjoint finite union $X = \sqcup X_i$ of Γ_i-graphs (and a Γ-forest is a union of Γ_i-trees).

We introduce now the outer space of a splitting (see [11, 17, 13] for details). Let G be a group and $\mathcal{G}$ be a splitting of G. The (projectivized) outer space of G, relative to the splitting $\mathcal{G}$, consists of (projective) classes of minimal simplicial metric $\mathcal{G}$-trees X with no redundant vertex (i.e. free and two-valent) and such that the G-action is by isometries.[2]

We use the notation $\mathcal{O}(G;\mathcal{G})$ or simply $\mathcal{O}(\mathcal{G})$ to indicate the outer space of G relative to $\mathcal{G}$. We use $\mathbb{P}\mathcal{O}(G;\mathcal{G})$ (or simply $\mathbb{P}\mathcal{O}(\mathcal{G})$) to indicate the projectivized outer space. For $X \in \mathcal{O}(\mathcal{G})$ we define its volume $\operatorname{vol}(X)$ as the sum of lengths of edges in $G\backslash X$. This is often referred to also as co-volume. The volume-one slice of $\mathcal{O}(\mathcal{G})$ is indicated by $\mathcal{O}_1(\mathcal{G})$[3].

[2]If $\mathcal{G}$ is the trivial splitting $G = F_n$, then $\mathcal{O}(\mathcal{G}) = CV_n$.

[3]We stress that the distinction between $\mathcal{O}(\mathcal{G})$ and $\mathbb{P}\mathcal{O}(\mathcal{G})$ is not crucial in our setting as we will mainly work with scale-invariant functions.

We defined $\mathcal{O}(\mathcal{G})$ as a space of **trees**, but we it will be often convenient to use **graphs** X so that $\widetilde{X} \in \mathcal{O}(\mathcal{G})$. Clearly the two viewpoints are equivalent. We introduce the following convention: when we want to consider spaces of graphs we add a "lower *gr*" to our notation:

$$\mathcal{O}_{\mathrm{gr}}(\mathcal{G}) = \{\mathcal{G}\text{-graph } X : \widetilde{X} \in \mathcal{O}(\mathcal{G})\}$$

The **spaces $\mathcal{O}(\mathcal{G})$ and $\mathcal{O}_{\mathbf{gr}}(\mathcal{G})$ are naturally identified via $X \leftrightarrow \widetilde{X}$.** In particular, they are completely interchangeable in all statements.

If X is a finite connected graph of groups with trivial edge-groups, and $\mathcal{S}$ is the splitting of $\pi_1(X)$ given by vertex-groups, then we set

$$\mathcal{O}(X) = \mathcal{O}(\pi_1(X); \mathcal{S}).$$

Let now $\mathcal{G}$ be a splitting of a group G, X be a $\mathcal{G}$-graph, and $\Gamma = \sqcup_i \Gamma_i$ be a sub-graph of X whose connected components Γ_i have non-trivial fundamental groups (as graphs of groups). Then Γ induces a sub-splitting $\mathcal{S}$ of $\mathcal{G}$ where the factor-groups H_j are either

- the fundamental groups $\pi_1(\Gamma_i)$, or
- the vertex-groups of non-free vertices in $X \setminus \Gamma$.

In this case will use the notation

$$\mathcal{O}(X/\Gamma) := \mathcal{O}(G; \mathcal{S}) \qquad \mathcal{O}(\Gamma) := \Pi_i \mathcal{O}(\Gamma_i)$$

(and similarly for $\mathcal{O}_{\mathrm{gr}}$). We tacitly identify $X = (X_1, \ldots, X_k) \in \mathcal{O}(\Gamma)$ with the labelled disjoint union $X = \sqcup_i X_i$. So an element of $\mathcal{O}(\Gamma)$ can be interpreted as a metric Γ-forest. The quotient of $\mathcal{O}(\Gamma)$ by the natural action of $\mathbb{R}^+$ is the projective outer space of Γ, and it is denoted by $\mathbb{P}\mathcal{O}(\Gamma)$. (Thus $\mathbb{P}\mathcal{O}(\Gamma)$ is not the product of the $\mathbb{P}\mathcal{O}(\Gamma_i)$'s.) The notion of volume extends to Γ-trees: If $X = (X_1, \ldots, X_k) \in \mathcal{O}(\Gamma)$ we set $\mathrm{vol}(X) = \sum_i \mathrm{vol}(X_i)$, and $\mathcal{O}_1(\Gamma)$ denotes the volume-one slice of $\mathcal{O}(\Gamma)$. We extend our notation and define define $\mathcal{O}(X/A)$ and $\mathcal{O}(A)$ also to the case where X is a non connected Γ-graph and $A \subset X$ is a sub-graph whose components have non-trivial fundamental groups.

Notation 3.7. In what follows we use the following convention:

- G will always be a group with a splitting $\mathcal{G} = (\{G_1, \ldots, G_p\}, F_n)$;
- $\Gamma = \sqcup \Gamma_i$ will always mean that Γ is a finite disjoint union of finite graphs of groups Γ_i, each with trivial edge-groups and non-trivial fundamental group $H_i = \pi_i(\Gamma_i)$, each H_i being equipped with the splitting given by the vertex-groups.

We set

$$\mathrm{rank}(\Gamma) = \sum_i \mathrm{rank}(\Gamma_i).$$

Remark 3.8. One should think G-statements as referring to classical outer space CV_n, $\mathcal{G}$-statements as referring to its simplicial bordification and deformation spaces of free products, and Γ-statements as general statements about more general deformation spaces, that come into play along the way of our rank-inductive strategy. More precisely, any Γ-statement specialises to a $\mathcal{G}$-statement (in the case where Γ is connected), to a G-statement (Γ connected and trivial splitting), and to a CV_n-statement (Γ is connected, the splitting is trivial, and $G = F_n$).

For this reason, the paper will contain mainly Γ-statements.

Notation 3.9. We will also consider moduli spaces with marked points. The moduli space of $\mathcal{G}$-trees with k marked points $p_1, \ldots, p_k$ (not necessarily distinct) is denoted by $\mathcal{O}(G; \mathcal{G}, k)$ or simply $\mathcal{O}(\mathcal{G}, k)$. If $\Gamma = \sqcup_{i=1}^s \Gamma_i$, given $k_1, \ldots, k_s \in \mathbb{N}$ we set

$$\mathcal{O}(\Gamma, k_1, \ldots, k_s) = \Pi_i \mathcal{O}(\Gamma_i, k_i).$$

We introduce now the group $\mathrm{Aut}(\Gamma)$. The group of automorphisms of G that preserve the set of conjugacy classes of the G_i's is denoted by $\mathrm{Aut}(G;\mathcal{G})$. We set $\mathrm{Out}(G;\mathcal{G}) = \mathrm{Aut}(G;\mathcal{G})/\mathrm{Inn}(G)$

The group $\mathrm{Aut}(G;\mathcal{G})$ acts on $\mathcal{O}(\mathcal{G})$ by changing the marking (i.e. the action), and $\mathrm{Inn}(G)$ acts trivially. Hence $\mathrm{Out}(G;\mathcal{G})$ acts on $\mathcal{O}(\mathcal{G})$. If $X \in \mathcal{O}(\mathcal{G})$ and $[\phi] \in \mathrm{Out}(G;\mathcal{G})$ then ϕX is the same metric tree as X, but the action is $(g,x) \to \phi(g)x$. The action is simplicial and continuous with respect to both simplicial and equivariant Gromov topologies. (See Section 3.2 for details on simplicial structures). We remark that despite the left notation, this is a right-action.

We now extend the definition of $\mathrm{Aut}(G;\mathcal{G})$ to the case of $\Gamma = \sqcup_i \Gamma_i$. We denote by $\mathfrak{S}_k$ the group of permutations of k elements.

Let G and H be two isomorphic groups endowed with splitting $\mathcal{G} : G = G_1 * \dots G_p * F_n$ and $\mathcal{H} : H = H_1 * \dots H_p * F_n$. The set of isomorphisms from G to H that maps each G_i to a conjugate of one of the H_i's is denoted by $\mathrm{Isom}(G,H;\mathcal{G},\mathcal{H})$. If splittings are clear from the context we write simply $\mathrm{Isom}(G,H)$.

Definition 3.10. For $\Gamma = \sqcup_{i=1}^k \Gamma_i$ as in Notation 3.7, we set

$$\mathrm{Aut}(\Gamma) = \{\phi = (\sigma, \phi_1, \dots, \phi_k) : \ \sigma \in \mathfrak{S}_k \text{ and } \phi_i \in \mathrm{Isom}(H_i, H_{\sigma_i})\}.$$

$$\mathrm{Inn}(\Gamma) = \{(\sigma, \phi_1, \dots, \phi_k) \in \mathrm{Aut}(\Gamma) : \sigma = id, \phi_i \in \mathrm{Inn}(H_i)\}$$

$$\mathrm{Out}(\Gamma) = \mathrm{Aut}(\Gamma)/\mathrm{Inn}(\Gamma).$$

The composition of $\mathrm{Aut}(\Gamma)$ is component by component defined as follows. Given $\phi = (\sigma, \phi_1, \dots, \phi_k)$ and $\psi = (\tau, \psi_1, \dots, \psi_k)$ we have

$$\psi\phi = (\tau\sigma, \psi_{\sigma(1)}\phi_1, \dots, \psi_{\sigma(k)}\phi_k).$$

The group $\mathrm{Aut}(\Gamma)$ acts on $\mathcal{O}(\Gamma)$ in the natural way with kernel $\mathrm{Inn}(\Gamma)$, so $\mathrm{Out}(\Gamma)$ acts on $\mathcal{O}(\Gamma)$. More precisely, if $X = (X_1, ..., X_k) \in \mathcal{O}(\Gamma)$ then $X_{\sigma(i)}$ becomes an H_i-tree — denoted $\phi_i X_{\sigma(i)}$ — via the pre-composition of $\phi_i : H_i \to H_{\sigma(i)}$ with the $H_{\sigma(i)}$-action. We set $\phi X = (\phi_1 X_{\sigma(1)}, \dots, \phi_k X_{\sigma(k)})$. (See [13, Section 2] for more details).

3.2. **Simplicial structure of outer spaces and its bordification (**[13, Sections 2.5 and 2.6]**).** The simplicial structure we are going to use is the usual one, that is, (open) simplices are defined as follows: for any $X \in \mathcal{O}(\mathcal{G})$, the set of $\mathcal{G}$-trees obtained from X by varying edge-orbit-lengths in $(0,\infty)$, is an open simplex of $\mathcal{O}(\mathcal{G})$, that we refer to as the open simplex of X, and denote by Δ_X. We notice that $\Delta_X \cap \mathcal{O}_1(\mathcal{G})$ is the standard open simplex in $\mathbb{R}^{\text{number of edge-orbits}}$, while Δ_X is its positive cone (which is topologically still an open simplex, just one dimension bigger). On any open simplex we put the **Euclidean** sup-distance $d_\Delta^{Euclid}(X,Y)$ ($d_\Delta(X,Y)$ or $d(X,Y)$ for short)

$$d_\Delta^{Euclid}(X,Y) = d_\Delta(X,Y) = \max_{e \text{ edge}} |L_X(e) - L_Y(e)|.$$

Such definitions naturally extend to the case of $\Gamma = \sqcup_i \Gamma_i$. (Note, however, that the simplicial structure of $\mathbb{P}\mathcal{O}(\Gamma)$ is not the product of the structures of $\mathbb{P}\mathcal{O}(\pi_1(\Gamma_i))$.)

Remark 3.11. The identification of a simplex Δ with a subset of $\mathbb{R}^m$, induces the notion of linear combination $sX + tY$ for any $X, Y \in \Delta$ and $s,t \geq 0$. In particular, the convex combination $tX + (1-t)Y$ is well defined for any $t \in [0,1]$. We refer to the set $\overline{XY} := \{tX + (1-t)Y, t \in [0,1]\}$ as the Euclidean segment between X and Y.

Simplicial faces of a simplex Δ come in two flavours: finitary faces and faces at infinity.

More precisely, given $\underline{X} \in \mathcal{O}_{\mathrm{gr}}(\Gamma)$, a *finitary face* of Δ_X corresponds to the collapse a forest in $\underline{X}$ whose components have trivial core, so that the resulting graph of groups induces the same splitting of the fundamental group (as graph of groups). We denote the finitary faces just *faces*. We define the *closed* simplex $\overline{\Delta}$ as the closure of Δ in $\mathcal{O}(\Gamma)$, that is:

$$\overline{\Delta} = \Delta \cup \{\text{all faces of } \Delta\} = \Delta \cup \{\text{all finitary faces of } \Delta\}.$$

The *finitary boundary* of X is the set of its proper finitary faces:

$$\partial_{\mathcal{O}}\Delta = \partial_{\mathcal{O}}\overline{\Delta} = \overline{\Delta} \setminus \Delta.$$

A non-finitary simplicial face of an open simplex Δ_X, corresponds to the collapse of sub-graph $\underline{A} \subset \underline{X}$ with at least a component with non-trivial core, and belongs to the outer space $\mathcal{O}(X/A)$, (instead of $\mathcal{O}(X)$). However, if $Y = X/A$, the simplicial topology naturally defines a topology on $\overline{\Delta_X} \cup \overline{\Delta_Y}$, which we still call the simplicial topology. Such a simplicial face will be called a *face at infinity of* $\overline{\Delta_X}$, and if all components of $\underline{A}$ are core-graphs, we call it a *face at infinity of* Δ_X. So, with this terminology, any simplicial face of Δ is either a finitary face of Δ, or a face at infinity of some finitary (not necessarily proper) face of Δ. We refer to [13, Section 2] for a more detailed discussion.

We define the *boundaries at infinity* of a simplex Δ by

$$\partial_\infty\Delta = \{\text{faces at infinity of } \Delta\} \qquad \text{(collapsing of only core sub-graphs)}$$

$$\partial_\infty\overline{\Delta} = \{\text{faces at infinity of } \overline{\Delta}\} \quad \text{(more general collapsing)}$$

and the *closure at infinity by*

$$\overline{\Delta}^\infty = \overline{\Delta} \cup \partial_\infty\overline{\Delta}.$$

If we denote by $\partial\Delta$ the simplicial boundary of Δ, we have

$$\partial\Delta = \partial_\infty\overline{\Delta} \cup \partial_{\mathcal{O}}\overline{\Delta}$$

and

$$\partial_\infty\overline{\Delta} = \bigcup_{F=\text{face of } \Delta} \partial_\infty F$$

(where the union is over all faces of Δ, Δ included.) Moreover, the simplicial closure of Δ is just $\overline{\Delta}^\infty$.

We define the *boundary at infinity* and the *simplicial bordification* of $\mathcal{O}(\Gamma)$ as

$$\partial_\infty\mathcal{O}(\Gamma) = \bigcup_{\Delta \text{ simplex}} \partial_\infty\Delta \qquad \text{and} \qquad \overline{\mathcal{O}(\Gamma)} = \overline{\mathcal{O}(\Gamma)}^\infty = \mathcal{O}(\Gamma) \cup \partial_\infty\mathcal{O}(\Gamma).$$

Remark 3.12. We note that when Γ is just a topological graph with $\pi_1(\Gamma) = F_n$ (all vertex groups are trivial) then $\mathcal{O}(\Gamma)$ is simply the Culler-Vogtmann Outer space CV_n, and the bordification $\overline{\mathcal{O}(\Gamma)}$ is the free splitting complex $\mathcal{FS}_n$. (See [13] for more details).

3.3. **Horoballs and regeneration (**[13, Section 2.7]**).** We keep Notation 3.7.

Definition 3.13 (Horoballs)**.** Given $X \in \partial_\infty\mathcal{O}(\Gamma)$, $\mathrm{Hor}(X)$ is the set of marked metric trees $Y \in \mathcal{O}(\Gamma)$ such that $\underline{X}$ is obtained from $\underline{Y}$ by collapsing a proper family of core sub-graphs. By convention, we set $\mathrm{Hor}(X) = X$ when $X \in \mathcal{O}(\Gamma)$ (and we use $\mathrm{Hor}(\underline{X})$ for graphs). In other words, $Y \in \mathrm{Hor}(X)$ if $\underline{X}$ is obtained by setting to zero the edge-lengths of a proper family of core sub-graphs (note that this implies that Δ_X is a simplicial face of Δ_Y).

$\mathrm{Hor}(X)$ can be regenerated from X as follows[4]

Lemma 3.14. *Suppose $X \in \partial_\infty \mathcal{O}_{gr}(\Gamma)$. Let $Y \in \mathcal{O}_{gr}(\Gamma)$ and $A = \sqcup_i A_i \subset Y$ be a family of core-graphs such that $X = Y/A$. Then, for some k_i, we have*

$$\mathrm{Hor}(\widetilde{X}) = \Pi_i \mathcal{O}(\widetilde{A_i}, k_i).$$

In particular, $\mathrm{Hor}(X) = \mathrm{Hor}(\widetilde{X})$ *is path connected.*

Remark 3.15. Note that we are using the tilde notation here, despite the objects being equivalent, to emphasise that the marked points are points in the trees.

Proof. Let v_i be the non-free vertex of X corresponding to A_i. In order to recover a generic point $Z \in \mathrm{Hor}(X)$, we need to replace each v_i with an element $V_i \in \mathcal{O}_{\mathrm{gr}}(A_i)$. Moreover, in order to completely define the marking on Z, we need to know where to attach - to V_i - the edges of X incident to v_i, and this choice has to be done in the universal covers $\widetilde{V}_i$. No more is needed. Therefore, if k_i denotes the valence of the vertex v_i in X, we have

$$\mathrm{Hor}(\widetilde{X}) = \Pi_i \mathcal{O}(\widetilde{A_i}, k_i).$$

(Note that some k_i could be zero, e.g. if A_i is a connected component of Y.)

Each of the spaces $\mathcal{O}(\widetilde{A_i}, k_i)$ is path connected. Indeed, the map that 'forgets' the marked points is a continuous map to a path connected space whose fibers are connected; since each A_i is connected, we can continuously deform any marked k-tuple of points to another, as we do not insist that they are distinct.

The last statement now follows since a product of path connected spaces is path connected. $\square$

Remark 3.16. With above notation, the forgetting of marked points, gives a well-defined projection $\mathrm{Hor}(X) \to \mathcal{O}(A) = \Pi_i \mathcal{O}(A_i)$. In what follows we will be mainly interested in the composition of such map with the projection $\mathcal{O}(A) \to \mathbb{P}\mathcal{O}(A)$. We therefore give a name to such projection, defining

$$\pi : \mathrm{Hor}(X) \to \mathbb{P}\mathcal{O}(A).$$

(Here $\mathrm{Hor}(X)$ is intended to be not projectivized).

Remark 3.17. Note that *the same* tree X can be considered as a point at infinity of *different* spaces. If we need to specify in which space we work we write $\mathrm{Hor}_\Gamma(X)$.

3.4. **Displacement function, optimal maps and train tracks.** For any $g \in G$ and $X \in \mathcal{O}(\mathcal{G})$, the translation length $L_X(g)$ is defined as $\inf_p = d_X(gp, p)$. Elements with zero translation length correspond to vertex stabilisers, and are called *elliptic*; others have the infimum realised along an axis, and are referred to as *hyperbolic* elements. (Note that an element being elliptic or hyperbolic depends only on $\mathcal{G}$ and not on X). The same happens in $\mathcal{O}(\Gamma)$ componentwise (that is for $g \in \cup_i H_i$, where H_i is as in Notation 3.7). In this section we consider only hyperbolic elements.

Given $X, Y \in O(\Gamma)$, we can compute the translation length of any hyperbolic $g \in \cup_i H_i$ in both X and Y, and we define

$$\Lambda(X, Y) = \sup_g \frac{L_Y(g)}{L_X(g)} = \inf\{\mathrm{Lip}(f) \ : \ f : X \to Y \text{ Lipschitz equivariant map}\}$$

[4]Lemma 3.14, even if implicitly contained and proved in [13, Section 2.7], it is not explicitly stated there. We state and prove it here for future reference.

where $\mathrm{Lip}(f)$ denotes the best Lipschitz constant for f.

It turns out that above second inequality is indeed true, and that sup and inf are in fact max and min (Theorem 4.1([13, Theorem 3.7]), and Theorem 4.2([13, Theorem 3.15])). $\Lambda(X, Y)$ can be computed by means of **straight maps**; that is to say equivariant Lipschitz maps with constant speeds on edges. Given a straight map, the **tension graph** $X_{\max}(f)$ (or simply $X_{\max}$) is the union of edges that are maximally stretched by f. A straight map that realises the above minimum is called **weakly optimal map**, and it is **optimal** if the tension graph has no one-gated vertex (we refer to [13] for further details on gate-structures). An optimal map is **minimal** if the tension graph coincides with the union of the axes of all maximally stretched elements.

Remark 3.18. We could also take the following point of view: given $X, Y \in \overline{\mathcal{O}(\Gamma)}$, let $Hyp(X)$ denote the set of hyperbolic elements in X, and similarly for $Hyp(Y)$. Note that if $X, Y \in \mathcal{O}(\Gamma)$, then $Hyp(X) = Hyp(Y)$. One can then define,

$$\Lambda(X,Y) = \sup_{g \in Hyp(X)} \frac{L_Y(g)}{L_X(g)} = \inf\{\mathrm{Lip}(f) \ : \ f : X \to Y \text{ Lipschitz equivariant map}\}$$

where $\mathrm{Lip}(f)$ denotes the best Lipschitz constant for f, as long as $Hyp(Y) \subseteq Hyp(X)$. That is, as long as elliptic elements of X are also elliptic in Y. If this is not the case, we set $\Lambda(X, Y) = \infty$.

For any automorphism $[\phi] \in \mathrm{Out}(\Gamma)$ we define the displacement function

$$\lambda_\phi : \mathcal{O}(\Gamma) \to \mathbb{R} \qquad \lambda_\phi(X) = \Lambda(X, \phi X)$$

If Δ is a simplex of $\mathcal{O}(\Gamma)$ we define

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

If there is no ambiguity we write simply λ instead of λ_ϕ. Finally, we set

$$\lambda(\phi) = \inf_{X \in \mathcal{O}(\Gamma)} \lambda_\phi(X)$$

In [13] the behaviour of the displacement near points in $\partial_\infty(\mathcal{O}(\Gamma))$ is extensively studied. In particular, it is proven that if $X_\infty \in \partial_\infty(\mathcal{O}(\Gamma))$ is the limit of a sequence of points $X_i \in \mathcal{O}(\Gamma)$ such that $\lambda_\phi(X_i)$ is bounded above, then X_∞ and ϕX_∞ have the same elliptic elements, and ϕ induces an element of $\mathrm{Out}(X_\infty)$. Therefore, for those points, the expression $\lambda_\phi(X_\infty) = \Lambda(X_\infty, \phi X_\infty)$ still makes sense. For other points $T \in \partial_\infty(\mathcal{O}(\Gamma))$ we set $\lambda_\phi(T) = \infty$.

Remark 3.19. Observe that $\Lambda(X_\infty, \phi(X_\infty))$ is finite, according to Remark 3.18, if and only if the set of elliptic elements of X_∞ is ϕ-invariant. If $X_\infty \in \partial_\infty(\mathcal{O}(\Gamma))$ has finite ϕ-displacement, then we can regenerate X_∞ to a point $X \in \mathcal{O}(\Gamma)$ such that X admits an invariant core subgraph, A, which (as a forest) is a union of the axes of the elliptic elements of X_∞ which are hyperbolic in X. X_∞ is obtained from X by collapsing A. Then, by Lemma 4.6, there will be a sequence $X_i \in \mathcal{O}(\Gamma)$ such that $X_i \to X$ and $\lambda_\phi(X_i)$ is bounded above.

Moreover, if $[\phi]$ is irreducible, then *every* $X \in \partial_\infty(\mathcal{O}(\Gamma))$ has infinite ϕ-displacement since no point in $\mathcal{O}(\Gamma)$ admits an invariant core graph.

The displacement function of an automorphism is not continuous at the bordification. Given $[\phi]$, we say that $X \in \overline{\mathcal{O}(\Gamma)}$ has **not jumped** if there is a sequence $X_i \to X$ of points in $\mathcal{O}(\Gamma)$ such that $\lambda_\phi(X_i) \to \lambda_\phi(X)$. Given a simplex Δ with X in the boundary at infinity of Δ, we say that X has **not jumped in** Δ if the above condition holds with $X_i \in \Delta$.

Definition 3.20 ($\mathcal{O}$-maps)**.** For $X, Y \in O(\mathcal{G})$, a map $f : X \to Y$ is called $\mathcal{O}$-map if it is Lipschitz continuous and G-equivariant. Let now $X = (X_1, \dots, X_k)$ and $Y = (Y_1, \dots, Y_k)$ be two elements of $\mathcal{O}(\Gamma)$. A map $f = (f_1, \dots, f_k) : X \to Y$ is called an $\mathcal{O}$-map if for each i the map f_i is an $\mathcal{O}$-map from X_i to Y_i. (No index permutation here).

Let $[\phi] = [(\sigma, \phi_1, \dots, \phi_k)]$ be an element of $\mathrm{Out}(\Gamma)$ - see Definition 3.10.

Definition 3.21 (Maps representing $[\phi]$)**.** Let $X \in \mathcal{O}(\Gamma)$. We say that a map $f : X \to X$ *represents* $[\phi]$[5] if f maps each X_i to $X_{\sigma(i)}$, and such that, if we denote by f_i the map $f|_{X_i} : X_i \to X_{\sigma(i)}$, then f_i is Lipschitz and equivariant with respect to the isomorphism $\phi_i : H_i \to H_{\sigma(i)}$, that is $f_i(hx) = \phi_i(h) f_i(x)$. Note that a map representing $[\phi]$ can be viewed as an $\mathcal{O}$-map $f : X \to \phi X$. We say that f is optimal if each f_i is optimal.

If X is a Γ-graph, then a map $f : X \to X$ represents $[\phi]$ if it has a lift $\widetilde{f} : \widetilde{X} \to \widetilde{X}$ representing $[\phi]$.

Notation 3.22. For notational coherence with [13], if not otherwise specified, if $X, Y \in \mathcal{O}(\Gamma)$ and $f : X \to Y$, when we say that f is straight we understand that it is also an $\mathcal{O}$-map.

In [13, Section 4] we introduced the notion of partial train tracks and partial train tracks at infinity. Roughly: given $[\phi]$, a **partial train track** for $[\phi]$ on $X \in \mathcal{O}(\Gamma)$ is a straight map $f : X \to X$ representing $[\phi]$ such that X has a f-invariant sub-graph to which the restriction of f is a train track; a **partial train track at infinity** is when $X \in \partial_\infty(\mathcal{O}(\Gamma))$.

The deep link between partial train track maps and displacement function, is fully studied and exploited in [13]. In this paper we use results from [13], but we don't need to directly involve partial train tracks. And in fact the words "train track" will appear only in Section 4, where we quote literally statements form [13].

For completeness of exposition we just recall that, as proved in [13], given $[\phi]$, the minimally displaced set of $[\phi]$, that is to say the set of trees T such that $\lambda_\phi(T) = \lambda(\phi)$, coincides with the set of points admitting a partial train track map. For reducible automorphisms, the minimally displaced set may be empty in $\mathcal{O}(\Gamma)$, but if one includes partial train tracks at infinity (partial train tracks for a point at the bordification where the displacement does not jump) then the set of points admitting these partial train tracks is always non-empty and is contained in the minimally displaced set (of points at infinity). Notationally, $\mathrm{Min}(\phi) = \mathrm{TT}(\phi)$ is the minimally displaced set in $\mathcal{O}(\Gamma)$, which coincides with the set of points supporting a partial train track map.

4. Results needed from [13]

In what follows, we will need to quote many lemmas and results from [13]. For the ease of the reader we collected the statements we need from [13] in this section. We decided to quote them exactly as they appear in [13], paying the price that some of them may look a little redundant or overstated here. The reader can safely skip this section now, coming back here when a needed result is cited.

Theorem 4.1 (Sausage Lemma [13, Theorem 3.7])**.** *Let $X, Y \in \mathcal{O}_{gr}(\Gamma)$. The stretching factor $\Lambda(X, Y)$ is realized by a loop $\gamma \subset X$ having has one of the following forms:*

- *Embedded simple loop O;*
- *embedded "infinity"-loop ∞;*
- *embedded barbel O— O;*

[5]In [13] we used f *represents* ϕ. Such notation appears in Section 4 where we quote results from [13].

- *singly degenerate barbel* $\bullet$—O;
- *doubly degenerate barbel* $\bullet$—$\bullet$.

(the $\bullet$ *stands for a non-free vertex.) Such loops are usually named "candidates".*

Theorem 4.2 ([13, Theorem 3.15])**.** *Let $X, Y \in \mathcal{O}(\Gamma)$ and let $f : X \to Y$ be a straight map. There is a map*[6] $\mathrm{weakopt}(f) : X \to Y$ *which is weakly optimal and such that*

$$d_\infty(f, \mathrm{weakopt}(f)) \leq \mathrm{vol}(X)(\lambda(f) - \Lambda(X, Y))$$

Moreover, for any weakly optimal map $\varphi : X \to Y$ and for any $\varepsilon > 0$ there is an optimal map $g : X \to Y$ such that $d_\infty(g, \varphi) < \varepsilon$.

Definition 4.3 (Exit points, [13, Definition 4.19])**.** Let $[\phi] \in \mathrm{Out}(\Gamma)$. A point $X \in \mathcal{O}(\Gamma)$ is called an *exit point* of Δ_X if for any neighbourhood U of X in $\mathcal{O}(\Gamma)$ there is an optimal map $f : X \to X$, representing ϕ, a point $X_E \in U$, and a folding path ([13, Definition 3.21]) directed by f, $X = X_0, X_1, \ldots, X_m = X_E$ in U, such that Δ_{X_i} is a finitary face of $\Delta_{X_{i+1}}$, Δ_X is a proper face of Δ_{X_E}, and such that

$$\lambda_\phi(X_E) < \lambda_\phi(X) \qquad \text{(a strict inequality).}$$

Lemma 4.4 ([13, Lemma 4.20])**.** *Let $[\phi] \in \mathrm{Out}(\Gamma)$ and $X \in \mathcal{O}(\Gamma)$ such that $\lambda_\phi(X)$ is a local minimum for λ_ϕ in Δ_X. Suppose $X \notin \mathrm{TT}(\phi)$.*

Then, for any open neighbourhood U of X in $\mathcal{O}(\Gamma)$, there is an optimal map $f : X \to X$, representing ϕ, points $Z, X' \in U$, and a folding path, $X = X_0, \ldots, X_m = Z, X_{m+1}, \ldots, X_n = X'$, directed by f and such that:

- $X_0, \ldots, X_m \in U \cap \Delta_X$,
- $\lambda_\phi(Z) = \lambda_\phi(X)$,
- Δ_X *is a proper face of* $\Delta_{X'}$,
- $\lambda_\phi(X') < \lambda_\phi(X)$.

In particular X is an exit point of Δ_X.

Theorem 4.5 ([13, Theorem 5.8], lower semicontinuity of λ)**.** *Fix $\phi \in \mathrm{Aut}(\Gamma)$ and $X \in \mathcal{O}_{gr}(\Gamma)$. Let $(X_i)_{i\in\mathbb{N}} \subset \Delta_X$ be a sequence such that there is C such that for any i, $\lambda_\phi(X_i) < C$. Suppose that $X_i \to X_\infty \in \partial_\infty \Delta_X$ which is obtained from X by collapsing a sub-graph $A \subset X$. Then ϕ induces an element of* $\mathrm{Aut}(X/A)$*, still denoted by ϕ.*

Moreover $\lambda_\phi(X_\infty) \leq \liminf_{i\to\infty} \lambda_\phi(X_i)$, and if strict inequality holds, then there is a sequence of minimal optimal maps $f_i : X_i \to X_i$ representing ϕ such that eventually on i we have $(X_i)_{\max} \subseteq \mathrm{core}(A)$[7].

Lemma 4.6 ([13, Lemma 5.12], regeneration of optimal maps)**.** *Fix $\phi \in \mathrm{Aut}(\Gamma)$ and $X \in \mathcal{O}_{gr}(\Gamma)$. Let $X_\infty \in \partial_\infty \Delta_X$ be obtained from X by collapsing a ϕ-invariant core sub-graph A. Then, for any straight map $f_A : A \to A$ representing $\phi|_A$, and for any $\varepsilon > 0$ there is $X_\varepsilon \in \Delta_X$ such that*

$$\lambda_\phi(X_\varepsilon) \leq \max\{\lambda_\phi(X_\infty) + \varepsilon, \mathrm{Lip}(f_A)\}.$$

More precisely, for any $Y \in \mathbb{P}\mathcal{O}_{gr}(A)$ and map $f_Y : Y \to Y$ representing $\phi|_A$, for any map $f : X_\infty \to X_\infty$ representing ϕ, for any $\widehat{X} \in \pi^{-1}(Y)$ [8]*, and for any $\varepsilon > 0$; there is*

[6] We describe an algorithm to find the map $\mathrm{weakopt}(f)$, but the algorithm will depend on some choices, hence the map $\mathrm{weakopt}(f)$ may be not unique in general.

[7] By [13, Proposition 5.6] we know that $\mathrm{core}(A)$ is ϕ-invariant

[8] See Remark 3.16 for an explanation of the map π.

$0 < \delta = \delta(f, f_Y, X_\infty, \Delta_{\widehat{X}})$, *such that for any* $Z \in \Delta_{\widehat{X}} \cap \pi^{-1}(Y)$, *if* $\mathrm{vol}_Z(Y) < \delta$ *there is a straight map* $f_Z : Z \to Z$ *representing* ϕ *such that* $f_Z = f_Y$ *on* Y *and*

$$\mathrm{Lip}(f_Z) \leq \max\{\lambda_\phi(X_\infty) + \varepsilon, \mathrm{Lip}(f_Y)\}$$

(hence the optimal map $\mathrm{opt}(f_Z)$ *satisfies the same inequality*[9]*).*

Theorem 4.7 ([13, Corollary 5.14]). *Let* $\phi \in \mathrm{Aut}(\Gamma)$. *Let* $X \in \mathcal{O}_{gr}(\Gamma)$ *containing an invariant sub-graph* A. *Let* $X_\infty = X/A$ *and* $C = \mathrm{core}(A)$. *Then*

$$\lambda_{\phi|_C}(\Delta_C) \leq \lambda_\phi(\Delta_X).$$

Moreover the following are equivalent:

(1) X_∞ *has not jumped in* Δ_X*;*
(2) $\lambda_\phi(X_\infty) \geq \lambda_\phi(\Delta_X)$*;*
(3) $\lambda_\phi(X_\infty) \geq \lambda_{\phi|_C}(\Delta_C)$.

In particular, $\lambda_\phi(X_\infty)$ *cannot belong to the (potentially empty) interval* $(\lambda_{\phi|_C}(\Delta_C), \lambda_\phi(\Delta_X))$. *Moreover, points realising* $\lambda_\phi(\Delta_X)$ *do not jump in* Δ_X.

Corollary 4.8 ([13, Corollary 5.17]). *Let* $\phi \in \mathrm{Aut}(\Gamma)$. *Let* Δ *be a simplex of* $\mathcal{O}_{gr}(\Gamma)$. *Then there is a min-point* $X_{\min}$ *in* $\overline{\Delta}^\infty$ *(*i.e. a point so that $\lambda_\phi(X_{\min}) = \lambda_\phi(\Delta)$; note that $X_{\min}$ does not jump in Δ by Theorem 4.7).

Moreover, suppose that $X_{\min}$ is *maximal* in the following sense: if $X' \in \overline{\Delta}^\infty$ such that $\lambda_\phi(X') = \lambda_\phi(X_{\min}) = \lambda_\phi(\Delta)$, and $\Delta_{X_{\min}} \subseteq \overline{\Delta_{X'}}^\infty$, then $\Delta_{X_{\min}} = \Delta_{X'}$. ($X_{\min}$ is maximal with respect to the partial order induced by the faces of Δ). Then:

- $\lambda_\phi(X_{\min}) = \lambda_\phi(\Delta_{X_{\min}}) = \lambda_\phi(\Delta)$;
- any point P, such that $\Delta_{X_{\min}} \subseteq \overline{\Delta_P}^\infty \subseteq \overline{\Delta}^\infty$, satisfies $\lambda_\phi(P) \geq \lambda_\phi(\Delta)$ (hence does not jump in Δ by Theorem 4.7);
- for any $\epsilon > 0$, there exist points Z, W such that:
  - $Z \in \Delta$,
  - $\Delta_{X_{\min}} \subseteq \overline{\Delta_W}^\infty \subseteq \overline{\Delta}^\infty$,
  - $\lambda_\phi(W), \lambda_\phi(Z) \leq \lambda_\phi(\Delta) + \epsilon$,
  - λ_ϕ is continuous along the Euclidean segments, ZW and $WX_{\min}$, and any point P along these segments satisfies the following: $\lambda_\phi(\Delta) \leq \lambda_\phi(P)$.

(We allow degeneracies, meaning that X_{min} could equal W, or even Z).

Lemma 4.9 ([13, Lemma 6.2]). *For any* $[\phi] \in \mathrm{Out}(\Gamma)$ *and for any open simplex* Δ *in* $\mathcal{O}(\Gamma)$ *the function* $\lambda = \lambda_\phi$ *is quasi-convex*[10] *on segments of* Δ. *Moreover, for* $A, B \in \Delta$, *if* $\lambda(A) > \lambda(B)$ *then* λ *is strictly monotone near* A[11].

Lemma 4.10 ([13, Lemma 6.3]). *Let* $[\phi] \in \mathrm{Out}(\Gamma)$, *let* $\lambda = \lambda_\phi$, *and let* Δ *be a simplex in* $\mathcal{O}(\Gamma)$. *Let* $A, B \in \overline{\Delta}^\infty$ *be two points that have not jumped in* Δ. *Then for any* $P \in \overline{AB}$

$$\lambda(P) \leq \max\{\lambda(A), \lambda(B)\}$$

Moreover, if $\lambda(A) \geq \lambda(B)$, *then* $\lambda|_{\overline{AB}}$ *is continuous at* A.

Theorem 4.11 ([13, Theorem 7.2]). *For any* Γ *the global simplex-displacement spectrum*

$$\mathrm{spec}(\Gamma) = \Big\{\lambda_\phi(\Delta) : [\phi] \in \mathrm{Out}(\Gamma), \Delta \textit{ a simplex of } \overline{\mathcal{O}(\Gamma)}^\infty \textit{such that } \lambda_\phi(\Delta) < +\infty\Big\}$$

[9] We notice that while $f_Z = f_Y$ on Y, this is no longer true for $\mathrm{opt}(f_Z)$

[10] A function $f : [a, b] \to \mathbb{R}$ is quasi-convex if for any $a \leq x \leq t \leq y \leq b$ we have $f(t) \leq \sup\{f(x), f(y)\}$.

[11] In this statement A, B are points of Δ, and monotonicity is referred to the restriction of λ to the segment joining A, B.

is well-ordered as a subset of $\mathbb{R}$*. In particular, for any* $[\phi] \in \mathrm{Out}(\Gamma)$ *the spectrum of possible minimal displacements*

$$\mathrm{spec}(\phi) = \left\{ \lambda_\phi(\Delta) : \Delta \text{ a simplex of } \overline{\mathcal{O}(\Gamma)}^\infty \text{ such that } \lambda_\phi(\Delta) < +\infty \right\}$$

is well-ordered as a subset of $\mathbb{R}$*.*

Theorem 4.12 ([13, Theorem 7.3])**.** *Let* Γ *be as in Notation 3.7. Let* $[\phi]$ *be any element in* $\mathrm{Out}(\Gamma)$*. Then there exists* $X \in \overline{\mathcal{O}(\Gamma)}^\infty$ *that has not jumped and such that*

$$\lambda_\phi(X) = \lambda(\phi).$$

Lemma 4.13 ([13, Lemma 7.7])**.** *Let* $\phi \in \mathrm{Aut}(\Gamma)$*. Let* $X_\infty \in \overline{\mathcal{O}_{gr}(\Gamma)}$ *which has not jumped. Suppose that there is a loop* $\gamma \in X_\infty$ *and* $k > 0$ *and such that* $L_{X_\infty}(\phi^n(\gamma)) \geq k^n L_{X_\infty}(\gamma)$ *for all* $n \in \mathbb{N}$*. Then,*

$$k \leq \lambda(\phi).$$

In particular, if X_∞ *is a train track for* ϕ *as an element of* $\mathrm{Aut}(X_\infty)$*, then it is a minpoint for* ϕ *as an element of* $\mathrm{Aut}(\Gamma)$*.*

Theorem 4.14 ([13, Theorem 7.8])**.** *Let* $\phi \in \mathrm{Aut}(\Gamma)$*. Let* $X \in \mathcal{O}(\Gamma)$ *and* X_∞ *be such that* $\underline{X_\infty}$ *is obtained from* $\underline{X}$ *by collapsing a* ϕ*-invariant core sub-graph* $\underline{A}$*. Then*

$$\lambda(\phi|_A) \leq \lambda(\phi).$$

Moreover, if $\lambda(\phi|_A) = \lambda_\phi(X_\infty)$*, then*

$$\lambda(\phi) = \lambda(\phi|_A).$$

In particular X_∞ *has not jumped if and only if*

$$\lambda(\phi) \leq \lambda(X_\infty).$$

Remark 4.15. We note that if a point has not jumped, this simply means that there is some sequence converging to it, whose displacements tend to the displacement of that point. In general this will not hold for all sequences tending to the point.

Theorem 4.16 ([13, Theorem 7.13])**.** *Let* $\phi \in \mathrm{Aut}(\Gamma)$*. Let* X *be a* Γ*-graph having a* ϕ*-invariant core sub-graph* A*. Then there is* $Z \in \overline{\mathcal{O}(X/A)}^\infty$ *and* $W \in \mathrm{Hor}_{\mathcal{O}(\Gamma)}(Z)$ *such that the simplex* Δ_W *contains a minimising sequence for* λ*. Moreover if* Y *is the graph used to regenerate* W *from* Z*, then the minimising sequence can be chosen with straight maps* f_i *such that* $f_i(Y) = Y$ *and* $\mathrm{Lip}(f_i) \to \lambda(\phi)$*.*

5. Statement of the connectedness theorem and regeneration of paths in the bordification

We recall here Notation 3.7 (as a courtesy for readers who skipped the first sections).

- $\Gamma = \sqcup \Gamma_i$ will always mean that Γ is a finite disjoint union of finite graphs of groups Γ_i, each with trivial edge-groups and non-trivial fundamental group $H_i = \pi_i(\Gamma_i)$, each H_i being equipped with the splitting given by the vertex-groups. We set $\mathrm{rank}(\Gamma) = \sum_i \mathrm{rank}(\Gamma_i)$.

We also recall that for any $[\phi] \in \mathrm{Out}(\Gamma)$ we defined the displacement function

$$\lambda_\phi : \mathcal{O}(\Gamma) \to \mathbb{R} \qquad \lambda_\phi(X) = \Lambda(X, \phi X)$$

If Δ is a simplex of $\mathcal{O}(\Gamma)$ we define

$$\lambda_\phi(\Delta) = \inf_{X \in \Delta} \lambda_\phi(X).$$

If there is no ambiguity we write simply λ instead of λ_ϕ. Finally, we set

$$\lambda(\phi) = \inf_{X \in \mathcal{O}(\Gamma)} \lambda_\phi(X)$$

By convention (see Section 3.4) we extend the function λ to points in $X_\infty \in \partial_\infty(\mathcal{O}(\Gamma))$ for which there is a sequence of points $X_i \in \mathcal{O}(\Gamma)$ such that $X_i \to X_\infty$ with $\lambda(X_i)$ bounded above, and we set $\lambda = \infty$ on other points.

Finally, we recall that outer space comes in two flavours: trees and graphs. We will chose which one we use on a case-by-case basis, depending on which is more convenient. For that purpose we introduced the notation "$\mathcal{O}(\Gamma)$" for trees and "$\mathcal{O}_{\mathrm{gr}}(\Gamma)$" for graphs. Clearly $\mathcal{O}_{\mathrm{gr}}(\Gamma)$ and $\mathcal{O}(\Gamma)$ are isomorphic via $X \leftrightarrow \widetilde{X}$, and thus in all statements they are completely interchangeable.

Definition 5.1. Let $X, Y \in \overline{\mathcal{O}(\Gamma)}^\infty$. A *simplicial path* Σ between X, Y is given by:

(1) A finite sequence of points $X = X_0, X_1, \dots, X_k = Y$, called vertices, such that $\forall i = 1, \dots, k$, there is a simplex Δ_i such that $\Delta_{X_{i-1}}$ and Δ_{X_i} are both simplicial[12] faces of Δ_i. We allow one of $\Delta_{X_{i-1}}, \Delta_{X_i}$, or even both, to coincide with Δ_i.
(2) Euclidean segments $\overline{X_{i-1}X_i} \subseteq \overline{\Delta_i}^\infty$, called edges. We require the interior of $\overline{X_{i-1}X_i}$ (i.e. $\overline{X_{i-1}X_i} \setminus \{X_{i-1}, X_i\}$) to be contained in Δ_i.
(3) We say that Σ is *alternating* if for every i either Δ_{X_i} is a simplicial face of $\Delta_{X_{i-1}}$ or $\Delta_{X_{i-1}}$ is a simplicial face of Δ_{X_i}. Note that any simplicial path can be made alternating just by adding some extra vertex.

Definition 5.2. We say that a set χ is *connected by simplicial paths* if for any $x, y \in \chi$ there is a simplicial path between x and y which is entirely contained in χ.

Theorem 5.3 (Level sets are connected). *Let* $[\phi] \in \mathrm{Out}(\Gamma)$. *For any* $\varepsilon > 0$ *the set*

$$\{X \in \mathcal{O}(\Gamma) : \lambda_\phi(X) \leq \lambda(\phi) + \varepsilon\}$$

is connected in $\mathcal{O}(\Gamma)$ *by simplicial paths. The set*

$$\{X \in \overline{\mathcal{O}(\Gamma)}^\infty : \lambda_\phi(X) = \lambda(\phi)\}$$

is connected by simplicial paths in $\overline{\mathcal{O}(\Gamma)}^\infty$.

Moreover, connecting paths can be chosen so that the displacement λ_ϕ *is continuous along them.*

The main goal of the paper is the proof of Theorem 5.3. The rough strategy is to prove that paths in the bordification can regenerate to paths in $\mathcal{O}(\Gamma)$ without increasing λ too much. Then, the first claim will follow from the second, which we will prove via a peak-reduction argument. Proofs proceed via induction on the rank of Γ. This is part of the reason that we need to fundamentally deal with the case where Γ is disconnected. We remind that Theorem 5.3, if specialised to the case where Γ is connected and vertex-groups are trivial, is a CV_n-statement about connectedness of level sets of, not necessarily irreducible, automorphisms of the free group F_n.

Corollary 5.4. *Let* $[\phi] \in \mathrm{Out}(\Gamma)$ *be irreducible. Then the set*

$$\{X \in \mathcal{O}(\Gamma) : \lambda_\phi(X) = \lambda(\phi)\}$$

is connected in $\mathcal{O}(\Gamma)$ *by simplicial paths.*

[12] We remind that simplicial faces include faces at infinity. That is to say, $\Delta_{X_{i-1}}$ and Δ_{X_i} are both faces of $\overline{\Delta_i}^\infty$.

Proof. This is an immediate consequence of Theorem 5.3, since by Remark 3.19, the irreducibility of $[\phi]$ implies that every point on the boundary, $\partial_\infty(\mathcal{O}(\Gamma))$, has infinite displacement. $\square$

Remark 5.5. Theorem 5.3 is trivially true if $\operatorname{rank}(\Gamma) = 1$, because in that case either $\mathcal{O}(\Gamma)$ or $\mathbb{P}\mathcal{O}(\Gamma)$ is a single point.

Lemma 5.6 (Regeneration of segments)**.** *Fix $[\phi] \in \operatorname{Out}(\Gamma)$. Let $X_\infty, Y_\infty \in \overline{\mathcal{O}(\Gamma)}^\infty$ such that Δ_{Y_∞} is a (not necessarily proper) simplicial face of Δ_{X_∞}. Suppose that $\lambda(X_\infty) \geq \lambda(\phi)$. Then there is an open simplex Δ of $\mathcal{O}(\Gamma)$ such that for any $\varepsilon > 0$ there is $Y \in \operatorname{Hor}(Y_\infty) \cap \overline{\Delta}$ and $X \in \operatorname{Hor}(X_\infty) \cap \Delta$ such that*[13]

$$\lambda_\phi(Y), \lambda_\phi(X) < \max\{\lambda_\phi(Y_\infty), \lambda_\phi(X_\infty)\} + \varepsilon.$$

Moreover, such an inequality holds for any $T \in \overline{XX_\infty}$ and any $S \in \overline{YY_\infty}$.

Proof. For this proof will be more convenient to work in $\mathcal{O}_{\mathrm{gr}}$ rather than $\mathcal{O}$. Let X_∞ be obtained by collapsing a ϕ-invariant core-subgraph A from a Γ-graph $\widehat{X}$. Since $\lambda_\phi(X_\infty) \geq \lambda(\phi)$, by Theorem 4.14 $\lambda(\phi|_A) \leq \lambda_\phi(X_\infty)$. By Theorems 4.12 and 4.16, there is a simplex in $\mathcal{O}_{\mathrm{gr}}(A)$ that contains a minimising sequence for $\lambda(\phi|_A)$. Let A_ε be a point in that simplex such that $\lambda(A_\varepsilon) < \lambda(\phi|_A) + \varepsilon$. The required simplex Δ is obtained by inserting a copy of A_ε in place of A in X_∞. We note that such a Δ is not unique. By Lemma 4.6 there is a point $X \in \Delta \cap \operatorname{Hor}(X_\infty)$ such that $\lambda_\phi(X) \leq \lambda_\phi(X_\infty) + \varepsilon$.

Consider now the points in $\overline{\Delta} \cap \operatorname{Hor}(Y_\infty)$. By hypothesis there is a ϕ-invariant $B \subseteq X_\infty$ such that as a graph (that is, forgetting the metric), Y_∞ is obtained from X_∞ by collapsing B. B has a pre-image in X still denoted by B. Let T be the forest $(A \cup B) \setminus \operatorname{core}(A \cup B)$. If $Y' = X/T$, as a graph, $Y_\infty = X/(A \cup B) = Y'/\operatorname{core}(A \cup B)$.

Thus the finitary face $\Delta_{Y'}$ of Δ obtained by the collapse of T intersects $\operatorname{Hor}(Y_\infty)$.

Let $f : X \to X$ be an optimal map representing $[\phi]$. Since $\operatorname{core}(A \cup B)$ is ϕ-invariant, $f(\operatorname{core}(A \cup B)) \subset \operatorname{core}(A \cup B)$ up to homotopy. It follows that there is a straight map $g : \operatorname{core}(A \cup B) \to \operatorname{core}(A \cup B)$ representing $[\phi|_{A\cup B}]$ such that $\operatorname{Lip}(g) \leq \lambda_\phi(X) \leq \lambda_\phi(X_\infty) + \varepsilon$. By Lemma 4.6 there is a point $Y \in \operatorname{Hor}(Y_\infty) \cap \Delta_{Y'}$ such that $\lambda_\phi(Y) \leq \max\{\lambda_\phi(Y_\infty) + \varepsilon, \operatorname{Lip}(g)\} \leq \max\{\lambda_\phi(Y_\infty) + \varepsilon, \lambda_\phi(X_\infty) + \varepsilon\}$. The last claim also follows by Lemma 4.6, since the volume of A (or B) is strictly decreasing on the Euclidean segment $\overline{XX_\infty}$ (or $\overline{YY_\infty}$), and the invariant subgraph is being scaled uniformly. $\square$

Now we can plug in the inductive hypothesis in the proof of Theorem 5.3.

Lemma 5.7 (Regeneration of horoballs)**.** *Suppose that Theorem 5.3 is true in any rank less than* $\operatorname{rank}(\Gamma)$. *Let $[\phi] \in \operatorname{Out}(\Gamma)$. Let $T \in \mathcal{O}_{gr}(\Gamma)$ be a Γ-graph having a proper ϕ-invariant core sub-graph S. Let $X \in \partial_\infty \mathcal{O}_{gr}(\Gamma)$ be the graph obtained from T by collapsing S, and let $A, B \in \operatorname{Hor}(X) \subset \mathcal{O}_{gr}(\Gamma)$. Let m_A and m_B be the supremum of λ_ϕ on the Euclidean segments $\overline{AX}$ and $\overline{BX}$ respectively. Then, for any $\varepsilon > 0$ there is a simplicial path Σ between A and B, and in* $\operatorname{Hor}(X)$*, such that for any vertex Z of Σ we have*

$$\lambda_\phi(Z) < \max\{m_A, m_B\} + \varepsilon.$$

Proof. Let $L = \max\{m_A, m_B\}$. The displacement $\lambda_\phi(T)$ is a finite number just because λ_ϕ is a well-defined function on $\mathcal{O}_{\mathrm{gr}}(\Gamma)$. For any group element $g \in \cup_i H_i$, and for any $t \in [0, 1)$, the translation length of g in $T_t = tX + (1-t)T$ is bounded by $L_T(g)$. Moreover, as T is a finite graph, there is $\delta > 0$ such that for any reduced loop γ in T, either $\gamma \subseteq S$

[13]Note that the fact that $\operatorname{Hor}(Y_\infty) \cap \overline{\Delta} \neq \emptyset$ implies that Δ_{Y_∞} is a simplicial face of Δ. The same holds true for X_∞.

or the length of γ in any T_t is at least δ. Therefore, since S is ϕ-invariant, and by using Theorem 4.1, we see that $\lambda_\phi(T_t)$ is bounded by some constant C, uniform on t.

By Theorem 4.5 we have that $\lambda_\phi(X)$ is finite, and by Lemma 4.6 both m_A and m_B are finite[14].

Recall that if $X = T/S$ as graphs of groups, then we denote by $\pi : \mathrm{Hor}(X) \to \mathbb{P}\mathcal{O}_{\mathrm{gr}}(S)$ the projection that associates to a point in $\mathrm{Hor}(X)$ its collapsed part (see Section 3.3).

For any $Y \in \mathrm{Hor}(X)$, $\lambda_\phi(Y)$ can be computed by tacking the supremum of stretching factors of candidates given by Theorem 4.1. Those may or may not be contained in S, and clearly the supremum over all candidates is bigger or equal to the supremum over candidates contained in S. Since S is ϕ-invariant, this implies that $\lambda_\phi(\pi(Y)) \le \lambda_\phi(Y)$ for any $Y \in \mathrm{Hor}(X)$; so

$$\lambda_\phi(\pi(A)) \le \lambda_\phi(A) \quad \lambda_\phi(\pi(B)) \le \lambda_\phi(B)$$

hence, $\lambda_\phi(\pi(A)), \lambda_\phi(\pi(B)) \le L$. The rank of S is strictly smaller than $\mathrm{rank}(\Gamma)$ because it is a proper sub-graph of T. Hence Theorem 5.3 holds for $\mathcal{O}_{\mathrm{gr}}(S)$. Therefore, the induction hypothesis produces a finite simplicial path $(Y_i) \in \mathcal{O}_{\mathrm{gr}}(S)$ between $\pi(A)$ and $\pi(B)$ such that $\lambda_\phi(Y_i) < L + \varepsilon$. Hence, by Lemma 3.14, there is a finite simplicial path in $\mathrm{Hor}(X)$ between A and B whose vertices are points $\widehat{T}_j$ such that for any j there is i such that $\pi(\widehat{T}_j) = Y_i$. By Lemma 4.6 there is a simplicial path in $\mathrm{Hor}(X)$ whose vertices are points $Z_j \in \Delta_{\widehat{T}_j}$ such that $\pi(Z_j) = \pi(\widehat{T}_j) = Y_i$ and $\lambda_\phi(Z_j) < L + \varepsilon$. □

We recall that we are using the notation of Definition 5.1.

Theorem 5.8 (Regeneration of paths)**.** *Suppose that Theorem 5.3 is true in any rank less than* $\mathrm{rank}(\Gamma)$. *Let* $[\phi] \in \mathrm{Out}(\Gamma)$. *Let* $\Sigma = (X_i)_{i=1}^m$ *be an alternating simplicial path in* $\overline{\mathcal{O}(\Gamma)}^\infty$, *and let* L *be a positive real number.*

Suppose that for any point X_i *we have*

$$\lambda(\phi) \le \lambda_\phi(X_i) \le L.$$

Then, for any $\varepsilon > 0$ *there exists a simplicial path* $\Sigma_\varepsilon = (W_j)_{j=1}^k$ *in* $\mathcal{O}(\Gamma)$, *such that for any point* P *of* Σ_ε, $\lambda(P) \le (L + \varepsilon)$.

Moreover, we can choose the path so that $W_1 \in \mathrm{Hor}(X_1)$, $W_k \in \mathrm{Hor}(X_m)$, *each* W_j *belongs to the horoball of some* X_i; *and so that* X_1 *and* X_m *do not jump in* Δ_{W_1} *and* Δ_{W_k} *respectively.*

Proof. By Lemmas 4.9 and 4.10, and since the displacement is continuous in $\mathcal{O}(\Gamma)$, it suffices to check displacement on vertices of Σ_ε.

For any $i < m$, we apply Lemma 5.6 to the i^{th} pair of consecutive points X_i, X_{i+1}. This produces points $A_i \in \mathrm{Hor}(X_i)$ and $B_{i+1} \in \mathrm{Hor}(X_{i+1})$ whose displacement is less than $L + \varepsilon$. Note that ε is arbitrary. In particular Theorem 4.7 implies that X_1 does not jump in Δ_{A_1} and X_m does not jump in Δ_{B_m}. Moreover, A_i, B_{i+1} are in the same closed simplex of $\mathcal{O}(\Gamma)$ (so there is a Euclidean segment joining them).

Additionally Lemma 5.6 tells us that the displacement of points along the segments, $\overline{A_i X_i}, \overline{B_i X_i}$ is bounded by $L + \varepsilon$.

Note that A_i is defined for $1 \le i < m$ and B_i for $1 < i \le m$. By Lemma 5.7, for $1 < i < m$, there is a simplicial path Y_{ij} between B_i and A_i such that $Y_{ij} \in \mathrm{Hor}(X_i)$ and $\lambda_\phi(Y_{ij}) \le L + \varepsilon$. The path Σ_ε is now defined by the concatenation of such paths and the segments $\overline{A_i B_{i+1}}$. □

[14]One has to apply Lemma 4.6 as follows: X here plays the role of X_∞ of lemma; T here is X in lemma, S here is A in lemma; A, B here play the role of $\widehat{X}$ in lemma (for suitable Y).

6. Calibration of paths

We keep Notation 3.7. For the remaining of the section we fix $[\phi] \in \mathrm{Out}(\Gamma)$. We recall that for simplices $\Delta \subset \overline{\mathcal{O}(\Gamma)}^{\infty}$ we are using the notation $\lambda(\Delta) = \lambda_\phi(\Delta) = \inf_{X \in \Delta} \lambda_\phi(X)$.

Our aim is to run a peak reduction argument to prove Theorem 5.3, by starting with a simplicial path and locally modifying it near peaks. Theorem 5.8 provides simplicial paths with bounded displacement, however, for our purposes we need paths, that possibly touch the boundary at infinity, where the displacement is continuous. (The displacement is not in general continuous on $\overline{\mathcal{O}(\Gamma)}^{\infty}$.)

In this section we describe a procedure for *calibrating* simplicial paths (see below precise definitions).

Definition 6.1. Let Σ be a (simplicial) path in $\overline{\mathcal{O}(\Gamma)}^{\infty}$. We set $\lambda(\Sigma)$, the *displacement* of Σ, to be the supremum of displacements of points along Σ.

Definition 6.2. Let L be a positive real number. A simplicial path $\Sigma = (X_i)_{i=0}^k$ in $\overline{\mathcal{O}(\Gamma)}^{\infty}$ is said to be *L-calibrated* if:

(i) λ is continuous on Σ;
(ii) $\lambda(\Sigma) \leq L$;
(iii) no point P of Σ jumps (which, by Theorem 4.14, is equivalent to $\lambda(\phi) \leq \lambda(P)$);
(iv) for any point P, in the interior of Σ and that realises the maximum $\lambda(\Sigma)$, we have $\lambda(P) = \lambda(\Delta_P)$ (*i.e.* P is minimising in its simplex). Note that this implies that $\lambda(\Sigma) \in \mathrm{spec}(\phi) \cup \{\lambda(X_0), \lambda(X_k)\}$ (see Theorem 4.11 for definition of $\mathrm{spec}(\phi)$).

Remark 6.3. If A, B are two consecutive vertices of an L-calibrated path then, by the continuity of λ, neither point can have jumped in the simplex they span. Hence by Lemma 4.10 and property (ii) of Definition 6.2, for any P in the segment $\overline{AB}$ we have,

$$\lambda(\phi) \leq \lambda(P) \leq \max\{\lambda(A), \lambda(B)\} \leq L.$$

Theorem 6.4 (Calibration)**.** *Suppose Theorem 5.3 is true in any rank less than* $\mathrm{rank}(\Gamma)$. *Let Σ be a simplicial path in $\overline{\mathcal{O}(\Gamma)}^{\infty}$ with finite displacement and such that no point of Σ jumps. Then in $\overline{\mathcal{O}(\Gamma)}^{\infty}$ there exists a $\lambda(\Sigma)$-calibrated simplicial path Σ_o with the same endpoints as Σ.*

Proof. We outline the strategy of this proof to aid the reader.

- First we regenerate Σ to a path Σ_1 which lives inside $\mathcal{O}(\Gamma)$. This is basically an application Theorem 5.8 in its full generality, which, in particular, requires the inductive hypothesis on the rank. This is the only place of this section where such hypothesis is needed. We also note that in case ϕ is irreducible, then any path with finite displacement is in $\mathcal{O}(\Gamma)$ (see Remark 3.19) so this step (and hence inductive hypothesis) is not necessary in this case.
- Next, we define a simplicial path Σ_2 in $\overline{\mathcal{O}(\Gamma)}$, obtained from Σ_1 by, essentially, replacing each vertex with a point that minimizes the displacement in the corresponding simplex.
- Finally, we add extra points to Σ_2 in order to obtain a simplicial path Σ_o to ensure that λ is continuous along the path.
- Along the way, we verify that we maintain control of the displacements of our paths, exploiting both quasi-convexity and the fact that $\mathrm{spec}(\phi)$ is well ordered (Theorem 4.11).

Let $\Sigma = (X_i)_{i=1}^m$. Up to possibly adding extra vertices belonging to segments of Σ, we may assume that it is alternating. (Note that this does not change the displacement of Σ).

Let $M = \min\{x \in \mathrm{spec}(\phi) : x > \lambda(\Sigma)\}$, which exists because $\lambda(\Sigma) < +\infty$ and $\mathrm{spec}(\phi)$ is well ordered (Theorem 4.11). Let $\varepsilon > 0$ so that $\lambda(\Sigma) + \varepsilon < M$.

We start by invoking Theorem 5.8 to produce a simplicial path $\Sigma_1 = (W_j)_{j=1}^k$ in $\mathcal{O}(\Gamma)$, so that $\lambda(\Sigma_1) \leq \lambda(\Sigma) + \varepsilon < M$ and so that W_1 and W_k do not jump in in Δ_{X_1} and Δ_{X_m} respectively. (Note that Δ_{X_1} is a face of Δ_{W_1}, and Δ_{X_m} is of Δ_{W_k}).

We define a new simplicial path, Σ_2, as follows:

(1) For any j, if $\Delta_{W_{j-1}}$ and Δ_{W_j} are both proper faces of some Δ_j, then we add to the path a new point, $\widehat{W_j} \in \Delta_j \cap \overline{W_{j-1}W_j}$. We note that $\lambda(\widehat{W_j}) \leq \lambda(W_{j-1})$ and $\lambda(W_j)$, by quasi-convexity (Lemma 4.9).

(2) We renumber the sequence of vertices, denoting them by $(W_j)_{j=1}^l$ (for some $l \geq k$). We now have a simplicial path which is alternating.

(3) For any $1 \leq j \leq l$, we use Corollary 4.8 and replace W_j by a point $Y_j \in \overline{\Delta_{W_j}}^\infty$, chosen so that $\lambda(Y_j) = \lambda(\Delta_{W_j}) = \lambda(\Delta_{Y_j})$, and requiring Y_j to be maximal in the sense of Corollary 4.8.

(4) We add endpoints $Y_0 = X_1$ and $Y_{l+1} = X_m$.

(5) If two consecutive points coincide, then we identify them and we renumber the sequence accordingly (and removing the corresponding segment). We call the resulting alternating simplicial path Σ_2.

Lemma 6.5. *For any vertex, $Y_j \in \Sigma_2$, we have that $\lambda(\Sigma) \geq \lambda(Y_j) \in \mathrm{spec}(\phi) \cup \{\lambda(X_1), \lambda(X_m)\}$.*

Proof. The statement is obvious for endpoints. For other points, by construction, we have $M > \lambda(\Sigma) + \varepsilon \geq \lambda(W_j) \geq \lambda(Y_j) \in \mathrm{spec}(\phi)$, and our choice of M implies $\lambda(Y_j) \leq \lambda(\Sigma)$. $\square$

Remark 6.6. In Step (4) we have $\lambda(Y_0) \geq \lambda(Y_1)$ and $\lambda(Y_{l+1}) \geq \lambda(Y_l)$. This is because by definition Y_1 is the point in $\overline{\Delta_{X_1}}^\infty$ that realises $\lambda(\Delta_{X_1}) = \inf_{T \in \Delta_{X_1}} \lambda(T)$, and $Y_0 = X_1$. The same argument works for Y_l.

Lemma 6.7. *Let A, B be two consecutive vertices of Σ_2. Then,*

(a) For any point P of $\overline{AB}$ we have $\lambda(P) \geq \lambda(\phi)$.

(b) if $\lambda(A) = \lambda(B)$, then λ is constant on the segment $\overline{AB}$;

(c) if $\lambda(A) > \lambda(B)$, then there exists a simplex $\Delta \subset \mathcal{O}(\Gamma)$ and points C, D so that:

- *$A, B, C, D \in \overline{\Delta}^\infty$;*
- *$\lambda(A) < \lambda(C), \lambda(D) < \lambda(B)$;*
- *λ is continuous on Euclidean segments $\overline{AC}$, $\overline{CD}$, and $\overline{DB}$.*

Proof. Either $\lambda(A) \geq \lambda(B)$ or vice versa. Without loss of generality, up to possibly switch the names of A, B, we may assume that $\lambda(A) \geq \lambda(B)$.

By how Σ_2 is defined, A, B are introduced either in Step (3) or in Step (4). Suppose first that both come from Step (3). Since they are consecutive in Σ_2, by Step (5) we may assume that the pair $\{A, B\}$ comes, in Step (3), from a pair $\{W_j, W_{j+1}\}$ of two consecutive vertices of the Step (2)-path. Since the path of Step (2) is alternating, and contained in $\mathcal{O}(\Gamma)$, either Δ_{W_j} is a finitary face of $\Delta_{W_{j+1}}$ or vice versa: let Δ_0 be the one which is face of the other, and let Δ be the other. (We may have $\Delta = \Delta_0$.) Moreover, from Step (3), either $\lambda(A) = \lambda(\Delta_0)$ and $\lambda(B) = \lambda(\Delta)$, or vice versa.

Since Δ_0 is a face of Δ we have $\lambda(\Delta_0) = \inf_{T \in \Delta_0} \lambda(T) \geq \inf_{T \in \Delta} \lambda(T) = \lambda(\Delta)$. Since we are assuming $\lambda(A) \geq \lambda(B)$, w.l.o.g. we may also assume $A \in \overline{\Delta_0}^\infty$ and $B \in \overline{\Delta}^\infty$, thus $\lambda(A) = \lambda(\Delta_0), \lambda(B) = \lambda(\Delta)$.

Suppose now that A is introduced in Step (4). Then it is an endpoint, say $A = Y_0 = X_1$. In this case necessarily $B = Y_1$ is obtained in Step (3) from $X_1 = A$. Whence $\lambda(A) > \lambda(\Delta_A)$ and $\lambda(B) = \lambda(\Delta_A)$. In this case we set $\Delta_0 = \Delta = \Delta_A$.

The same reasoning would work if B is introduced in Step (4), but then we would get $\lambda(B) > \lambda(A)$ contradicting $\lambda(A) \geq \lambda(B)$. Therefore this latter situation cannot happen. In particular B is always introduced in Step (3).

In any case, there there exists an open simplex Δ in $\mathcal{O}(\Gamma)$, with a (not necessarily proper) finitary face, Δ_0, such that $A \in \overline{\Delta_0}^{\infty}$, $B \in \overline{\Delta}^{\infty}$ and so that $\lambda(\Delta_0) \leq \lambda(A)$, $\lambda(\Delta) = \lambda(B)$. Thus both A and B belong to $\overline{\Delta}^{\infty}$.

Now let Δ_1 be the simplicial face of Δ spanned by A and B (which may be different from Δ). Both $\lambda(A), \lambda(B)$ are finite. So, topologically, A and B are obtained from a graph, X, by collapsing invariant subgraphs C_A and C_B, respectively. Therefore the points in Δ_1 are obtained from X by collapsing $C_A \cap C_B$, which is also invariant and hence all points in Δ_1 have finite displacement.

By the maximality of the dimension of Δ_B (Step (3)), and Theorem 4.8, no point in Δ_1 has jumped in Δ. Hence, by Theorem 4.7 and Lemma 4.10, for any point P, on the segment from A to B,

$$\lambda(\phi) \leq \lambda(\Delta) = \lambda(B) \leq \lambda(P) \leq \max\{\lambda(A), \lambda(B)\}.$$

This in particular proves (a). Moreover, if $\lambda(A) = \lambda(B)$, we deduce that the previous inequalities - except the first - are all equalities, thus proving (b).

Finally, suppose that $\lambda(A) > \lambda(B)$. Since λ is continuous in Δ_1, and since A has not jumped in Δ_1 by Theorem 4.7, we deduce - by Lemma 4.10 - that λ is continuous along the segment from A to B except, possibly, at B.

If λ is continuous in $\overline{AB}$, there is nothing to prove. Otherwise, we use the fact that B is defined in Step (3) by applying Corollary 4.8. Our points C, D correspond then to points Z, W of Corollary 4.8, which can be chosen with displacement arbitrarily close to $\lambda(\Delta) = \lambda(B)$, in particular so that $\lambda(A) > \lambda(C), \lambda(D)$. The fact that $\lambda(C), \lambda(D) > \lambda(B) = \lambda(\Delta)$ follows from maximality condition of B (Step (3)). Corollary 4.8 also provides the continuity of λ on the segments $\overline{CD}$ and $\overline{DB}$. The continuity of $\overline{AC}$ follows from Lemma 4.10 because A has higher displacement.

□

We are now in position to finish the proof of Theorem 6.4. Having Σ_2, we build Σ_o by using Lemma 6.7 to add points C, D between consecutive vertices where λ is not continuous. In particular, λ is continuous on Σ_o, and condition (i) of Definition 6.2 is satisfied. Point (a) of Lemma 6.7 gives condition (iii).

Note that added vertices are never point of maximum. Therefore Lemma 6.5 provides condition (iv). Finally Lemma 6.5 and Lemma 4.10 imply that $\lambda(\Sigma_o) \leq \lambda(\Sigma)$, so also condition (ii) of Definition 6.2 is fulfilled with $L = \lambda(\Sigma)$. Thus Σ_o is $\lambda(\Sigma)$-calibrated. □

7. Preparation to peak reduction

We keep Notation 3.7. For the remaining of the section we fix $[\phi] \in \mathrm{Out}(\Gamma)$. We recall that for simplices $\Delta \in \overline{\mathcal{O}(\Gamma)}^{\infty}$ we are using the notation $\lambda(\Delta) = \lambda_\phi(\Delta) = \inf_{X \in \Delta} \lambda_\phi(X)$. In this section we prove some preliminary result needed to perform reduction of peaks.

We start by stating a (technical) fact that can be informally phrased as follows[15]:

[15]We recall that by definition $\overline{\mathcal{O}(\Gamma)}^{\infty} = \overline{\mathcal{O}(\Gamma)}$ and that the symbol ∞ is just to put emphasis on the fact that we are considering the simplicial bordification of the outer space obtained by adding all simplices at infinity.

Given $X \in \overline{\mathcal{O}(\Gamma)}^{\infty}$ and $f : X \to X$ an optimal map representing $[\phi]$, if Y is sufficiently close to X for the Euclidean metric, then any fold in X directed by f can be closely read in Y.

Theorem 7.1. *Let $X, Y \in \overline{\mathcal{O}_{gr}(\Gamma)}$. Suppose that Δ_X is a simplicial face of Δ_Y. Thus as graphs, X is obtained from Y by collapsing a sub-graph A. Suppose that* core(A) *is ϕ-invariant. For $t \in [0,1]$ let $Y_t = (1-t)X + tY$ be a parametrization of the Euclidean segment from X to Y. Let $\sigma_t : Y_t \to X$ be the map obtained by collapsing A and by linearly rescaling the edges in $Y \setminus A$.*

Let $f : X \to X$ be an optimal map representing $[\phi]$. Then for any $\varepsilon > 0$ there is $t_\varepsilon > 0$ such that $\forall 0 \le t < t_\varepsilon$ there is an optimal map $g_t : Y_t \to Y_t$ representing $[\phi]$ such that

$$d_\infty(\sigma_t \circ g_t, f \circ \sigma_t) < \varepsilon.$$

Proof. The proof of this theorem relies on accurate (but boring) estimates. For the happiness of the reader we postpone the proof to the appendix. □

Remark 7.2. Note that when $Y \in \overline{\mathcal{O}(\Gamma)}$, we may regard $\mathcal{O}(Y)$ as a subset of $\overline{\mathcal{O}(\Gamma)}$. Moreover, if $\lambda(Y) < \infty$, as is our usual assumption, then the same is true for all points in $\mathcal{O}(Y)$, since all points in this space share the same vertex groups which are necessarily invariant, by consequence of the fact that $\lambda(Y) < \infty$. Note also that λ is continuous on $\mathcal{O}(Y)$, because in general the displacement is continuous in the interior of *any* outer space.

Remark 7.3. Consider the situation given by the hypotheses of 7.1. The ϕ-invariance of core(A) allows us to build a straight map, $g : Y \to Y$, representing $[\phi]$ which leaves core(A) invariant. This map might not be optimal, but its Lipschitz constant provides an upper bound on the displacement of Y.

Now, along the path Y_t, we have the same topological trees (graphs of groups) except at the endpoint, X. We can thus re-scale edges but use the same topological straight map, g, to provide straight maps for all points Y_t except for X. From the invariance of core(A), one easily sees that there is a constant, C, so that $\lambda(Y_t) < C$ for all points on the path. (We can include X as well in this last statement).

The hypotheses of Theorem 4.5 therefore apply and we may deduce that $\lambda(X) \le \liminf_{t\to 0} \lambda(Y_t)$.

Corollary 7.4. *Let $X, Y \in \overline{\mathcal{O}(\Gamma)}$. We use the notation and hypotheses of Theorem 7.1. (In particular Δ_X is a simplicial face of Δ_Y). Let $f : X \to X$ be an optimal map representing $[\phi]$. Suppose further that τ is an f-illegal turn of X. Let Δ^τ be the simplex obtained by folding τ and let $X^\tau \in \Delta^\tau$ be the a point obtained from X by folding τ.*

Given $\varepsilon > 0$, there exists $t_\varepsilon > 0$, so that for all t smaller than t_ε, there exists an alternating simplicial path $\Sigma_t = (Z_i^t)_{i=0}^m$ in $\mathcal{O}(Y)$ from $Z_0^t = Y_t$ to a point $Z_m^t = Z_t$, so that

- *$\Delta_{Z_i^t}$ has Δ_X as a simplicial face for any $i = 0, \dots, m-1$,*
- *Δ_{Z_t} has Δ^τ as a simplicial face,*
- *for any point P of Σ_t we have $\lambda(X) - \varepsilon < \lambda(P) \le \lambda(Y_t)$;*
- *for $s \in [0,t]$ the map $s \mapsto Z_s$ parametrizes the segment from X^τ to Z_t.*

Proof. For this proof we will work entirely with trees. So Y will denote a Γ-forest, A an equivariant family of sub-trees — that is to say, the full pre-image in Y of an invariant subgraph $\underline{A} \subseteq \underline{Y}$ — and so on.

The map σ_t is that introduced in the statement of Theorem 7.1, and g_t is the map provided by Theorem 7.1. Also, the t_ϵ is that provided by Theorem 7.1.

We denote by A_t the metric copy of A in Y_t. By hypothesis there are two different segments α_τ, β_τ incident at the same vertex v in X such that f overlaps α_τ and β_τ. If $v \notin \sigma_t(A_t)$ then, for any small enough ε and $t < t_\varepsilon$, also g_t must overlap $\alpha = \sigma_t^{-1}(\alpha_\tau)$ and $\beta = \sigma_t^{-1}(\beta_\tau)$, and the claim follows by (equivariantly) performing the corresponding simple fold directed by g_t. Thus in this case the folding path consists of two points: $Z_0^t = Y_t$ and $Z_1^t = Z_t$. The inequality "$\leq \lambda(Y_t)$" follows because the fold is directed by an optimal map, the inequality "$> \lambda(X) - \varepsilon$" follows by lower semicontinuity of λ.

Otherwise, α and β are segments incident to the same component of A_t. If α and β are incident to the same point, then we proceed as above, so we can suppose that they are incident to different points of A_t.

For small enough ε and $t < t_\varepsilon$ we have that g_t overlaps some open sub-segments of α and β. Let $a \in \alpha$ and $b \in \beta$ such that $g_t(a) = g_t(b)$ and such that a is the closest possible to A_t. Let a' be the point where α is attached to A_t, and b' the one where β is attached to A_t.

Let γ' be the segment from a' to b' in A_t, and let γ be the segment between a and b in Y_t. Clearly $\gamma = [a, a'] \cup \gamma' \cup [b', b]$, where $[a, a']$ is the sub-segment of α from a to a', and $[b', b]$ is the sub-segment of β from b' to b. Note that $[a, a'] \neq \alpha$ and $[b', b] \neq \beta$ because α and β are open and a is the closest possible to A_t.

On γ we put an extra simplicial structure given by the pull-back via g_t: we declare new vertices of γ the points whose g_t-image is a vertex of Y_t. $g_t(\gamma)$ is a tree because Y_t is. Moreover, since $g_t(a) = g_t(b)$, the restriction of g_t to γ cannot be injective. In particular, if $x \in \gamma$ is a point such that $d_{Y_t}(g_t(x), g_t(a))$ is maximal, then x is a vertex of γ, and the two sub-segments of γ incident to x are completely overlapped.

Let Z_1^t be the tree obtained by equivariantly identify such segments. Note that $s \mapsto Z_1^s$ parametrizes the segment from X to Z_1^t. Clearly, g_t induces a map $g_t^1 : Z_1^t \to Z_1^t$. Such map is continuous and not necessarily straight. However,

$$\mathrm{Lip}(g_t^1) \leq \mathrm{Lip}(g_t)$$

and $\mathrm{Str}(g_t^1)$ still represents $[\phi]$. Since $\mathrm{Lip}(\mathrm{Str}(g_t^1)) \leq \mathrm{Lip}(g_t^1)$ he have

$$\lambda(Z_1^t) \leq \lambda(Y_t).$$

Let A_t' be the union of A_t and the orbits of $[a, a']$ and $[b', b]$. Since $[a, a'] \neq \alpha$ and $[b', b] \neq \beta$, then the collapsing of A_t' produces a point of Δ_X. As our identification occurred in A_t', it follows that $\Delta_{Z_1^t}$ has Δ_X as a simplicial face.

Also, since Y_t parametrizes the segment from X to Y, as t varies Z_1^t parametrizes the segment from X to Z_1^t.

Note that a priori we may have $\Delta_{Z_1^t} = \Delta_Y$, but in any case $\Delta_{Z_1^t}$ is either a (non necessarily proper) simplicial face of Δ_Y or vice versa.

In Z_1^t we have a simple path γ_1 resulting from γ by the cancellation of the two identified segments at x. By construction g_t^1 is simplicial. If g_t^1 is not injective on γ_1, we can iterate the above procedure and define points Z_i^t with

$$\lambda(Z_i^t) \leq \mathrm{Lip}(g_t) = \lambda(Y_t)$$

and such that $\Delta_{Z_i^t}$ has Δ_X as a simplicial face. Moreover either $\Delta_{Z_i^t}$ has $\Delta_{Z_{i-1}^t}$ as a simplicial face or vice versa, so the simplicial path we are producing is alternating. Since γ has a finite number of vertices, we must stop, and we do when γ_i is a single point. At this stage, α and β are incident to the same point and we are reduced to the initial case. Note that any $Z_i^t \to X$ as $t \to 0$, thus so does any point in segment from Z_i^t to Z_{i+1}^t.

Therefore by lower semicontinuity of λ for any $\varepsilon > 0$, since we have finitely many Z_i^t's, for sufficiently small t we have that for any i

$$\lambda(X) - \varepsilon < \lambda(Z_i^t)$$

and the same inequality holds for points in the segments from Z_i^t to Z_{i+1}^t. Thus, up to possibly replacing t_ϵ with a smaller positive number, we get that inequality of third bullet in the statement, holds for any $t < t_\epsilon$. □

Remark 7.5. The length of the simplicial path produced by Corollary 7.4 is bounded a priori by a constant depending only on rank(Γ). More precisely, consider the sequence of simplices $\Delta_{Z_i^t}$. It may happens that two consecutive $\Delta_{Z_i^t}$ and $\Delta_{Z_{i+1}^t}$ are equal, due to the fact that, in the proof of Corollary 7.4, we subdivided γ. Up to cancel such consecutive repetitions, the length of the sequence of $\Delta_{Z_i^t}$ is bounded by a constant depending on the complexity of A_t, hence on rank(Γ).

Corollary 7.6. *Let $X, Y \in \overline{\mathcal{O}(\Gamma)}$ and suppose that Δ_X is a simplicial face of Δ_Y. Suppose that $\lambda(X) > \lambda(Y)$.*

Moreover, suppose that X is an exit point for Δ_X[16]*, and let X_E be as Definition 4.3, chosen so that $\lambda(X_E) \geq \lambda(Y)$.*

Then there is a simplicial path $\Sigma = (W_i)$ in $\overline{\mathcal{O}(Y)}$, starting at Y and ending at X_E, with $W_i \in \mathcal{O}(Y)$ except possibly for the point X_E, such that for any point P of Σ we have

$$\lambda(Y) \leq \lambda(P) \leq L < \lambda(X)$$

for some $L < \lambda(X)$.

Proof. We inductively use Corollary 7.4: suppose that the exit point, X_E, is obtained by successive folds, $\tau_1, \ldots, \tau_m$. (So that $\Delta_{X_E} = \Delta^{\tau_m}$.)

We parametrize the segment between X and Y by $Y_t = tY + (1-t)X$. Lemma 4.9 and Lemma 4.10 imply that on the Euclidean segment from X to Y, the displacement is continuous, quasi-convex and strictly monotone near X. Hence, there exists a t (which can be taken to be arbitrarily small), such that Y_t satisfies $\lambda(X) - \varepsilon < \lambda(Y_t) < \lambda(X)$. We then plug this in to Corollary 7.4, to find a point Z_t, whose displacement satisfies $\lambda(X) - \varepsilon < \lambda(Z_t) < \lambda(X)$, and a simplicial path, in $\mathcal{O}(Y)$, from Y_t to Z_t, where all points met have the same displacement inequality, where the path starts at Δ_Y and ends at Δ^{τ_1}. Since $s \mapsto Z_s$ parametrizes the segment from X to Z_t, we are in position to apply Corollary 7.4 again to the point Z_t, noting that Δ_X is a simplicial face of Δ_{Z_t} and that $\lambda(Z_t) < \lambda(X)$.

We continue inductively.

Concatenating our paths, and adding the points Y and X_E, yields the result; the constant L is simply the maximum displacement of points of our paths. By construction the displacement is a number strictly less than $\lambda(X)$ on vertices. Since $\Sigma \subset \mathcal{O}(Y)$ except possibly for its last point X_E, the displacement is continuous and quasi-convex (Lemma 4.9) on Σ except possibly at X_E where it may jump, but still lower-semicontinuity is preserved (Theorem 4.5). This implies that $L < \lambda(X)$.

□

Remark 7.7. As in Remark 7.5, up to repetitions, the simplicial length of the path Σ provided by Corollary 7.6 is bounded a priori by a constant depending only on rank(Γ). This is because of Remark 7.5 and because the length of the path from X to X_E is bounded by the dimension of $\mathcal{O}(\Gamma)$.

[16] See Definition 4.3

8. End of the proof of Theorem 5.3: peak reduction on simplicial paths

We fix Γ as in Notation 3.7 and $[\phi] \in \mathrm{Out}(\Gamma)$. Let $\lambda = \lambda_\phi$. We will prove:

Lemma 8.1. *For any $L \geq \lambda(\phi)$, the level set*

$$\{X \in \overline{\mathcal{O}(\Gamma)}^\infty : \lambda(\phi) \leq \lambda_\phi(X) \leq L\}$$

is connected by L-calibrated simplicial paths.

This in particular gives the second claim of Theorem 5.3 (when $L = \lambda(\phi)$). Moreover, if Σ is any L-calibrated path (hence in the above level set), then, by possibly adding some extra vertices to Σ we obtain a path in the same level set, and that in addition is alternating. So Theorem 5.8 applies and Σ can be regenerated to $\mathcal{O}(\Gamma)$, and this proves first claim of Theorem 5.3.

We will proceed by induction and assume that Theorem 5.3 is true in any rank less than $\mathrm{rank}(\Gamma)$.

From now on we fix $A, B \in \overline{\mathcal{O}(\Gamma)}^\infty$ such that $\lambda(A), \lambda(B) \geq \lambda(\phi)$. For any $L \geq \max\{\lambda(A), \lambda(B)\}$ we denote by $\Sigma_L(A, B)$ the set of L-calibrated simplicial paths from A to B.

Lemma 8.2. *For some L, $\Sigma_L(A, B) \neq \emptyset$.*

Proof. Since $\lambda(A), \lambda(B) \geq \lambda(\phi)$, they have not jumped. Let $A' \in \mathrm{Hor}(A)$ and $B' \in \mathrm{Hor}(B)$, so that A has not jumped in $\Delta_{A'}$ and B has not jumped in $\Delta_{B'}$. Since $A', B' \in \mathcal{O}(\Gamma)$, which is connected, there is a simplicial path in $\mathcal{O}(\Gamma)$ between A', B'. We can therefore use Theorem 6.4 to obtain an element of Σ_L (where the L is the maximum displacement along such a path). □

Definition 8.3. For any calibrated path $\Sigma = (X_i)$ we say that X_i is a *peak* if $\lambda(X_i) = \lambda(\Sigma)$. A pair of two consecutive peaks X_{i-1}, X_i is called a *flat peak*. A peak is *strict* if it is not part of a flat peak.

To any Σ we can associate the triple $(\lambda(\Sigma), p, p_f) \in \mathrm{spec}(\phi) \times \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0}$ where p is the number of peaks, and p_f that of flat peaks. We order $\mathrm{spec}(\phi) \times \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0}$ with lexicographic order, from left to right. That is, $(\lambda, p, p_f) > (\lambda', p', p'_f)$ means:

- $\lambda > \lambda'$, or
- $\lambda = \lambda'$ and $p > p'$, or
- $\lambda = \lambda'$ and $p = p'$ and $p_f > p'_f$.

Lemma 8.4. *There exists $\Sigma_0 = (X_i) \in \Sigma_L(A, B)$, a calibrated path from A to B, which minimises (λ, p, p_f). Namely, Σ_0 minimizes, in order:*

(1) $\lambda(\Sigma)$;
(2) the number peaks;
(3) the number of flat peaks.

Proof. By Theorem 4.11 the set $\mathrm{spec}(\phi)$ is well-ordered, so $\mathrm{spec}(\phi) \times \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0}$ is lexicographically well-ordered. Therefore every minimising sequence must eventually realise the minimum. □

From now on we fix such a minimising Σ_0.

Note that if X is a strict peak of a path Σ, then λ is locally strictly monotone near X, on both sides of X in Σ. (By Lemma 4.9.)

Once again, we need the inductive hypothesis.

Lemma 8.5. *Suppose that Theorem 5.3 is true in any rank less than* $\mathrm{rank}(\Gamma)$*. Then* Σ_0 *has no strict peaks in its interior.*

Proof. Suppose that $\lambda(X_{i-1}) < \lambda(X_i) > \lambda(X_{i+1})$. Set $X = X_i$, $Y = X_{i-1}, Z = X_{i+1}$, so that $\lambda(Y), \lambda(Z) < \lambda(X)$.

By calibration, X minimizes λ in its simplex, hence Δ_X is a proper face of both Δ_Y and Δ_Z.

Since X is not a ϕ-minimally displaced point, by Lemma 4.13 $X \notin \mathrm{TT}(\phi) \subset \mathcal{O}(X)$. By Lemma 4.4, X is an exit point. Let X_E be as in Definition 4.3. Since X_E can be chosen arbitrarily close to X, we chose one so that $\lambda(X_E) \geq \max\{\lambda(Y), \lambda(Z)\}$.

Now we invoke Corollary 7.6 to get a simplicial path Σ in $\overline{\mathcal{O}(Y)}$ from Y to X_E, the displacement of whose points is between $\lambda(Y)$ and L, for some $L < \lambda(X)$. In particular $\lambda(\Sigma) < \lambda(X)$.

We now interpret this as a simplicial path in $\overline{\mathcal{O}(\Gamma)}$. Since $\lambda(Y) \geq \lambda(\phi)$ no point of such path jumps. We apply Theorem 6.4 to obtain a calibrated path Σ_Y from Y to X_E, whose displacement is less than $\lambda(X)$. By symmetry, we get a calibrated path Σ_Z from X_E to Z whose displacement is less than $\lambda(X)$. Let Σ_1 be the simplicial path obtained by following Σ_0 till Y, then Σ_Y, then Σ_Z and then again Σ_0 till its end. Since $\lambda(\Sigma_Y), \lambda(\Sigma_Z) < \lambda(X) = \lambda(\Sigma_0)$, we have $\lambda(\Sigma_1) \leq \lambda(\Sigma_0)$.

If $\lambda(\Sigma_1) < \lambda(\Sigma_0)$, we apply Theorem 6.4 and contradict the minimality of Σ_0. Otherwise, paths Σ_Y and Σ_Z do not contain peaks of Σ_1. Therefore Σ_1 is a $\lambda(\Sigma_0)$-calibrated which has fewer strict peaks than Σ_0, contradicting minimality. □

Lemma 8.6. Σ_0 *has no flat peaks unless* λ *is constant on* Σ_0 *and* $\lambda(\Sigma_0) = \lambda(\phi)$*.*

Proof. If the function λ is not constantly equal to $\lambda(\phi)$ on Σ_0, then in particular λ is strictly bigger than $\lambda(\phi)$ on peaks. Suppose that there is Y, X two consecutive vertices of Σ_0 with

$$\lambda(X) = \lambda(Y) = \lambda(\Sigma_0) > \lambda(\phi).$$

The idea is to find a third point Z to add between Y and X in order to destroy the flat peak. If there is a point Z in the interior of the segment YX, with $\lambda(\phi) \leq \lambda(Z) < \lambda(X) = \lambda(Y)$, then we add it.

Otherwise, λ is constant on $\overline{XY}$. Let W be a point in the interior of the segment $\overline{XY}$. If W is not a local minimum for λ in Δ_W, then near W we find Z with the above properties. We add it.

If W is a local minimum for λ in Δ_W then, by Lemma 4.13 and Lemma 4.4, near W in $\mathcal{O}(W)$ there is a point Z with the above properties and such that Δ_W is a finitary face of Δ_Z in $\mathcal{O}(W)$. We add Z.

In each case, we have added a point, Z, such that Δ_X and Δ_Y are faces of Δ_Z, and since the original path was calibrated, we can verify - using Theorem 4.7 - in each case that X, Y did not jump in Δ_Z. Hence we can add Z to the path. By Lemma 4.10, the new path is still a calibrated path (continuity at Z is automatic, since λ is continuous in $\mathcal{O}(W)$), with the same displacement as Σ_0, and the same number of peaks, but with one less flat peak, contradicting the minimality of Σ_0. □

It follows that the maximum displacement of points of Σ_0 is reached at endpoints. Thus Σ_0 is a calibrated simplicial path in the requested level set, proving Lemma 8.1. To finish the proof of Theorem 5.3, simply observe that we have shown that we can connect any two points in $\{X \in \overline{\mathcal{O}(\Gamma)}^{\infty} : \lambda_\phi(X) = \lambda(\phi)\}$ by a calibrated simplicial path with no peaks, either strict or flat, unless the displacement is constant. This immediately implies that the displacement is constant along the path. □

The following is an observation that may be helpful for algorithmic purposes.

Remark 8.7. If ϕ is irreducible, there exists a constant K, depending only on rank(Γ), such that, given a L-calibrated alternating simplicial path Σ having some peak in its interior, and such that the displacement is not constant along Σ, there exists a L-calibrated alternating simplicial path Σ' with either less displacement or one peak less, and whose simplicial length is increased at most by K.

This is because we can remove a strict peak from Σ as in Lemma 8.5 — if Σ contains no strict peak, we create one as in Lemma 8.6, without changing the global number of peaks nor $\lambda(\Sigma)$, and increasing the length of Σ by 1 —. The control on simplicial length comes from the use of Corollary 7.6 and Theorem 6.4 in the proof of Lemma 8.5:

By Remark 7.7 any use of Corollary 7.6 increase the simplicial length by a fixed amount, and since ϕ is irreducible, every calibrated path is in $O(\Gamma)$; therefore the calibration process Theorem 6.4 does not involve regeneration of paths, nor continuity issues, (so the alternating Σ_2 is already calibrated in the proof of Theorem 6.4), and it is readily checked that in this case calibration increases the length by a fixed amount.

9. Applications

In this section we show how the connectedness of the level sets gives a solution to some decision problems. Namely we will prove Theorems 2.9 and 2.5 and some generalisations. We will work with **graphs in the volume-one slice of CV_n.**

Recall that a point, X, of CV_n is called ε-thin if there is a homotopically non-trivial loop in X of length at most ε. Conversely, X is called ε-thick if it is not ε-thin.

Proposition 9.1 ([3, Proposition 10]. See also [13, Proposition 5.5], and [11, Section 8])**.** *Let $X \in CV_n$ (that is, X is a volume-one marked metric graph) and $f : X \to X$ a straight map representing some automorphism of F_n. Let $\lambda = Lip(f)$, let N equal the maximal length of chains of topological subgraphs of any graph in CV_n (this is clearly a finite number) and let μ be any real number greater than λ. Then if X is $1/((3n-3)\mu^{(N+1)})$-thin, then it has a nontrivial core sub-graph which is f-invariant up to homotopy, in particular the automorphism represented by f is reducible. For instance, one can take $N = 3n - 3$.*

Definition 9.2. A uniform rose in CV_n is a rose-graph (*i.e.* a bouquet of circles) whose edges all have the same length. Let $X \in CV_n$. Then we call R an adjacent uniform rose if it obtained by collapsing a maximal tree in X and then rescaling so that all edges in R have the same length.

Proposition 9.3. *Let $X \in CV_n$ be a point which is ε-thick and let R be any adjacent uniform rose (both of volume 1). Then, $\Lambda(X, R) \leq 1/\varepsilon$ and $\Lambda(R, X) \leq n$.*

Proof. By Theorem 4.1, we can look at candidates that realise the stretching factor. Since, topologically, one passes from X to R by collapsing a maximal tree, we get that a candidate in X, when mapped to R, crosses every edge at most twice. In fact the candidate crosses every edge of R at most once in the case of an embedded simple loop or an infinity-symbol loop. This gives the first inequality, on taking into account that X is ε-thick and that barbells have length at least 2ε.

For the second inequality note that an embedded loop in R is an edge and has length $1/n$ and lifts to an embedded loop in X, of length at most 1. An infinity-symbol loop in R consists of two distinct edges, has length $2/n$ and lifts to a loop in X which goes through every edge at most twice. (Barbells are not present in R). $\square$

Corollary 9.4. *Let $X \in CV_n$ be ε-thick and let R be an adjacent uniform rose. Consider $[\phi] \in \mathrm{Out}(F_n)$. Then $\Lambda(R, \phi R) \leq \frac{n}{\varepsilon}\Lambda(X, \phi X)$.*

Now, we use connectedness of level sets (Theorem 5.3) for deducing the following result.

Proposition 9.5. *Let R, R_∞ be two points in CV_n which are both uniform roses. Let $[\phi] \in \mathrm{Out}(F_n)$ be irreducible and suppose that μ is any real number greater than:*

$$\max\{\Lambda(R, \phi R), \Lambda(R_\infty, \phi R_\infty)\}.$$

Then there exist $R_0 = R, R_1, R_2, \ldots, R_k = R_\infty$, which are all uniform roses in CV_n such that:

- *For each i, there exists a simplex Δ_i such that Δ_{R_i} is a rose-face of both Δ_i and Δ_{i+1}.*
- *$\Lambda(R_i, \phi R_i) \leq \frac{n}{\varepsilon}\mu$, where $\varepsilon = 1/((3n-3)\mu^{(N+1)})$.*

Proof. This follows from Theorem 5.3, using Definition 5.1, since each pair Δ_i and Δ_{i+1} have a (at least one) common rose face; just take any uniform adjacent rose in any common rose face. The remaining point follows from Corollary 9.4 and Proposition 9.1. □

Proof of Theorem 2.5: We clearly have an algorithm which terminates (Remark 2.10), and it is apparent that if $\psi \in S_\phi$ then these automorphisms are conjugate. It remains to show the converse; that if they are conjugate, then $\psi \in S_\phi$.

Let R be the uniform rose corresponding to the basis B. If ψ were conjugate to ϕ, then there would be a conjugator, some $[\tau] \in \mathrm{Out}(F_n)$ such that $\psi = \tau\phi\tau^{-1}$. Let $R_\infty = \tau R$. Remind that the $\mathrm{Out}(F_n)$-action on CV_n is a right-action, namely $\phi(\psi(X)) = (\psi\phi)X$. In particular,

(1) $\quad ||\psi||_B = \Lambda(R, \psi R) = \Lambda(R, (\tau\phi\tau^{-1})R) = \Lambda(\tau^{-1}(\tau R), \tau^{-1}(\phi(\tau R))) = \Lambda(R_\infty, \phi R_\infty).$

Now we use Proposition 9.5 to find a sequence $R = R_0, R_1, \ldots, R_k = R_\infty$, such that each consecutive pair are incident to a common simplex and $\Lambda(R_i, \phi R_i) \leq n(3n-3)\mu^{3n-1} = K$.

Let τ_i so that $R_i = \tau_i R$. Since R_i and R_{i+1} are both incident to a common simplex, there exists a CMT automorphism ζ_i such that $\tau_i(\zeta_i(\tau_i^{-1}(R_i))) = R_{i+1}$. Thus

$$\zeta_i\tau_i R = \tau_i(\zeta_i(R)) = \tau_i(\zeta_i(\tau_i^{-1}(R_i))) = R_{i+1} = \tau_{i+1}R,$$

and up possibly compose ζ_i with a graph-automorphism of R, we may assume $\tau_{i+1} = \zeta_i\tau_i$. Therefore $\tau_{i+1} = \zeta_i \ldots \zeta_0$ (and we set $\tau_0 = id$).

Now let $\phi_i = \tau_i\phi\tau_i^{-1}$. Clearly $\phi_0 = \phi$ and $\phi_k = \psi$.

Since $\phi_{i+1} = \zeta_i\phi_i\zeta_i^{-1}$, to finish the proof we just need that $||\phi_i||_B \leq K$. This follows since, as in (1)

$$||\phi_i||_B = \Lambda(R, \phi_i R) = \Lambda(R_i, \phi R_i) \leq K.$$

□

We prove now Theorem 2.9. First a lemma,

Lemma 9.6. *Let X be a core graph and f a homotopy equivalence on X, having a proper subgraph X_0, with nontrivial fundamental group, such that $f(X_0) = X_0$. Then there is a maximal tree, T, such that the automorphism induced by f on the rose X/T is visibly reducible.*

Proof. Choose X_0 to be minimal. Therefore it will have components, $X_1, \ldots, X_k$ such that $f(X_i) = X_{i+1}$ with subscripts taken modulo k. Take a maximal tree for each X_i and extend this to a maximal tree, T, for X. It is then clear that if we take B_i to be the set of edges in X/T coming from X_i, that f_* will be visibly reducible as witnessed by $B_1, \ldots, B_k$. (Note each subgroups generated by each B_i are only permuted/preserved up to conjugacy, since the X_i are disjoint and so one cannot choose a common basepoint). $\square$

Proof of Theorem 2.9: The algorithm clearly terminates (Remark 2.10), and if there is a ψ in S^+ which is visibly reducible, then ϕ is reducible. It remains, therefore, to show that if ϕ is reducible, then there is some $\psi \in S^+$ which is visibly reducible.

We proceed much as in the proof of Theorem 2.5, but here we do not know that the points in CV_n we encounter will remain uniformly thick.

Let R be the uniform rose corresponding to the basis B. By Theorem 4.16, there exists an $X \in CV_n$ with a core invariant subgraph and such that $\Lambda(X, \phi(X)) < \mu$. By Theorem 5.3, there exists a simplicial path from R to X, whose vertices are points, $X_0 = R, X_1, \ldots, X_k = X$, such that $\Lambda(X, \phi(X_i)) < \mu$. Choose the maximal index, M, such that $X_0, X_1, \ldots, X_M$ are all ε-thick, where $\varepsilon = 1/((3n-3)\mu^{(N+1)})$ as in Proposition 9.1.

If $M < k$, then X_{M+1} is ε-thin, and by Proposition 9.1, we have that X_{M+1} has an optimal representative for $[\phi]$ which admits an invariant subgraph. Therefore, up to replacing X with X_{M+1}, we may assume that X_i is ε-thick for $i = 0, \ldots, k-1$.

Since X_k has an invariant subgraph, by Lemma 9.6, we may find an adjacent uniform rose face, R_k, so that the representative of $[\phi]$ at R_k is visibly reducible.

Now, for each $i \leq k-1$, we find a uniform rose R_i which is adjacent to both X_i and X_{i+1}, which exist by definition of simplicial path (Definition 5.1). Note that since $X_0 = R$ is a rose, then $R_0 = R$. Moreover, by Corollary 9.4 we have $\Lambda(R_i, \phi R_i) < K$ for any $i = 0, \ldots, k-1$.

We now conclude exactly as in the proof of Theorem 2.5: Let $[\tau] \in \mathrm{Out}(F_n)$ be such that $R_k = \tau R$, and let $\psi = \tau\phi\tau^{-1}$. Find CMT automorphisms ζ_i such that $\tau_i = \zeta_{i-1} \ldots \zeta_0$ satisfies $R_i = \tau_i R$ and $\tau_k = \tau$. Define $\tau_0 = id$ and $\phi_i = \tau_i \phi \tau_i^{-1}$, so that $\phi_0 = \phi$, $\phi_k = \psi$, and $\phi_{i+1} = \zeta_i \phi_i \zeta_i^{-1}$.

Since each $\Lambda(R_i, \phi R_i) < K$, as in (1), we get that each $\phi_i \in S_i$ for $i \leq k-1$. Hence $\psi \in S^+$ and is visibly reducible, as desired. $\square$

9.1. **Generalisations.** Our algorithms work in some more general setting that just free groups. For instance, consider the case of a group G equipped with a splitting $\mathcal{G} = (\{G_i\}, n)$ where the factor groups G_i are finite groups. In this case $\mathcal{O}_{\mathrm{gr}}(\mathcal{G})$ is a deformation space of finite graphs of groups with trivial edge-groups and finite vertex groups.

This leads to Theorem 2.11, which we now explain how to prove.

Theorems 2.5 and 2.9 generalise as follows. As above, we work in the volume-one slice of $\mathcal{O}_{\mathrm{gr}}(\mathcal{G})$. Instead of uniform roses one can use uniform "hairy roses", that is to say, graph $X \in \mathcal{O}_{\mathrm{gr}}(\mathcal{G})$ obtained from a rose by attaching, to the unique vertex, edges each ending with a non-free vertex. Uniform here means that all edges have the same length.

Any $X \in \mathcal{O}_{\mathrm{gr}}(\mathcal{G})$ is a face of a simplex containing a hairy rose simplex: to see this, first, for any non-free vertex v which is not a leaf, fold a little all edges at v; then, once all non-free vertices are leaves, collapse a maximal tree in the sub-graph consisting of edges incident only at free vertices. We say that a uniform hairy rose is adjacent to X if obtained in this way, plus a rescaling of edges.

Now define a CMT automorphisms of a hairy rose as a change of marking between two hairy roses 'adjacent' to a common point. More precisely, let Δ_1, Δ_2 be simplices in $\mathcal{O}_{\mathrm{gr}}(\mathcal{G})$, having a common face; let $X \in \overline{\Delta_1} \cap \overline{\Delta_2}$, and let R_1, R_2 be uniform hairy roses in $\overline{\Delta_1}, \overline{\Delta_2}$ respectively. Then we call R_1 and R_2 adjacent.

Then, letting R be a fixed marked hairy rose, we define,

$$CMT_R(G) = \{[\phi] \in \mathrm{Out}(G) \ : \ \phi(R) \text{ is adjacent to } R\}.$$

Remark 9.7. We note that this slightly different to the notion of adjacency in CV_n, but the idea is very similar. We start with an alternating simplicial path and want to replace each vertex along that path with a hairy rose. In CV_n, one can do this by replacing each point with a rose in such a way that consecutive roses are in faces of a common simplex. In this situation, moving to a hairy rose involves inserting 'stems' and then collapsing a maximal tree (ignoring the stems). However, there are several (although finitely many) ways of introducing these stems since the vertex groups are non-trivial. This means each vertex in the original simplicial path gives rise to two hairy roses - one can insert stems consistently between consecutive points, but not necessarily for three consecutive points - and in the resulting sequence of hairy roses, consecutive hairy roses are adjacent in the sense described above.

There are finitely many CMT automorphisms since the finiteness of the vertex groups implies that the stabiliser of any point is finite, and also that the deformation space $\mathcal{O}_{\mathrm{gr}}(\mathcal{G})$ is locally finite (and so there are only finitely many hairy roses adjacent to a given one). Moreover, since $\mathcal{O}_{\mathrm{gr}}(\mathcal{G})$ is connected, the CMT automorphisms generate $\mathrm{Out}(\mathcal{G})$.

Now we can build algorithms exactly as in Theorems 2.5 and 2.9. The fact that vertex groups are finite implies that Remark 2.4 holds true. So the set S in the statements is finite, and algorithms stop in finite time. The fact that there are finitely many CMT automorphisms implies that the set S^+ in Theorem 2.9 is finite.

The proof that these algorithms work now goes *mutatis mutandis* as in the case of CV_n. In particular, the conjugacy problem for irreducible automorphisms and the detection of reducibility are solvable in $\mathrm{Out}(\mathcal{G})$.

10. Appendix: proof of Theorem 7.1

In this section we give the proof of Theorem 7.1, which we restate for convenience (recall we are using Notation 3.7 and $[\phi] \in \mathrm{Out}(\Gamma)$).

Theorem (Theorem 7.1). *Let $X, Y \in \overline{\mathcal{O}_{gr}(\Gamma)}$. Suppose that Δ_X is a simplicial face of Δ_Y. Thus as graphs, Y is obtained by collapsing a sub-graph A. Suppose that* core(A) *is ϕ-invariant. For $t \in [0,1]$ let $Y_t = (1-t)X + tY$ be a parametrization of the Euclidean segment from X to Y. Let $\sigma_t : Y_t \to X$ be the map obtained by collapsing A and by linearly rescaling the edges in $Y \setminus A$.*

Let $f : X \to X$ be an optimal map representing $[\phi]$. Then for any $\varepsilon > 0$ there is $t_\varepsilon > 0$ such that $\forall 0 \leq t < t_\varepsilon$ there is an optimal map $g_t : Y_t \to Y_t$ representing $[\phi]$ such that

$$d_\infty(\sigma_t \circ g_t, f \circ \sigma_t) < \varepsilon.$$

Proof. We split the proof in two sub-cases. First when A is itself a core graph, and then the case when core(A) is trivial. Clearly the disjoint union of the two cases implies the mixed case.

We will work at once with graphs and trees, by using Notation 3.5.

Lemma 10.1 (When A is a core graph). *Let $X, Y \in \overline{\mathcal{O}_{gr}(\Gamma)}$. Suppose that as graphs of groups, X is obtained from Y by collapsing a ϕ-invariant core sub-graph $A = \sqcup A_i$. Then the conclusion of Theorem 7.1 holds.*

Proof. We begin by fixing some notation. First of all, we will use the symbol λ to denote any of the displacement functions of ϕ (i.e. $\lambda_\phi, \lambda_{\phi|_A}, \dots$). If x is a point in a metric space, we denote by $B_r(x)$ the open metric ball centered at x and radius r. For any i, we denote by v_i the non-free vertex of X obtained by collapsing A_i. For any t we denote by A^t the metric copy of A in Y_t. Note that A is uniformly collapsed in Y_t, that is to say, $[A^t] \in \mathbb{P}\mathcal{O}(A)$ is the same element for any $0 < t \leq 1$, and we have $\mathrm{vol}(A^t) = t\,\mathrm{vol}(A^1)$.

By lower semicontinuity of λ (Theorem 4.5) we have that

$$\forall \varepsilon_0 > 0 \exists t_{\varepsilon_0} > 0 \text{ such that } \forall t < t_{\varepsilon_0} \text{ we have } \lambda(Y_t) > \frac{\lambda(X)}{1+\varepsilon_0}. \tag{2}$$

A priori f may collapse some edge, in any case $\forall \varepsilon_1 > 0 \exists f_1 : X \to X$ a straight map representing $[\phi]$ such that f_1 does not collapse any edge, and

$$d_\infty(f, f_1) < \varepsilon_1 \quad \text{and} \quad \mathrm{Lip}(f_1) < \mathrm{Lip}(f)(1+\varepsilon_1) = \lambda(X)(1+\varepsilon_1). \tag{3}$$

Moreover $\exists 0 < \rho_0 = \rho_0(X, f_1)$ such that $\forall \rho < \rho_0$

- $B_\rho(x)$ is star-shaped for any $x \in X$ (i.e. it contains at most one vertex);
- for any i, each connected component of $f_1^{-1}(B_\rho(v_i))$ is star-shaped and contains exactly one pre-image of v_i;
- for any i, j the connected components of $f_1^{-1}(B_\rho(v_i))$ and those of $f_1^{-1}(B_\rho(v_j))$ are pairwise disjoint.

We fix an optimal map $\varphi : A^1 \to A^1$ representing $[\phi|_A]$. Since $[A^t] \in \mathbb{P}\mathcal{O}(A)$ does not depend on t, $\varphi : A^t \to A^t$ is an optimal map for any $t \in (0,1]$ and the Lipschitz constant does not change. Clearly (by Sausage Lemma 4.1)

$$\mathrm{Lip}(\varphi) \leq \lambda(Y_t) \quad \text{for any } t. \tag{4}$$

The natural option is to define g_t by using $\sigma_t^{-1} \circ f_1 \circ \sigma_t$. Hence, we need to deal with places where σ_t^{-1} is not defined. (We have to understand how to deal with arcs in X whose f_1-image crosses some v_i.)

We fix lifts $\widetilde{\varphi}$ of φ and $\widetilde{f_1}$ of f_1. For any v_i, and any $x \in f_1^{-1}(v_i)$, to any germ of edge α at x we associate a path $\gamma_\alpha \in Y$ as follows. We do two different constructions: one in case x is one of the v_j's, and another for the case when x is different from others v_j's.

Case 1. Suppose $x = v_k$ and $f_1(x) = v_i$ for some k, i (not necessarily different). Let α be a germ of edge at x. First of all we **choose** a lift $\widetilde{\alpha}$ of α. All subsequent choices of lifts of objects, made during the definition of γ_α, will depend on, and will be uniquely determined by, the choice of $\widetilde{\alpha}$. After having defined γ_α, we forget about such choices of lifts.

The germ α corresponds to a germ $\alpha_Y (= \sigma_t^{-1}(\alpha))$ in Y incident to A_k at a point that we denote by p_α. The lift $\widetilde{\alpha}$ corresponds to a germ $\widetilde{\alpha}_Y$ incident to $\widetilde{p}_\alpha \in \widetilde{A}_k$, where $\widetilde{p}_\alpha$ is a preimage of p_α and $\widetilde{A}_k$ is the component of the preimage of A_k containing $\widetilde{p}_\alpha$. (See Figure 1.)

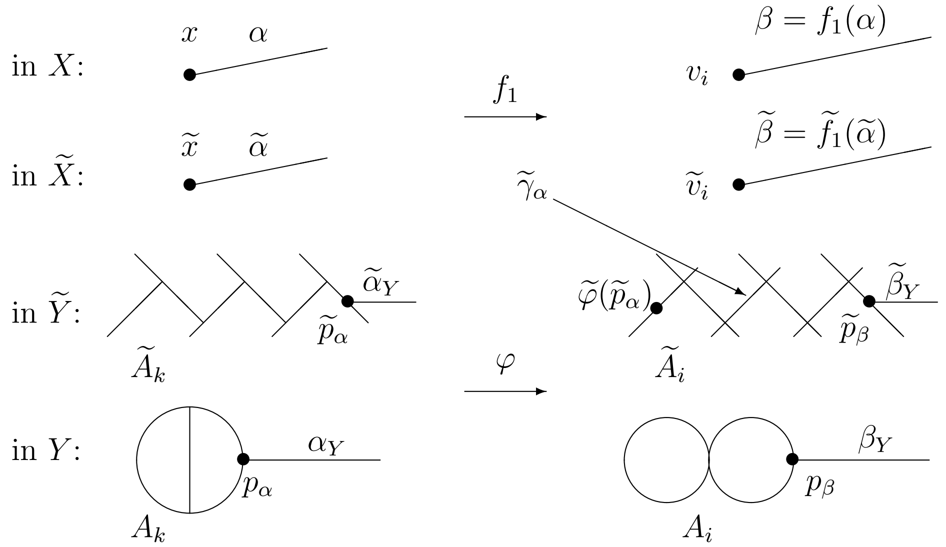

FIGURE 1. How to choose the paths $\widetilde{\gamma}_\alpha$

Let $\beta = f_1(\alpha)$ and choose $\widetilde{\beta}$ to be the lift of β so that $\widetilde{f_1}(\widetilde{\alpha}) = \widetilde{\beta}$. Note that in case $f_1(\alpha) = \alpha$, $\widetilde{\beta}$ is not necessarily equal to $\widetilde{\alpha}$ (it is only in the same orbit).

Clearly $\widetilde{\beta}$ emanates from a lift $\widetilde{v}_i$ of v_i so that $\widetilde{f_1}(\widetilde{x}) = \widetilde{v}_i$. The germ $\widetilde{\beta}$ corresponds to a germ $\widetilde{\beta}_Y$ incident to $\widetilde{A}_i$ at a point $\widetilde{p}_\beta$, where $\widetilde{A}_i$ is the component of the preimage of A_i so that $\widetilde{\varphi}(\widetilde{A}_k) = \widetilde{A}_i$. We define $\widetilde{\gamma}_\alpha$ as the unique geodesic path in $\widetilde{A}_i$ connecting $\widetilde{\varphi}(p_\alpha)$ to $\widetilde{p}_\beta$. Now we define γ_α as the projection to Y of $\widetilde{\gamma}_\alpha$. It is a path from $\varphi(p_\alpha)$ to p_β.

Remark 10.2. We chose a path $\widetilde{\gamma}_\alpha$ for any germ α in X, which is a finite graph. Therefore we have only finitely many such $\widetilde{\gamma}_\alpha$'s. We can then complete that family of paths by equivariance.

Remark 10.3. If we use $\widetilde{\alpha}$ instead $g\widetilde{\alpha}$, then bot $\widetilde{\varphi}(\widetilde{p}_\alpha)$ and $\widetilde{p}_\beta$ — and therefore also $\widetilde{\gamma}_\alpha$ — are translated by $\phi(g)$, hence the path γ_α is actually independent on the choice of the lift $\widetilde{\alpha}$.

Case 2. Let $x \in X$ be such that $f_1(x) = v_i$ for some i, but x is not one of the v_j's. (In case x is not a vertex, up to add x to the simplicial structures of X and Y, so we can consider it as a vertex.) For any germ of edge α at x we define γ_α as follows.

First, we fix a base-point $x_i \in A_i$, and for any component $\widetilde{A}_i$ (of the preimage of A_i) we choose a lift $\widetilde{x}_i \in \widetilde{A}_i$. Any germ of edge α at x corresponds to a germ α_Y is Y. For any such α we choose a lift $\widetilde{\alpha}$. Since f_1 does not collapse edges, $\widetilde{f_1}(\widetilde{\alpha})$ is a germ of edge $\widetilde{\beta}$ at some lift $\widetilde{v}_i$ of v_i, and corresponds to a germ $\widetilde{\beta}_Y$ at $\widetilde{A}_i$ in $\widetilde{Y}$. Let $\widetilde{\gamma}_\alpha$ be the unique path in $\widetilde{A}_i$ connecting $\widetilde{x}_i$ and $\widetilde{\beta}_Y$. We finally define γ_α as the projection to Y of $\widetilde{\gamma}_\alpha$.

Remark 10.4. As above we chose only finitely many such $\widetilde{\gamma}_\alpha$'s and we can complete the choices equivariantly.

Remark 10.5. The path γ_α actually depends on the choices of x_i and $\widetilde{x}_i$, but for any pair of germs α_1, α_2 at x, the reduced version of the concatenation $\gamma_{\alpha_1}^{-1}\gamma_{\alpha_2}$ does not depend on such choices.

Note that, as germs, $\alpha_Y = \sigma_t^{-1}(\alpha)$ and $\beta_Y = \sigma_t^{-1}(\beta) = \sigma_t^{-1}(f_1(\alpha))$. Now we have a path $\gamma_\alpha \subset A$ for any pre-image of germs at the v_i's, chosen independently on t. Let $t \in (0,1]$. We define a map

$$\overline{g}_t : Y_t \to Y_t$$

representing $[\phi]$ as follows:

- in $\sigma_t^{-1}\big(X \setminus f_1^{-1}(\sqcup_i B_\rho(v_i))\big)$ we just set $\overline{g}_t = \sigma_t^{-1} \circ f_1 \circ \sigma_t$;
- in $\sigma_t^{-1}\big(f_1^{-1}(\sqcup_i B_\rho(v_i))\big) \setminus A^t$ we use the paths γ_α. More precisely, let N be a connected component of $f_1^{-1}(B_\rho(v_i))$ and let $x \in N$ such that $f_1(x) = v_i$. For any edge $\alpha \in N$ emanating from x we define $\overline{g}_t(\sigma_t^{-1}(\alpha))$ by mapping linearly[17] $\sigma_t^{-1}(\alpha)$ to the path given by the concatenation of $\beta_Y = \sigma_t^{-1}(f_1(\alpha))$ and γ_α. Note that $\overline{g}_t|_{\sigma_t^{-1}(\alpha)} = \mathrm{Str}(\overline{g}_t|_{\sigma_t^{-1}(\alpha)})$.
- in A^t we set $\overline{g}_t = \varphi$;

finally, we set

$$g_t = \mathrm{opt}(\mathrm{Str}(\overline{g}_t))$$

where straightening and optimization are made with respect to the metric structure of Y_t. We now estimate the Lipschitz constant of $\overline{g}_t$. Clearly we have the lower bound

$$\lambda(Y_t) = \mathrm{Lip}(g_t) \leq \mathrm{Lip}(\overline{g}_t).$$

Moreover, since on edges of $Y_t \setminus A^t$ the map σ_t is just a rescaling of edge-lengths, for any $\varepsilon_2 > 0$ there is $t_{\varepsilon_2} > 0$ such that $\forall t < t_{\varepsilon_2}$

$$\mathrm{Lip}(\sigma_t) < 1 + \varepsilon_2 \qquad \mathrm{Lip}(\sigma_t^{-1}) < 1 + \varepsilon_2. \tag{5}$$

Now we compute an upper bound for $\mathrm{Lip}(\overline{g}_t)$. As $\overline{g}_t$ is defined in three different regions, namely

- $\Omega_1 = \sigma_t^{-1}\big(X \setminus f_1^{-1}(\sqcup_i B_\rho(v_i))\big)$,
- $\Omega_2 = \sigma_t^{-1}\big(f_1^{-1}(\sqcup_i B_\rho(v_i))\big) \setminus A^t$,
- $\Omega_3 = A^t$;

we will estimate $\mathrm{Lip}(\overline{g}_t)$ on these three regions separately.

In Ω_1 we have $\overline{g}_t = \sigma_t^{-1} \circ f_1 \circ \sigma_t$. Then

$$\mathrm{Lip}(\overline{g}_t|_{\Omega_1}) \leq \mathrm{Lip}(\sigma_t^{-1})\,\mathrm{Lip}(f_1)\,\mathrm{Lip}(\sigma_t).$$ [18]

Hence, by (3), (5), and by setting $(1+\varepsilon_2)^2(1+\varepsilon_1) = 1+\varepsilon_3$, we have

$$\mathrm{Lip}(\overline{g}_t|_{\Omega_1}) \leq (1+\varepsilon_2)^2\lambda(X)(1+\varepsilon_1) = (1+\varepsilon_3)\lambda(X). \tag{6}$$

Now, we switch to Ω_2. Let N be a connected component of $f_1^{-1}(\sqcup_i B_\rho(v_i))$. Let $x \in N$ such that $f_1(x) = v_i$ and let α be an edge of N emanating from x. By definition $\overline{g}_t$ is

[17] I.e. at constant speed

[18] Note that $\mathrm{Lip}(\sigma_t)$ and $\mathrm{Lip}(\sigma_t^{-1})$ are not the inverse of each other because different edges are stretched by σ_t by a priori different amounts.

linear on $\sigma_t^{-1}(\alpha)$, thus in order to estimate its Lipschitz constant we need to know only the lengths of $\sigma_t^{-1}(\alpha)$ and its image. Clearly

$$L_X(\alpha) = L_X(\sigma_t(\sigma_t^{-1}(\alpha))) \leq \mathrm{Lip}(\sigma_t) L_{Y_t}(\sigma_t^{-1}(\alpha)) \quad \text{and thus} \quad L_{Y_t}(\sigma_t^{-1}(\alpha)) \geq \frac{L_X(\alpha)}{\mathrm{Lip}(\sigma_t)}.$$

Moreover, since we have $L_X(f_1(\alpha)) = \rho$, we get

$$\rho \leq \mathrm{Lip}(f_1) L_X(\alpha) \qquad \text{and so} \qquad L_X(\alpha) \geq \frac{\rho}{\mathrm{Lip}(f_1)}$$

whence, by (5) and (3), we obtain

$$L_{Y_t}(\sigma_t^{-1}(\alpha)) \geq \frac{\rho}{\mathrm{Lip}(\sigma_t)\,\mathrm{Lip}(f_1)} > \frac{\rho}{(1+\varepsilon_2)\,\mathrm{Lip}(f_1)} > \frac{\rho}{\lambda(X)(1+\varepsilon_1)(1+\varepsilon_2)}.$$

Since γ_α is the same path in A for every t, its length in A^t depends linearly on t, namely here is a constant C_α such that

$$L_{Y_t}(\gamma_\alpha) = C_\alpha t$$

whence, setting $C = \max_\alpha C_\alpha$,

$$\begin{aligned}
\mathrm{Lip}(\overline{g}_t|_{\sigma_t^{-1}(\alpha)}) &\leq \frac{L_{Y_t}(\sigma_t^{-1}(f_1(\alpha)) + L_{Y_t}(\gamma_\alpha)}{L_{Y_t}(\sigma_t^{-1}(\alpha))} \leq \frac{\mathrm{Lip}(\sigma_t^{-1})\rho + tC}{L_{Y_t}(\sigma_t^{-1}(\alpha))} \\
&< ((1+\varepsilon_2)\rho + tC)\frac{\lambda(X)(1+\varepsilon_1)(1+\varepsilon_2)}{\rho} \\
&< (1+\varepsilon_2)(\rho + tC)\frac{\lambda(X)(1+\varepsilon_1)(1+\varepsilon_2)}{\rho} \\
&= \lambda(X)(1+\varepsilon_3)(1+\frac{tC}{\rho}).
\end{aligned}$$

Therefore $\forall \varepsilon_4 > 0 \exists t_{\varepsilon_4} > 0$ such that $\forall t < t_{\varepsilon_4}$ and for any α as above, we have $\mathrm{Lip}(\overline{g}_t|_{\sigma_t^{-1}(\alpha)}) < \lambda(X)(1+\varepsilon_4)$ and so

$$\mathrm{Lip}(\overline{g}_t|_{\Omega_2}) = \sup_\alpha \mathrm{Lip}(\overline{g}_t|_{\sigma_t^{-1}(\alpha)}) < \lambda(X)(1+\varepsilon_4). \tag{7}$$

Finally, on $\Omega_3 = A^t$ we have $\overline{g}_t = \varphi$ and so $\mathrm{Lip}(\overline{g}_t|_{A^t}) = \mathrm{Lip}(\varphi)$. Thus, by 4

$$\mathrm{Lip}(\overline{g}_t|_{\Omega_3}) \leq \lambda(Y_t). \tag{8}$$

Since by (2) $\lambda(X) \leq \lambda(Y_t)(1+\varepsilon_0)$, by putting together (6), (7), and 8 we have that for any $\varepsilon_5 > 0$ there is $t_{\varepsilon_5} > 0$ such that for any $t < t_{\varepsilon_5}$ we have

$$\mathrm{Lip}(\overline{g}_t) \leq \lambda(Y_t)(1+\varepsilon_5).$$

We are now in position to obtain the inequality claimed in the statement. Since g_t is optimal, $\mathrm{Lip}(g_t) = \lambda(Y_t)$, and by Theorem 4.2

$$d_\infty(g_t, \overline{g}_t) < \mathrm{vol}(Y_t)(\mathrm{Lip}(\overline{g}_t) - \lambda(Y_t)) < \mathrm{vol}(Y_t)\lambda(Y_t)\varepsilon_5.$$

We first estimate

$$d_\infty(\sigma_t \circ \overline{g}_t, f_1 \circ \sigma_t).$$

In $\sigma_t^{-1}\big(X \setminus f_1^{-1}(\sqcup_i B_\rho(v_i))\big)$ we have $\overline{g}_t = \sigma_t^{-1} \circ f_1 \circ \sigma_t$ so here the distance is zero. On A^t, since $\overline{g}_t(A) = A$, for any i there is j such that we have $\sigma_t(\overline{g}_t(A_i)) = \sigma_t(A_j) = v_j = f_1(v_i)$, hence also in A^t the distance is zero. Finally, let N be a connected component of $f_1^{-1}(\sqcup_i B_\rho(v_i))$. Let $x \in N$ such that $f_1(x) = v_i$ and let α be an edge of N emanating from x. The path $\overline{g}_t(\sigma_t^{-1}(\alpha))$ is given by the concatenation of $\sigma_t^{-1}(f_1(\alpha))$ with γ_α. The latter is collapsed by σ_t, and the image of the former is just $f_1(\alpha) = f_1 \circ \sigma_t(\sigma_t^{-1}(\alpha))$. Since the length of γ_α in A^t is bounded by tC we have that

$$d_\infty(\sigma_t \circ \overline{g}_t, f_1 \circ \sigma_t) \to 0 \qquad \text{as } t \to 0.$$

In particular $\forall \varepsilon_6 \exists t_{\varepsilon_6}$ such that $\forall t < t_{\varepsilon_6}$ we have

$$d_\infty(\sigma_t \circ \overline{g}_t, f_1 \circ \sigma_t) < \varepsilon_6.$$

Finally,

$$\begin{aligned}
& d_\infty(\sigma_t \circ g_t, f \circ \sigma_t) \\
\leq\ & d_\infty(\sigma_t \circ g_t, \sigma_t \circ \overline{g}_t) + d_\infty(\sigma_t \circ \overline{g}_t, f_1 \circ \sigma_t) + d_\infty(f_1 \circ \sigma_t, f \circ \sigma_t) \\
\leq\ & \mathrm{Lip}(\sigma_t) d_\infty(g_t, \overline{g}_t) + \varepsilon_6 + d_\infty(f_1, f) \\
<\ & (1 + \varepsilon_2)\, \mathrm{vol}(Y_t) \lambda(Y_t) \varepsilon_5 + \varepsilon_6 + \varepsilon_1
\end{aligned}$$

which is arbitrarily small for $t \to 0$. □

Lemma 10.6 (When core(A) is trivial)**.** *Let $X, Y \in \overline{\mathcal{O}_{gr}(\Gamma)}$. Suppose that as graphs of groups, X is obtained from Y by collapsing a sub-forest $A = \sqcup A_i$ whose tree A_i each contains at most one non-free vertex. Then the conclusion of Theorem 7.1 holds.*

Proof. Except the definition of g_t, the proof goes exactly as that of Lemma 10.1, and it is even simpler. So let's define g_t. As above A^t denote the scaled version of A. Let v_i be the vertex of X resulting from the collapse of A_i. The function λ is now continuous

$$\lambda(Y_t) \to \lambda(X).$$

As above, if f collapses some edge we find $f_1 : X \to X$ a straight map representing $[\phi]$ which collapses no edge and with

$$d_\infty(f, f_1) < \varepsilon_1 \quad \text{and} \quad \mathrm{Lip}(f_1) < \mathrm{Lip}(f)(1 + \varepsilon_1) = \lambda(X)(1 + \varepsilon_1).$$

We choose ρ so that $B_\rho(v_i)$ is star-shaped, the components of $f_1^{-1}(B_\rho(v_i))$ are star-shaped and contain a unique pre-image of v_i, and so that the components of $f_1^{-1}(B_\rho(v_i))$ and $f_1^{-1}(B_\rho(v_j))$ are pairwise disjoint. Finally we chose ρ small enough so that if $f(v_i) \notin \{v_j\}$, then $f(v_i) \notin \cup_j B_\rho(v_j)$.

For any i we choose a base vertex $x_i \in A_i$ which is the non-free vertex of A_i if any. For any $x \in X$ such that $f_1(x) = v_i$ and for any edge α in $f_1^{-1}(B_\rho(v_i))$ incident to x, let γ_α be the unique embedded path connecting $\sigma_t^{-1}(f_1(\alpha))$ to x_i. We define $\overline{g}_t : Y_t \to Y_t$ as follows:

- in $\sigma_t^{-1}\big(X \setminus f_1^{-1}(\sqcup_i B_\rho(v_i))\big)$ we just set $\overline{g}_t = \sigma_t^{-1} \circ f_1 \circ \sigma_t$;
- in $\sigma_t^{-1}\big(f_1^{-1}(\sqcup_i B_\rho(v_i))\big) \setminus A^t$ we use the paths γ_α. More precisely, let N be a connected component of $f_1^{-1}(B_\rho(v_i))$ and let $x \in N$ such that $f_1(x) = v_i$. For any edge $\alpha \in N$ emanating from x we define $\overline{g}_t(\sigma_t^{-1}(\alpha))$ by mapping linearly[19] $\sigma_t^{-1}(\alpha)$ to the path given by the concatenation of $\sigma_t^{-1}(f_1(\alpha))$ and γ_α. Note that $\overline{g}_t|_{\sigma_t^{-1}(\alpha)} = \mathrm{Str}(\overline{g}_t|_{\sigma_t^{-1}(\alpha)})$.
- in the components A_i^t so that $f_1(v_i) = v_j$, we set $g(A_i^t) = x_j$;

finally we set $g_t = \mathrm{opt}(\mathrm{Str}(\overline{g}_t))$. The estimates on Lipschitz constants and distances now follow exactly as in the proof of Lemma 10.1. □

[19]I.e. at constant speed

References

[1] Yael Algom-Kfir, *Strongly contracting geodesics in outer space*, Geom. Topol. **15** (2011), no. 4, 2181–2233. MR 2862155

[2] Yael Algom-Kfir and Mladen Bestvina, *Asymmetry of outer space*, Geom. Dedicata **156** (2012), 81–92. MR 2863547

[3] Mladen Bestvina, *A Bers-like proof of the existence of train tracks for free group automorphisms*, Fund. Math. **214** (2011), no. 1, 1–12. MR 2845630

[4] Mladen Bestvina, Mark Feighn, and Michael Handel, *The Tits alternative for* $\mathrm{Out}(F_n)$*. II. A Kolchin type theorem*, Ann. of Math. (2) **161** (2005), no. 1, 1–59. MR 2150382

[5] Mladen Bestvina and Michael Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. (2) **135** (1992), no. 1, 1–51. MR 1147956

[6] Mark Feighn and Michael Handel, *Algorithmic constructions of relative train track maps and cts.*

[7] ———, *The recognition theorem for* $\mathrm{Out}(F_n)$, Groups Geom. Dyn. **5** (2011), no. 1, 39–106. MR 2763779

[8] Stefano Francaviglia. Geodesic currents and length compactness for automorphisms of free groups. *Transactions of the American Mathematical Society*, 361(1):161–176, 2009.

[9] Stefano Francaviglia and Armando Martino, *Metric properties of outer space*, Publ. Mat. **55** (2011), no. 2, 433–473. MR 2839451

[10] ———, *The isometry group of outer space*, Adv. Math. **231** (2012), no. 3-4, 1940–1973. MR 2964629

[11] ———, *Stretching factors, metrics and train tracks for free products*, Illinois J. Math. **59** (2015), no. 4, 859-899.

[12] ———, *On the connectivity of level sets of automorphisms of free groups, with applications to decision problems*, arXiv:1703.09945.

[13] ———, *Displacements of automorphisms of free groups I: Displacement function, minpoints and train tracks*, Trans. of AMS **374** (2021), no. 5, 3215-3264.

[14] S. M. Gersten, *Addendum: "On fixed points of certain automorphisms of free groups"*, Proc. London Math. Soc. (3) **49** (1984), no. 2, 340–342. MR 748994

[15] ———, *On fixed points of certain automorphisms of free groups*, Proc. London Math. Soc. (3) **48** (1984), no. 1, 72–90. MR 721773

[16] ———, *Fixed points of automorphisms of free groups*, Adv. in Math. **64** (1987), no. 1, 51–85. MR 879856

[17] Vincent Guirardel and Gilbert Levitt, *The outer space of a free product*, Proc. Lond. Math. Soc. (3) **94** (2007), no. 3, 695–714. MR 2325317

[18] Michael Handel and Lee Mosher, *The free splitting complex of a free group, I: hyperbolicity*, Geom. Topol. **17** (2013), no. 3, 1581–1672. MR 3073931

[19] ———, *The free splitting complex of a free group, II: Loxodromic outer automorphisms*, Trans. Amer. Math. Soc. **372** (2019), no. 6, 4053–4105. MR 4009387

[20] Camille Horbez *Hyperbolic graphs for free products, and the Gromov boundary of the graph of cyclic splittings.* Journal of topology 9.2 (2016): 401-450.

[21] Ilya Kapovich, *Algorithmic detectability of iwip automorphisms*, Bull. Lond. Math. Soc. **46** (2014), no. 2, 279–290. MR 3194747

[22] Ilya Kapovich, *Detecting fully irreducible automorphisms: a polynomial time algorithm*, https://arxiv.org/abs/1609.03820.

[23] Jérôme E. Los, *On the conjugacy problem for automorphisms of free groups*, Topology **35** (1996), no. 3, 779–808, With an addendum by the author. MR 1396778

[24] Martin Lustig, *Conjugacy and centralizers for iwip automorphisms of free groups*, Geometric group theory, Trends Math., Birkhäuser, Basel, 2007, pp. 197–224. MR 2395795

[25] Sebastian Meinert, *The Lipschitz metric on deformation spaces of G-trees*, Algebr. Geom. Topol. **15** (2015), no. 2, 987–1029. MR 3342683

[26] John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565. MR 695906

[27] E. Ventura, *Fixed subgroups in free groups: a survey*, Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001), Contemp. Math., vol. 296, Amer. Math. Soc., Providence, RI, 2002, pp. 231–255. MR 1922276