

DISPLACEMENTS OF AUTOMORPHISMS OF FREE GROUPS I: DISPLACEMENT FUNCTIONS, MINPOINTS AND TRAIN TRACKS

STEFANO FRANCAVIGLIA

Dipartimento di Matematica of the University of Bologna

ARMANDO MARTINO

Mathematical Sciences, University of Southampton

ABSTRACT. This is the first of two papers in which we investigate the properties of the displacement functions of automorphisms of free groups (more generally, free products) on Culler-Vogtmann Outer space and its simplicial bordification - the free splitting complex - with respect to the Lipschitz metric. The theory for irreducible automorphisms being well-developed, we concentrate on the reducible case. Since we deal with the bordification, we develop all the needed tools in the more general setting of deformation spaces, and their associated free splitting complexes.

In the present paper we study the local properties of the displacement function. In particular, we study its convexity properties and the behaviour at bordification points, by geometrically characterising its continuity-points. We prove that the global-simplex-displacement spectrum of $Aut(F_n)$ is a well-ordered subset of \mathbb{R} , this being helpful for algorithmic purposes. We introduce a weaker notion of train tracks, which we call *partial train tracks* (which coincides with the usual one for irreducible automorphisms) and we prove that, for any automorphism, points of minimal displacement - minpoints - coincide with the marked metric graphs that support partial train tracks. We show that any automorphism, reducible or not, has a partial train track (hence a minpoint) either in the outer space or its bordification. We show that, given an automorphism, any of its invariant free factors is seen in a partial train track map.

In a subsequent paper we will prove that level sets of the displacement functions are connected, and we will apply that result to solve certain decision problems.

Note: the two papers were originally packed together in the preprint `arxiv:1703.09945`. We decided to split that paper following the recommendations of a referee.

Current date: June 17, 2020

E-mail addresses: stefano.francaviglia@unibo.it, A.Martino@soton.ac.uk.
1991 *Mathematics Subject Classification.* 20E06, 20E36, 20E08.

CONTENTS

1. Introduction	2
1.1. Overview	2
1.2. Anticipating the results	4
2. Setting, notation, and general definitions	6
2.1. Motivation for new definitions	6
2.2. Notation for free splittings	7
2.3. G -graphs and G -trees	8
2.4. Outer spaces	10
2.5. Simplicial structure	12
2.6. Simplicial bordification	13
2.7. Horoballs and regeneration	14
2.8. The groups $\text{Aut}(\Gamma)$ and $\text{Out}(\Gamma)$	14
3. Straight maps, gate structures, and optimal maps.	15
3.1. Straight maps	16
4. Displacement function and train track maps for automorphisms	26
5. Behaviour of λ at bordification points	34
6. Convexity properties of the displacement function	42
7. Existence of minimal displaced points and train tracks at the bordification	45
References	53

1. INTRODUCTION

1.1. **Overview.** Let F_n denote the free group of rank n and $\text{Out}(F_n) = \text{Aut}(F_n)/\text{Inn}(F_n)$ be the group of outer automorphisms. The natural space upon which this acts is CV_n , Culler-Vogtmann Space, which in turn admits a (non-symmetric) metric, the Lipschitz metric (see [7], [23]).

The motivation for this paper is to extend the main results of [9], and in particular the result that for irreducible elements, ϕ , of $\text{Out}(F_n)$, the points that are minimally displaced by ϕ in CV_n (with respect to the Lipschitz metric) coincide exactly with the points that support train track representatives for ϕ .

In order to do this, we need to extend our space to the free splitting complex, \mathcal{FS}_n - see [15] and [17] for more details, and a complete exposition. However, our methods rely on inductive arguments which require us to deal with a more general setting, of a deformation space and its free splitting complex, as follows.

Let Γ be a free product of groups with a specified free splitting. That is, abusing notation, Γ is a group $G_1 * \dots * G_k * F_r$, where the F_r is a free group and the G_i are specified. This is not, necessarily, the Grushko decomposition, and we allow the G_i to be freely decomposable, or even free. Equivalently, we can think of Γ as given by a specific graph of groups with trivial edge groups. (We shall suppress this here, but we also need to allow Γ to be disconnected in general).

The space $\mathcal{O}(\Gamma)$ is then the deformation space of this splitting. That is, the space of all the simplicial, edge-free trees on which Γ acts with the same elliptic elements as the defining splitting. We endow these trees with Γ -metrics, and identify two trees when there is an equivariant isometry between them (or an equivariant homothety if we don't insist on having volume 1). The space $\overline{\mathcal{O}(\Gamma)}$ is the free splitting complex associated to this; the space of all edge-free, simplicial metric trees whose elliptic elements contain all the elliptic elements from the original splitting (but may contain more). Again, trees are identified if they are equivariantly isometric.

More concretely, we work almost entirely with the natural simplicial structure on $\mathcal{O}(\Gamma)$, since each point is given by a graph of groups and induces a cone on an open simplex by varying the lengths of edges. Technically, varying the lengths of edges in the graph of groups produces a simplex if we impose the condition that the volume is one, or if we work projectively; these correspond to the spaces $\mathcal{O}_1(\Gamma)$ and $\mathbb{P}\mathcal{O}(\Gamma)$. If the number of edges in the graph of groups is k , this produces an open $k - 1$ simplex. In $\mathcal{O}(\Gamma)$ we get a cone on this open $k - 1$ simplex. (See 2.14 and 2.25.) Since the cone on an open simplex is homeomorphic to a open simplex of one dimension higher, we shall abuse notation and simply refer to simplicies in $\mathcal{O}(\Gamma)$.

One can then view $\overline{\mathcal{O}(\Gamma)}$ as the simplicial closure of $\mathcal{O}(\Gamma)$; a simplex in $\mathcal{O}(\Gamma)$ has faces which correspond to collapsing various subgraphs. When such a subgraph carries a hyperbolic element, the resulting quotient object defines a tree (graph of groups) with more elliptic elements, and hence a point of $\overline{\mathcal{O}(\Gamma)}$. All points of $\overline{\mathcal{O}(\Gamma)}$ arise in this way. We use the notation $\partial\mathcal{O}(\Gamma)$ to denote the points in $\overline{\mathcal{O}(\Gamma)}$ which are not in $\mathcal{O}(\Gamma)$. Since points of $\partial\mathcal{O}(\Gamma)$ can be considered as *points at infinity* of $\mathcal{O}(\Gamma)$, we often use the notations

$$\partial_\infty\mathcal{O}(\Gamma) \text{ and } \overline{\mathcal{O}(\Gamma)}^\infty.$$

(See Definitions 2.29 and 2.30 for more details). Any such Γ has an associated group, $\text{Out}(\Gamma)$ of (outer) automorphisms of the group which preserve the elliptic elements, and this groups acts on $\mathcal{O}(\Gamma)$ by isometries with respect to the Lipschitz metric.

In the case that $\Gamma = F_n$; that is, the trivial splitting of the free group where every non-identity element is hyperbolic, we obtain $\mathcal{O}(\Gamma) = CV_n$ and $\overline{\mathcal{O}(\Gamma)} = \mathcal{FS}_n$, and the associated automorphism group is $\text{Out}(F_n)$.

Any $\phi \in \text{Out}(F_n)$ acts on $\mathcal{O}(F_n)$ and induces a displacement function $\lambda_\phi : \mathcal{O}(F_n) \rightarrow [1, \infty)$

$$\lambda_\phi(X) = \Lambda(X, \phi X)$$

where Λ denotes the (multiplicative, non symmetric) Lipschitz distance.

We extend this function to $\overline{\mathcal{O}(F_n)}^\infty$.

If $X \in \mathcal{O}(F_n)$ exhibits a ϕ -invariant sub-graph A , the collapse of A defines a point $X/A \in \partial_\infty\mathcal{O}(F_n)$, whose displacement is finite, since $\phi(X/A)$ is a well defined point, again in $\partial_\infty\mathcal{O}(F_n)$ (and carrying the same set of hyperbolic/elliptic elements). In other words, if we let Γ denote the induced free splitting of F_n arising from the collapse of A , then both X/A and $\phi(X/A)$ belong to $\mathcal{O}(\Gamma)$, and hence are at finite Lipschitz distance.

By setting $\lambda_\phi(X/A) = \infty$ for those points $X/A \in \partial_\infty \mathcal{O}(F_n)$ whose collapsed part is not ϕ invariant, we have λ_ϕ defined on the whole $\overline{\mathcal{O}(F_n)}^\infty$ (although some points have infinite displacement).

The same process works for any Γ in place of F_n and we study these all at the same time. The advantage of this is that we can apply inductive arguments, which turn out to be key in understanding the properties of λ_ϕ . In particular, we prove that minimally displaced points are characterised in terms of (partial) train-track maps, and any automorphism has a minpoint in $\overline{\mathcal{O}(\Gamma)}^\infty$, though not necessarily in $\mathcal{O}(\Gamma)$ - Theorem 4.15 and Theorem 7.11.

We also study the (failure of the) continuity of the function λ_ϕ on $\overline{\mathcal{O}(\Gamma)}^\infty$ and characterize the points at which it is not continuous - the ‘jumping’ points, Theorem 5.14 and Theorem 7.8. We describe some of these results in more detail below.

1.2. Anticipating the results. The main tool for studying λ_ϕ is to use good representatives for ϕ . Namely, given $X \in \mathcal{O}(F_n)$ (or in any $\mathcal{O}(\Gamma)$), we need to find the best Lipschitz maps representing ϕ (that is to say $f : X \rightarrow X$ so that $f_* = \phi$ on $\pi_1(X)$). All maps we use will be *straight*, meaning that have constant speed on edges (hence they are determined by the image of vertices). It is classical that one may always find an optimal map $f : X \rightarrow \phi X$, whose Lipschitz constant satisfies

$$\text{Lip}(f) = \Lambda(X, \phi X).$$

However the usual proof, by means of Ascoli-Arzelà, is not constructive, nor quantitative. Our first result is Theorem 3.15 which can be stated as follows, and gives a constructive procedure - via a flow - for making a straight map optimal, and crucially adds a quantitative bound to the process.

Theorem (Optimization). Given $X, Y \in \mathcal{O}(F_n)$ and $f : X \rightarrow Y$ a Lipschitz map, there is a Lipschitz map $g : X \rightarrow Y$ so that $\text{Lip}(g) = \Lambda(X, Y)$ and so that $d_\infty(g, f) \leq \text{vol}(X)(\text{Lip}(f) - \Lambda(X, Y))$.

The estimate arising from this theorem will be crucial in many proofs. For any *straight* map $f : X \rightarrow Y$, the *tension graph* of f , denoted by X_{\max} , is the sub-graph of X whose edges are maximally stretched. We introduce the notion of *partial train track* map as a straight map $f : X \rightarrow \phi X$ such that there is an invariant sub-graph $A \subseteq X_{\max}$ (not necessarily proper) so that the restriction of f to A is a train track map in the usual sense. Our study of displacement functions is based on the use of partial train tracks. The first result on partial train tracks is that they characterise minimally displaced points (See Theorem 4.15 for a precise statement):

Theorem. For any automorphism ϕ , local minima for λ_ϕ are global minima and consist exactly of those points supporting a partial train track.

One of the main problems is that λ_ϕ is not continuous at the boundary points of $\mathcal{O}(F_n)$. We say that $X \in \partial_\infty \mathcal{O}(F_n)$ *has not jumped* if there is a sequence $X_i \rightarrow X$ of points $X_i \in \mathcal{O}(F_n)$, all contained in a single simplex, such that $\lambda_\phi(X_i) \rightarrow \lambda_\phi(X)$ (precise definitions are

given in Definition 5.13). In Sections 5 and 7 we give a complete description of jumping and non-jumping points. For instance, if we set $\lambda(\phi) = \inf_{X \in \mathcal{O}(F_n)} \lambda_\phi(X)$ we get (Theorem 7.8, see also Theorem 7.4 for related statements)

Theorem. $X \in \partial_\infty \mathcal{O}(F_n)$ has not jumped if and only if $\lambda_\phi(X) \geq \lambda(\phi)$.

In particular,

Theorem (Corollary 7.9). For any automorphism ϕ , if $X \in \partial_\infty \mathcal{O}(F_n)$ is a min-point for λ_ϕ (i.e. satisfies $\lambda_\phi(X) = \lambda(\phi)$) then X has not jumped.

For any ϕ we give the notion of *partial train track at infinity* as points $X \in \partial_\infty \mathcal{O}(F_n)$ which have not jumped **and** are partial train tracks for the induced automorphism in the deformation space of X . In Section 7 we prove that partial train tracks at infinity exist and are min-points. We prove in particular the existence of (non-jumping) min-points in the bordification of outer space.

Theorem (Theorem 7.11). Any automorphism has a partial train track in $\overline{\mathcal{O}(F_n)}^\infty$. Partial train tracks (at infinity or otherwise) are min-points for the displacement function. In particular, (non-jumping) min-points always exist.

The existence of partial train tracks also give information on invariant free factors:

Theorem (Theorems 7.13). For any automorphism ϕ , any ϕ -invariant free factor of F_n is visible in some partial train track.

And as in the irreducible case, existence of partial train tracks allows one to easily deduce that for any automorphism we have $\lambda(\phi^n) = \lambda(\phi)^n$. (Corollary 7.14.)

Remark (Connection with Relative Train Track Maps). There is a connection between relative train track maps and partial train track maps as follows: given the automorphism, ϕ , one constructs a relative train track map as in [5]. Suppose that λ is the maximum Perron-Frobenius eigenvalue for any stratum, and that the highest stratum in which it occurs is the r^{th} one. (That is, that the Perron-Frobenius eigenvalue is strictly greater than that of any higher stratum, and at least as great as that of any lower stratum). Now collapse the invariant subgraph G_{r-1} - the union of all the strata below the r^{th} stratum. This defines a point of the free splitting complex, where ϕ admits a representative supporting an invariant subgraph on which it is train track with expansion factor, λ . By making the volume of this subgraph sufficiently small, we can ensure that the Lipschitz constant of every other edge is strictly less than λ , and this is our partial train track at infinity. It then follows that $\lambda = \lambda(\phi)$.

However, the important difference between the two objects is that partial train tracks characterise exactly the minimally displaced set, Theorem 4.15, whereas relative train tracks do not.

The objects with which we work are usually not locally compact. This makes all convergence arguments technically difficult. For controlling the convergence and minimisation processes, in particular those of Section 7, we make crucial use of the following result on displacements. For any simplex of $\overline{\mathcal{O}(F_n)}^\infty$ define $\lambda_\phi(\Delta) = \inf_{X \in \Delta} \lambda_\phi(X)$. Then we prove;

Theorem(Theorem 7.2). For any F_n (and in fact for any deformation space) the global simplex-displacement spectrum

$$\text{spec}(F_n) = \left\{ \lambda_\phi(\Delta) : [\phi] \in \text{Out}(F_n), \Delta \text{ a simplex of } \overline{\mathcal{O}(F_n)}^\infty \text{ s.t. } \lambda_\phi(\Delta) < +\infty \right\}$$

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

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

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

Finally, we want to mention also Section 6, in which we give a detailed description of useful convexity properties of displacement functions, for instance proving that the displacement function is quasi-convex along Euclidean segments - see Lemma 6.2 and Lemma 6.3.

ACKNOWLEDGEMENTS: We would like to thank both the Università di Bologna and the Universitat Politècnica de Catalunya, for their hospitality during several visits.

We would also like to thank the referee of the earlier version of this paper (when it was a single paper together with [10]) for many very helpful comments, as well of the referee of the present paper for the very useful comments and suggestions.

2. SETTING, NOTATION, AND GENERAL DEFINITIONS

2.1. Motivation for new definitions. First, we wish to motivate our definitions and the general setting. Our aim is to study automorphisms of free groups which are possibly reducible. (Although our results will apply to free products more generally). If Γ is a marked graph with $\pi_1(\Gamma) = F$ a free group, and $\phi \in \text{Out}(F)$, then ϕ can be represented by a simplicial map (that is, a continuous map on the graph, sending vertices to vertices and edges to edge paths) $f : \Gamma \rightarrow \Gamma$. That is, f represents ϕ if there is an isomorphism $\tau : F_n \rightarrow \pi_1(\Gamma)$ such that $\phi = \tau^{-1} f_* \tau$. (The reason we are working with outer automorphisms is that we do not keep track of basepoints).

If ϕ is reducible, then it is possible that we may find a collection of disjoint connected sub-graphs $\Gamma_1, \dots, \Gamma_k$ such that f permutes the Γ_i 's. (We are guaranteed to find such a collection in *some* Γ). In order to study the properties of ϕ it may help to collapse such an invariant collection. (In other words, in the study of reducible automorphisms, we are naturally led to study the simplicial bordification of the Culler-Vogtmann Outer space CV_n .)

If we want to keep track of all the relevant information, we will be faced with the study of some particular kind of moduli spaces. Namely, moduli spaces of actions on trees with

possibly non-trivial vertex stabilizers (when we collapse the Γ_i 's) and the product of such spaces (when we consider the restriction to ϕ to the Γ_i 's.)

The typical topological object we are concerned with is a disjoint union of metric trees, where G acts with finitely many orbits, but with possibly non-trivial vertex-stabilizers. Therefore we will develop the paper in this general - free splitting - setting, but the reader is invited to restrict attention to the case of CV_n and its bordification.

2.2. Notation for free splittings. Let $G = G_1 * \cdots * G_p * F_n$ be any free product of groups, where F_n denotes the free group of rank n (we allow n to be zero, in that case we omit F_n). We do not assume that the G_i 's are indecomposable. The theory we are going to develop is general, but we are mainly interested in the case where G is itself a free group. (Thus, in general, this free product decomposition is not unique, since G has many different splittings as a free product.)

Definition 2.1 (Free Splittings). Given a group G , a *free splitting* \mathcal{G} of G is a pair $(\{G_i\}, n)$ where $\{G_i\}$ is a collection of subgroups of G and n is natural a number such that $G = G_1 * \cdots * G_p * F_n$. Two splittings $(\{G_1, \dots, G_p\}, n)$ and $(\{H_1, \dots, H_p\}, m)$ of G are considered to be *of the same type* if $m = n$ and, up to reordering factors, each H_i is conjugate to G_i .

Remark 2.2. 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 2.3. Free splittings are also referred to as *free factor systems* in the literature - originally introduced in [4], and also used in [15], [16] and [18].

Our point of view here is to take a fixed free factor system - a free splitting - and form the deformation space of that. This consists of trees equipped with edge-free actions whose vertex stabilisers are the (conjugacy classes of) the elements of the free factor system. That is, one can form the space of all trees which give rise to the same free factor system.

One can also form the what is known as the free splitting graph or complex, which consists of all possible free splittings (and one can also make this relative to a base free splitting). This relative version is what we have in mind when we come to define our *simplicial bordification* (see 2.28).

Note 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 2.4 (Sub-splittings). Let $\mathcal{G} = (\{G_1, \dots, G_p\}, n)$ and $\mathcal{S} = (\{H_1, \dots, H_q\}, r)$ be two free splittings of G . We say that \mathcal{S} is a *sub-splitting* of \mathcal{G} if each H_i decomposes as

$$H_i = G_{i_1} * \dots * G_{i_l} * F_{s_i},$$

and $r + \sum_i s_i = n$.

Definition 2.5 (Kurosh rank of a free splitting). The Kurosh rank of the splitting $G = G_1 * \cdots * G_p * F_n$ is $n + p$.

2.3. G -graphs and G -trees. Given a group G , a simplicial G -tree is a simplicial tree T endowed with a faithful simplicial action of G . T is minimal if it has no proper G -invariant sub-tree. In particular, if T is minimal then G acts without global fixed points and T has no leaves (valence one vertices).

We next define \mathcal{G} -trees and \mathcal{G} -graphs. For those familiar with Bass-Serre theory, these are the trees dual to a given splitting and the corresponding graphs of groups.

Definition 2.6 (\mathcal{G} -trees and \mathcal{G} -graphs). Let G be a group G and $\mathcal{G} = (\{G_1, \dots, G_p\}, n)$ be a splitting of G . A \mathcal{G} -tree is a metric simplicial G -tree T such that

- G acts isometrically on T ,
- For every G_i there is exactly one orbit of vertices whose stabilizer is conjugate to G_i . Such vertices are called *non-free*. Remaining vertices have trivial stabilizer and are called *free* vertices.
- T has trivial edge stabilizers.

A \mathcal{G} -graph is a finite connected metric graph of groups X such that

- X is *marked*; that is, there is a fixed isomorphism between the fundamental group of X (as a graph of groups) and G .
- X has trivial edge-groups;
- the fundamental group of X as a topological space is F_n ;
- the splitting given by the vertex groups is equivalent to \mathcal{G} .

We note that Bass-Serre theory gives a correspondence between \mathcal{G} -trees and \mathcal{G} -graphs.

Two \mathcal{G} -trees are equivalent if there is an equivariant isometry between them; there is an analogous equivalence at the level of \mathcal{G} -graphs, which is harder to state but comes down to what is called a graph of groups morphism, and arises as the quotient map one gets from an equivariant isometry.

Remark 2.7. Recall that for an action on a (simplicial) tree, every group element either fixes a point or has an axis of minimal displacement. In the former case the element is called elliptic, and in the latter case hyperbolic.

Notation 2.8. Throughout the paper, if G has a splitting \mathcal{G} which is clear from the context, then any G -tree is required to be a \mathcal{G} -tree. (And the same for graphs.)

Example 2.9. If X is a finite connected graph of groups with trivial edge-groups, then denote the splitting induced by the vertex groups of X by \mathcal{G} . It is clear that \mathcal{G} is a splitting for $\pi_1(X)$, X is a \mathcal{G} -graph, and the Bass-Serre tree associated to X is a \mathcal{G} -tree.

Definition 2.10 (Core graph). A *core-graph* is a graph of groups whose leaves have non-trivial vertex-group. Given a graph X we define $\text{core}(X)$ to be the maximal core sub-graph of X . (If the vertex groups are all trivial, so that X is simply a graph, then a core graph has no valence one vertices). Note that $\text{core}(X)$ is obtained by recursively cutting edges ending at leaves.

Given a splitting $\mathcal{G} = (\{G_i\}, n)$ of a group G and T a \mathcal{G} -tree, the quotient $X = G \backslash T$ is a connected \mathcal{G} -graph. T is minimal if and only if X is a finite core graph. Since in

the paper we are dealing with both \mathcal{G} -graphs and \mathcal{G} -trees, we introduce what we call the tilde-underbar notation.

Notation 2.11 (Tilde-underbar notation). Let \mathcal{G} be a splitting of a group G . If X is a \mathcal{G} -graph, then \tilde{X} denotes its universal covering, which is a \mathcal{G} -tree. As usual, if $x \in X$ then \tilde{x} will denote a lift of x in \tilde{X} .

We will also often want to lift “loops”; that is, given an element of $\pi_1(X)$ - the fundamental group of X as a graph of groups - we lift the loop to a line in the universal cover. Concretely, this requires a description of the loop as a sequence of edges and vertex group elements which can be mirrored in the Bass-Serre tree; note that edges “downstairs” are orbits of edges, so the vertex group element serves to determine which edge “upstairs” in the orbit to follow. Similarly, we may lift paths, using the same process.

Equivalently, a lift of a group element is the axis, and the collection of lifts is the orbit of this axis (if one forgets basepoints, then there is no favourite lift, and one just has the orbit, of which one chooses a component.)

For subsets, $A \subseteq X$, we will usually mean a collection of edges and vertex groups - that is, the entire vertex group is a part of A if at all. We can lift A by taking all the lifts of the loops realised in A , and if A is connected we usually refer to a component of this as a “lift” of A .

Conversely, if T is a minimal \mathcal{G} -tree we denote by \underline{T} the quotient \mathcal{G} -graph. We mirror this notation for points and subsets.

Hence, $\tilde{\underline{X}} = X$ for both graphs and trees.

Definition 2.12 (X -graphs, trees and forests). Let \mathcal{G} be a splitting of a group G .

- If X is a \mathcal{G} -graph (resp. \mathcal{G} -tree), then a X -graph (resp. X -tree) is just a \mathcal{G} -graph (resp. tree). Unless otherwise specified, given a finite connected graph of groups X with trivial edge-groups, an X -graph is a \mathcal{G} -graph (and an X -tree is a \mathcal{G} -tree).
- 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).
- A Γ -sub-forest \tilde{A} of a Γ -forest \tilde{X} is the lift of a sub-graph of a Γ -graph \underline{X} . Here sub-graphs and lifts are “full” in the sense that if a sub-graph contains a vertex, it contains the entire vertex group, as a sub-graph of groups — and “lift” means the full pre-image under the map from $\tilde{X} \rightarrow \underline{X}$.
- A sub-graph A of a Γ -graph is **non-trivial** if the fundamental group (as a graph-of groups) of any of its components contains a hyperbolic element. That is, if it contains a non-trivial element which is not just a vertex group element. A non-trivial Γ -sub-forest of a Γ -forest is a sub-forest obtained as a lift of a non-trivial sub-graph.

Definition 2.13 (Immersed loops). A path γ in a \mathcal{G} -graph X is called *immersed* if it has a lift $\tilde{\gamma}$ in \tilde{X} which is embedded. (Note that γ might not be topologically immersed in X near non-free vertices.)

2.4. Outer spaces. We briefly recall the definition of the outer space of a group G corresponding to a splitting \mathcal{G} , referring to [9, 14] for a detailed discussion of definitions and general properties.

Definition 2.14 (Outer space). Let G be a group and \mathcal{G} be a splitting of G . The (projectivized) outer space of G , relative to the splitting $\mathcal{G} = (\{G_1, \dots, G_p\}, n)$, consists of (projective) classes of minimal, simplicial, metric \mathcal{G} -trees, X with no redundant vertex (that is, no valence two vertex is allowed to be free) and such that the G -action is by isometries.

We use the notation $\mathcal{O}(G; \mathcal{G})$ or simply $\mathcal{O}(\mathcal{G})$ to denote 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 denote the projectivized outer space.

For $X \in \mathcal{O}(\mathcal{G})$ we define its (co-)volume $\text{vol}(X)$ as the sum of lengths of edges in $G \backslash X$. On occasion, we will need to work with the co-volume one slice of $\mathcal{O}(\mathcal{G})$, which we denote by $\mathcal{O}_1(\mathcal{G})$.

Remark 2.15. If \mathcal{G} is the trivial splitting of $G = F_n$, then $\mathcal{O}(\mathcal{G}) = CV_n$.

We stress here 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.

Remark 2.16. The equivalence relation that defines $\mathbb{P}\mathcal{O}(\mathcal{G})$ is the following: X and Y are equivalent if there is an homothety (isometry plus a rescaling by a positive number) $X \rightarrow Y$ conjugating the actions of G on X and Y . In particular, since G acts isometrically on a metric \mathcal{G} -tree, the inner automorphisms of G act trivially on $\mathcal{O}(\mathcal{G})$ and $\mathbb{P}\mathcal{O}(\mathcal{G})$.

Remark 2.17. If G has a the simple splitting $G = G_1$, then $\mathcal{O}(\mathcal{G})$ consists of a single element: a point stabilized by G_1 , and in this case the equivalence relation is trivial.

Remark 2.18. If $X \in \mathcal{O}(\mathcal{G})$, the quotient \underline{X} is a metric core \mathcal{G} -graph. Conversely, if X is a core metric \mathcal{G} -graph with no redundant vertex, then $\tilde{X} \in \mathcal{O}(\mathcal{G})$.

In the paper we will work with both graphs and trees. Strictly speaking we have defined $\mathcal{O}(\mathcal{G})$ as a space of trees, but it will be often convenient to use graphs X so that $\tilde{X} \in \mathcal{O}(\mathcal{G})$. Clearly the two viewpoints are equivalent and we shall have occasion to abuse notation and switch between graphs and trees. However, when we wish to make the distinction clear, we will add a “*gr*” subscript to indicate that we are working with graphs. To illustrate:

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

The spaces $\mathcal{O}_{\text{gr}}(\mathcal{G})$ and $\mathcal{O}(\mathcal{G})$ are naturally identified via $\underline{X} \leftrightarrow \tilde{X}$.

Notation 2.19. If \underline{X} is a finite connected graph of groups with trivial edge-groups, and \mathcal{S} is the splitting of $\pi_1(\underline{X})$ given by vertex-groups, then we set

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

Clearly, if \underline{X} is a metric core graph of groups with no redundant vertices, then $\tilde{X} \in \mathcal{O}(\underline{X})$.

Definition 2.20. We define Γ and $\mathcal{O}(\Gamma)$ as follows:

- $\Gamma = \sqcup \Gamma_i$ will always mean that Γ is a finite disjoint union of *connected* finite graphs of groups Γ_i , each with trivial edge-groups and non-trivial fundamental group $H_i = \pi_1(\Gamma_i)$, each H_i being equipped with the splitting given by the vertex-groups.
- Then $\mathcal{O}(\Gamma)$ is defined to be the product of the $\mathcal{O}(\Gamma_i)$. That is, a point in $\mathcal{O}(\Gamma)$ is a tuple of minimal, simplicial, edge-free isometric actions of the corresponding H_i on metric trees, up to equivariant isometry. The elliptic elements are precisely the vertex groups in Γ .
- There is a natural action of \mathbb{R}^+ on $\mathcal{O}(\Gamma)$ given by scaling each component by the same amount. The quotient of $\mathcal{O}(\Gamma)$ by such action 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 co-volume extends to Γ -trees: If $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$ we set $\text{vol}(X) = \sum_i \text{vol}(X_i)$, and $\mathcal{O}_1(\Gamma)$ denotes the co-volume 1 slice of $\mathcal{O}(\Gamma)$.
- We tacitly identify $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$ 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 (See Definition 2.12).

Remark 2.21. If X is a \mathcal{G} -tree, then $\mathcal{O}(X) = \mathcal{O}(\mathcal{G})$. In other words, $\mathcal{O}(\mathcal{G})$ is a particular case of $\mathcal{O}(\Gamma)$ with connected Γ . In the following we will therefore develop the theory for general $\mathcal{O}(\Gamma)$, as this includes the “connected” case $\mathcal{O}(\mathcal{G})$ (and in particular, the CV_n case).

Definition 2.22 (Rank). The Kurosh rank of a finite graph of groups with trivial edge-groups is the Kurosh rank of the splitting¹ induced on its fundamental group by the vertex-groups. If $\Gamma = \sqcup \Gamma_i$ we set

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

By definition, the rank is a natural number greater or equal to one. Note the the rank of a graph of groups X is simply the rank of its fundamental group as a topological space plus the number of non-free vertices.

We will also consider moduli spaces with marked points.

Notation 2.23. Let \mathcal{G} be a splitting of G . The moduli space of \mathcal{G} -trees with k labelled points p_1, \dots, p_k (not necessarily distinct) is denoted by $\mathcal{O}(G; \mathcal{G}, k)$ or simply $\mathcal{O}(\mathcal{G}, k)$. If X is a finite graph of groups with trivial edge-groups we set $\mathcal{O}(X, k) = \mathcal{O}(\pi_1(X), k)$. If $\Gamma = \sqcup_{i=1}^s \Gamma_i$, given $k_1, \dots, k_s \in \mathbb{N}$ we set

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

Notation 2.24. Let X be a finite, connected graph of groups with trivial edge-groups and let \mathcal{G} be the splitting of $\pi_1(X)$ given by vertex groups. Let $A \subseteq X$ be a non-trivial sub-graph (Definition 2.12). Then A induces a sub-splitting \mathcal{S}_A of \mathcal{G} where the factor groups are either the fundamental groups of the components of A , or vertex groups in $X \setminus A$.

¹See Definition 2.5.

- We denote by X/A the graph of groups where components of A are collapsed to points (a point for each component), and a vertex group is inserted - the fundamental group of the collapsed component.
- We denote by $\mathcal{O}(X/A)$ the outer space $\mathcal{O}(\pi_1(X), \mathcal{S}_A)$.
- We extend this definition to the case of $\Gamma = \sqcup_i \Gamma_i$ as in Definition 2.20 and $A \subseteq \Gamma$, by considering the splittings induced by A on any component Γ_i , and we use notations Γ/A and $\mathcal{O}(\Gamma/A)$ for the resulting (disjoint union of) graphs and its outer space.
- We borrow the same notation when speaking of Γ -forests (or \mathcal{G} -forests) and non-trivial sub-forests.

2.5. Simplicial structure. The simplicial structure we are going to use is the one familiar to experts - see [6] and [14]. Since we want to study the simplicial bordification of our outer spaces, we need to introduce faces “at infinity” and a suitable notation for distinguish them from usual finitary faces. Faces “at infinity” of $\mathcal{O}(\mathcal{G})$, will be in fact simplices in the outer space of some sub-splitting of \mathcal{G} .

Definition 2.25 (Open simplices). Given a \mathcal{G} -tree X , the open simplex Δ_X is the set of \mathcal{G} -trees equivariantly homeomorphic to X . If X is a \mathcal{G} -graph, then we agree that Δ_X is the set of graphs obtained by quotients of elements of $\Delta_{\tilde{X}}$. The Euclidean topology on Δ_X is given by assigning a G -invariant positive length $L_X(e)$ to each edge e of X . Therefore, if X has k orbits of edges, then Δ_X is isomorphic to the standard open $(k - 1)$ -simplex if we work in $\mathbb{P}\mathcal{O}(\mathcal{G})$ or $\mathcal{O}_1(\mathcal{G})$, and to the positive cone over it if we work on $\mathcal{O}(\mathcal{G})$.

Since the positive cone of a $k - 1$ -simplex is homeomorphic to an open k -simplex, we shall abuse notation and simply refer to simplices in $\mathcal{O}(\Gamma)$.

Given two elements X, Y in the same simplex $\Delta \subset \mathcal{O}(\mathcal{G})$ we define the **Euclidean** sup-distance $d_{\Delta}^{Euclid}(X, Y)$ ($d_{\Delta}(X, Y)$ for short)

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

Such definitions extend to the case of $\Gamma = \sqcup_i \Gamma_i$. We refer the Definition 2.20.

Definition 2.26 (Euclidean topology). If $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$, the simplex Δ_X is the set of Γ -forests equivariantly homeomorphic to X (component by component). The Euclidean topology and distance on Δ_X are defined by

$$d_{\Delta}(X, Y) = \sup_i d_{\Delta_{X_i}}(X_i, Y_i).$$

We note that the simplicial structure of $\mathbb{P}\mathcal{O}(\Gamma)$ is not the product of the structures of $\mathbb{P}\mathcal{O}(\pi_1(\Gamma_i))$.

Definition 2.27 (Faces and closed simplices). Let X be a Γ -graph and let $\Delta = \Delta_X$ be the corresponding open simplex. Let $F \subset X$ be a forest whose trees each contains at most one non-free vertex. The collapse of F in X produces a new Γ -graph, whence a simplex Δ_F . Such a simplex is called a *face* of Δ .

The *closed* simplex $\overline{\Delta}$ is defined by

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

2.6. Simplicial bordification. There are two natural topologies on $\mathcal{O}(\Gamma)$, the simplicial one and the equivariant Gromov topology, which are in general different. Here we will mainly use the simplicial topology. We notice that if Δ is an open simplex, then the simplex $\overline{\Delta}$ is not the standard simplicial closure of Δ , because not all its simplicial faces are *faces* according to Definition 2.27. This is because some simplicial faces of Δ are not in $\mathcal{O}(\Gamma)$ as defined. Such faces are somehow “at infinity” and describe limit points of sequences in $\mathcal{O}(\Gamma)$. We now give precise definitions to deal with these limit points.

It will be convenient to start with describing the free splitting complex, relative to Γ .

Definition 2.28 (Free Splitting Complex). Let Γ be as in Definition 2.20. That is, Γ is a finite disjoint union of graphs of groups. Then the free splitting complex relative to Γ is the set of tuples of minimal, simplicial, isometric, edge-free, actions on H_i -trees, where the set of elliptic elements *includes* all the vertex groups of Γ , and up to equivariant isometry.

We denote this set, $\overline{\mathcal{O}(\Gamma)}$.

As in, Definition 2.25, by varying the edge lengths on a given (tuple of) trees we can produce an simplex in $\overline{\mathcal{O}(\Gamma)}$. (Recall that we are abusing notation and calling the positive cone on an open simplex, an open simplex).

Definition 2.29 (Simplices and Faces). Let Δ be an open simplex in $\mathcal{O}(\Gamma)$.

- $\overline{\Delta}$ is the closure of Δ in $\mathcal{O}(\Gamma)$, as in Definition 2.27,
- $\partial_{\mathcal{O}}\Delta = \partial_{\mathcal{O}}\overline{\Delta}$ is the set-difference, $\overline{\Delta} \setminus \Delta$. We call this the *finitary boundary* of Δ . The *finitary faces* of Δ are the faces which appear in $\partial_{\mathcal{O}}\Delta$.
- $\overline{\Delta}^{\infty}$ is the closure of Δ in $\overline{\mathcal{O}(\Gamma)}$,
- $\partial_{\infty}\Delta = \partial_{\infty}\overline{\Delta}$ is the set-difference, $\overline{\Delta}^{\infty} \setminus \overline{\Delta}$.

Let X be a Γ -graph and $\Delta = \Delta_X$. Let A be a proper subgraph of X having at least a component which is not a tree with at most one non-free vertex. Equivalently, A contains the axis of a hyperbolic element.

Let Y be the graph of groups obtained by collapsing each component of A to a point (different components to different points). Then, $Y \in \mathcal{O}(X/A)$. The corresponding simplex Δ_Y is a simplicial face of Δ_X obtained by setting the edge-lengths of A to zero. Note that Δ_Y belongs to $\mathcal{O}(X/A)$ and not to $\mathcal{O}(X)$. However, the simplicial topology naturally defines a topology on $\Delta_X \cup \Delta_Y$, which we still name simplicial topology.

Also note that any point in $\partial_{\infty}\Delta$ is obtained by collapsing such a sub-graph, A . We may, in general, assume that all the components of A are core-graphs (each component contains the axis of a hyperbolic element), if we are willing to replace Δ with a finitary face of Δ .

Definition 2.30 (Boundary at infinity). We define the boundary at infinity and the simplicial bordification of $\mathcal{O}(\Gamma)$ as

$$\partial_{\infty}\mathcal{O}(\Gamma) = \bigcup_{\Delta \text{ simplex}} \partial_{\infty}\Delta.$$

Note that,

$$\overline{\mathcal{O}(\Gamma)}^{\infty} := \overline{\mathcal{O}(\Gamma)} = \mathcal{O}(\Gamma) \cup \partial_{\infty}\mathcal{O}(\Gamma).$$

Remark 2.31. Note that all these operations can be carried out in the projective spaces, with the definitions essentially unchanged.

Remark 2.32. We note that when $\Gamma = F_n$, that is the splitting of the free group where every non-trivial element is hyperbolic, then we get that $\mathcal{O}(\Gamma)$ is simply Culler-Vogtmann space, CV_n and the bordification, $\overline{\mathcal{O}(\Gamma)}$ is the free splitting complex, \mathcal{FS}_n .

2.7. Horoballs and regeneration. We recall Definition 2.20.

Definition 2.33 (Horoballs). Given $X \in \partial_\infty \mathcal{O}(\Gamma)$, $\text{Hor}(X)$ 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. We set $\text{Hor}(X) = X$ for $X \in \mathcal{O}(\Gamma)$, by convention (and use $\text{Hor}(\underline{X})$ for graphs).

In other words, a metric graph \underline{Y} is in the horoball of \underline{X} if \underline{X} is obtained from \underline{Y} by setting to zero the length of edges of a proper family of core sub-graphs. On the other hand, $\text{Hor}(X)$ can be regenerated from X as follows.

Suppose $X \in \partial_\infty \mathcal{O}(\Gamma)$. Thus there is a Γ -graph \underline{Y} and a sub-graph $\underline{A} = \sqcup_i \underline{A}_i \subset Y$ whose components \underline{A}_i are core-graphs, and such that $\underline{X} = \underline{Y}/\underline{A}$. Let v_i be the non-free vertex of \underline{X} corresponding to \underline{A}_i . In order to recover a generic point $Z \in \text{Hor}(X)$, we need to replace each v_i with an element $\underline{V}_i \in \mathcal{O}(\underline{A}_i)$. Moreover, in order to completely define the marking on \underline{Z} , we need to know where to attach to \underline{V}_i the edges of \underline{X} incident to v_i , and this choice has to be done in the universal covers \tilde{V}_i . No more is needed. Therefore, if k_i denotes the valence of the vertex v_i in X , we have

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

(Note that some k_i could be zero, e.g. if A_i is a connected component of Y .) There is a natural projection $\text{Hor}(X) \rightarrow \mathcal{O}(A)$ which forgets the marked points. We will be mainly interested in cases when we collapse A uniformly, for that reason we will use the projection to $\mathbb{P}\mathcal{O}(A)$:

$$\pi : \text{Hor}(X) \rightarrow \mathbb{P}\mathcal{O}(A)$$

where $\text{Hor}(X)$ is intended to be not projectivized.

Note that if $[P] \in \mathbb{P}\mathcal{O}(A)$, then $\pi^{-1}(P)$ is connected because it is just $\Pi_i (A_i^{k_i})$. Since $\mathcal{O}(A)$ is connected (as a product of connected spaces), then $\text{Hor}(X)$ is connected.

Remark 2.34. Note that *the same* graph of groups X can be considered as a point at infinity of *different* spaces. If we need to specify the space in which we work, we shall write $\text{Hor}_\Gamma(X)$ (or $\text{Hor}_G(X)$.)

2.8. The groups $\text{Aut}(\Gamma)$ and $\text{Out}(\Gamma)$. We are going to introduce the groups of automorphisms that preserve splittings, and their generalizations to the case of non-connected graphs.

Definition 2.35 (Automorphism-groups of splittings). Let G be endowed with the splitting $\mathcal{G} : G = G_1 * \cdots * G_p * F_n$. The group of automorphisms of G that preserve the set of conjugacy classes of the G_i 's is denoted by $\text{Aut}(G; \mathcal{G})$. We set $\text{Out}(G; \mathcal{G}) = \text{Aut}(G; \mathcal{G}) / \text{Inn}(G)^2$.

²Clearly $\text{Inn}(G) \subset \text{Aut}(G; \mathcal{G})$.

The group $\text{Aut}(G, \mathcal{G})$ acts on $\mathcal{O}(G)$ by changing the marking (i.e. the action), and $\text{Inn}(G)$ acts trivially. Hence $\text{Out}(G; \mathcal{G})$ acts on $\mathcal{O}(G; \mathcal{G})$. If $X \in \mathcal{O}(G; \mathcal{G})$ and $\phi \in \text{Out}(G; \mathcal{G})$ then ϕX is the same metric tree as X , but the action is $(g, x) \rightarrow \phi(g)x$. The action is simplicial and continuous w.r.t. both simplicial and equivariant Gromov topologies.

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

Definition 2.36 (Splitting isomorphism-groups). 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 map each G_i to a conjugate of one of the H_i 's is denoted by $\text{Isom}(G, H; \mathcal{G}, \mathcal{H})$. If splittings are clear from the context we write simply $\text{Isom}(G, H)$.

Definition 2.37 ($\text{Aut}(\Gamma)$). For $\Gamma = \sqcup_{i=1}^k \Gamma_i$ as in Definition 2.20, we set

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

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

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

Remark 2.38. Not all permutations appear. For instance, if the groups H_i are pairwise non-isomorphic, then the only possible σ is the identity.

Definition 2.39 ($\text{Inn}(\Gamma)$ and $\text{Out}(\Gamma)$). We set:

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

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

Example 2.40. If X is a \mathcal{G} -graph and $f : X \rightarrow X$ is a homotopy equivalence which leaves invariant a core subgraph A , then $f|_A$ induces an element of $\text{Aut}(A)$, and its free homotopy class an element of $\text{Out}(A)$.

The group $\text{Out}(\Gamma)$ acts on $\mathcal{O}(\Gamma)$ as follows. If $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$, then each X_i is an H_i -tree. If $(\sigma, \phi_1, \dots, \phi_k) \in \text{Aut}(\Gamma)$ then $X_{\sigma(i)}$ becomes an H_i -tree via the precomposition of $\phi_i : H_i \rightarrow H_{\sigma(i)}$ with the $H_{\sigma(i)}$ -action. We denote such an H_i -tree by $\phi_i X_{\sigma(i)}$. With that notation we have $\phi(X_1, \dots, X_k) = (\phi_1 X_{\sigma(1)}, \dots, \phi_k X_{\sigma(k)})$. (We remark that despite the left-positional notation, this is a right-action.) Since $\text{Inn}(\Gamma)$ acts trivially on $\mathcal{O}(\Gamma)$, then the $\text{Aut}(\Gamma)$ -action descends to an $\text{Out}(\Gamma)$ -action.

3. STRAIGHT MAPS, GATE STRUCTURES, AND OPTIMAL MAPS.

In this section we describe the theory of maps between trees (or graphs) representing points in outer spaces. We will simultaneously deal with the ‘‘connected’’ case $\mathcal{O}(\mathcal{G})$ (for instance the classical CV_n) and the general case $\mathcal{O}(\Gamma)$.

In this section G, \mathcal{G} and Γ will be as in Definitions 2.6 and 2.20.

3.1. Straight maps. Now we will mainly work with trees.

Definition 3.1 (\mathcal{O} -maps in $\mathcal{O}(\mathcal{G})$). Let $X, Y \in \mathcal{O}(\mathcal{G})$. A map $f : X \rightarrow Y$ is called an \mathcal{O} -map if it is Lipschitz-continuous and G -equivariant. The Lipschitz constant of f is denoted by $\text{Lip}(f)$.

We recall that we tacitly identify $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$ with the labelled disjoint union $\sqcup_i X_i$. Hence, if $X, Y \in \mathcal{O}(\Gamma)$, a continuous map $f : X \rightarrow Y$ is a collection of continuous maps $f_i : X_i \rightarrow Y_{\sigma(i)}$, where $\sigma \in \mathfrak{S}_k$.

Definition 3.2 (\mathcal{O} -maps in $\mathcal{O}(\Gamma)$). Let $X = (X_1, \dots, X_k)$ and $Y = (Y_1, \dots, Y_k)$ be two elements of $\mathcal{O}(\Gamma)$. A map $f = (f_1, \dots, f_k) : X \rightarrow Y$ 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. Compare with Definition 4.1.)

Definition 3.3 (Straight maps³). Let X, Y be two metric trees. A Lipschitz-continuous map $f : X \rightarrow Y$ is *straight* if it has constant speed on edges, that is to say, for any edge e of X there is a non-negative number $\lambda_e(f)$ such that for any $a, b \in e$ we have $d_Y(f(a), f(b)) = \lambda_e(f)d_X(a, b)$. If $X, Y \in \mathcal{O}(\mathcal{G})$ then we require any straight map to be an \mathcal{O} -map. A straight map between elements of $\mathcal{O}(\Gamma)$ is an \mathcal{O} -map whose components are straight. If X, Y are metric graphs, we understand that $f : X \rightarrow Y$ is a straight map if its lift to the universal covers is straight.

Remark 3.4. \mathcal{O} -map always exist and the images of non-free vertices is determined a priori by equivariance (see [9]). For any \mathcal{O} -map f there is a unique straight map denoted by $\text{Str}(f)$, which is homotopic, relative to vertices, to f . We have $\text{Lip}(\text{Str}(f)) \leq \text{Lip}(f)$.

Definition 3.5 (λ_{\max} and tension graph). Let $f : X \rightarrow Y$ be a Str-map. We set

$$\lambda(f) = \lambda_{\max}(f) = \max_e \lambda_e(f) = \text{Lip}(f).$$

We define the *tension graph* of f by

$$X_{\max}(f) = \{e \text{ edge of } X : \lambda_e(f) = \lambda_{\max}\}.$$

If there are no ambiguities on the map, we write λ_{\max} instead of $\lambda_{\max}(f)$ and X_{\max} for $X_{\max}(f)$.

Definition 3.6 (Stretching factors). For $X, Y \in \mathcal{O}(\Gamma)$ we define

$$\Lambda(X, Y) = \min_{f: X \rightarrow Y \text{ } \mathcal{O}\text{-map}} \text{Lip}(f)$$

The theory of stretching factors is well-developed in the connected case (i.e. for CV_n or general free products), but one can readily see that connectedness of trees plays no role, and the theory extends without modifications to the non-connected case. In fact, Λ is well-defined, (see [7, 9] for details) and it satisfies the multiplicative triangular inequality:

$$\Lambda(X, Z) \leq \Lambda(X, Y)\Lambda(Y, Z)$$

³In previous papers of the authors, a straight map is called *PL*-map. As a referee pointed out, piecewise linearity is a well-established notion in literature, which is slightly different from our notion (we don't allow subdivisions). For that reason we decided to change our previous terminology.

It can be used to define a non-symmetric metric $d_R(X, Y) = \log(\Lambda(X, Y))$ and its symmetrized version $d_R(X, Y) + d_R(Y, X)$ (see [7, 8, 9] for details) which induces the Gromov topology. The group $\text{Out}(\Gamma)$ acts by isometries on $\mathcal{O}(\Gamma)$.

Moreover, there is an effective way to compute Λ , via the so-called ‘‘sausage-lemma’’ (see [7, Lemma 3.14],[8, Lemma 2.16] for the classical case, and [9, Theorem 9.10] for the case of trees with non-trivial vertex-groups). We briefly recall here how it works.

Let X, Y be metric Γ -graphs. Any non-elliptic element $\gamma \in \pi_1(\Gamma)$ (i.e. an element not in a vertex-group) is represented by an immersed loop γ_X in X and one γ_Y in Y (see Definition 2.13). The loop γ_X (or, rather, its lift to \tilde{X}) is usually called **axis** of γ in X (or in \tilde{X}) and corresponds to the points of minimal translation of γ in \tilde{X} . The lengths $L_X(\gamma_X)$ and $L_Y(\gamma_Y)$ are then the minimal translation lengths of the element γ acting on \tilde{X} and \tilde{Y} , respectively. (So $L_X(\gamma_X) = L_X(\gamma)$ and $L_Y(\gamma_Y) = L_Y(\gamma)$.) We can define the stretching factor of γ as $L_Y(\gamma)/L_X(\gamma)$. Then $\Lambda(\tilde{X}, \tilde{Y})$ is the minimum of the stretching factors of all non-elliptic elements. (Recall we are using the tilde-underbar notation 2.11.)

Theorem 3.7 (Sausage Lemma [9, Theorem 9.10]). *Let $X, Y \in \mathcal{O}_{gr}(\Gamma)$. The stretching factor $\Lambda(X, Y)$ is realized by a loop $\gamma \subset X$ having one of the following forms:*

- *Embedded simple loop O ;*
- *embedded ‘‘infinity’’-loop ∞ ;*
- *embedded barbel $O - O$;*
- *singly degenerate barbel $\bullet - O$;*
- *doubly degenerate barbel $\bullet - \bullet$.*

(the \bullet stands for a non-free vertex.) Such loops are usually named ‘‘candidates’’.

Remark 3.8. The stretching factor $\Lambda(X, Y)$ is defined on $\mathcal{O}(\Gamma)$ (or in the co-volume slice $\mathcal{O}_1(\Gamma)$) and not in $\mathbb{P}\mathcal{O}(\Gamma)$. However, we will mainly be interested in computing factors of type $\Lambda(X, \phi X)$ (for $[\phi] \in \text{Out}(\Gamma)$) and that factor is scale invariant.

Definition 3.9 (Gate structures). Let X be any graph. A *turn* is a pair of germs of edges incident to the same vertex.

A *gate* structure on X is an equivalence relation on germs of edges, generated by (and in fact equal to) some collection of turns. Equivalence classes of germs are called *gates*. A turn is *illegal* if the two germs are in the same gate, it is *legal* otherwise. An immersed path in X is legal if it has only legal turns.

We can also consider gate structures on trees, where we require that the gate structures are equivariant, and hence descend to a gate structure on the quotient graph.

If $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$ we require the equivalence relation to be H_i -invariant on each X_i .

Any straight map induces a gate structure as follows.

Definition 3.10 (Gate structure induced by f). Given $X, Y \in \mathcal{O}(\Gamma)$ and a straight map $f : X \rightarrow Y$, the gate structure induced by f , denoted by

$$\sim_f$$

is defined by declaring equivalent two germs that have the same non-degenerate f -image.

Remark 3.11 (See [9]). Let $X, Y \in \mathcal{O}(\Gamma)$ and be $f : X \rightarrow Y$ a straight map. If v is a non-free vertex of X and e is an edge incident to v , then e and ge are in different gates for any $id \neq g \in \text{Stab}(v)$. (If e is collapsed by f , then it is not equivalent to any other edge by definition.)

Definition 3.12 (Optimal maps). Given $X, Y \in \mathcal{O}(\Gamma)$, a map $f : X \rightarrow Y$ is *weakly optimal* if it is straight and $\lambda(f) = \Lambda(X, Y)$.

A map $f : X \rightarrow Y$ is *optimal* if each vertex of X_{\max} has at least two gates in X_{\max} with respect to the gate structure induced by f .

Proposition 3.13. *A straight map between two Γ -forests is weakly optimal if and only if there is a periodic embedded legal line in the tension graph (i.e. a legal immersed loop in the quotient graph). In particular, optimal maps are weakly optimal.*

Proof. First note that the Lipschitz constant of any straight map f from X to Y provides an upper bound for the stretching factor of a loop. Hence, for any loop, γ ,

$$\frac{L_Y(\gamma)}{L_X(\gamma)} \leq \Lambda(X, Y) \leq \text{Lip}(f).$$

Let $f : X \rightarrow Y$ be our straight map. Suppose first that we have an embedded legal line, $L \subseteq \tilde{X}_{\max}$. To say that L is periodic means that L is the axis of a hyperbolic element, g . Moreover, the axis of g in Y is contained in $f(L)$, and since L is legal, the axis is exactly equal to $f(L)$ (as $f|_L$ is an embedding). Hence the stretching factor for g is exactly the Lipschitz constant for f . Thus,

$$\Lambda(X, Y) \leq \text{Lip}(f) = \frac{L_Y(g)}{L_X(g)} \leq \Lambda(X, Y).$$

Thus f is weakly optimal.

Conversely, suppose that $f : X \rightarrow Y$ is weakly optimal. By the Sausage Lemma 3.7, we may find a loop, γ , whose stretching factor equals $\Lambda(X, Y)$. Then,

$$\Lambda(X, Y) = \frac{L_Y(\gamma)}{L_X(\gamma)} = \text{Lip}(f).$$

Let L be the axis of γ in \tilde{X} . If either (i) L is not legal or, (ii), L is not a subset of \tilde{X}_{\max} , then $\frac{L_Y(\gamma)}{L_X(\gamma)} < \text{Lip}(f)$. Thus L is our required line. □

In general optimal maps are neither unique nor do they form a discrete set, even if $X_{\max} = X$, as the following example shows. (If $X_{\max} \neq X$ then one can use the freedom given by the lengths of edges not in X_{\max} to produce examples.)

Example 3.14 (A continuous family of optimal maps with $X_{\max} = X$). Consider $G = F_2$. Let X be a graph with three edges e_1, e_2, e_3 and two free vertices P, Q , as in Figure 1. Set the length of e_2 to be 2, name x the length of e_1 , and $1 + \delta$ that of e_3 . The parameters x, δ will be determined below. For any $t \in [0, 1]$ consider the point P_t at distance $1 + t$ from

P along e_2 , and the point Q_t at distance $1 - t$ from P along e_3 . P_t divides e_2 in oriented segments a_t, c_t . Q_t divides e_3 into b_t, d_t . Consider the straight map $f : X \rightarrow X$ defined as




FIGURE 1. A continuous family of optimal maps with $X_{\max} = X$. The red dashed line is $f(e_1)$ and the blue line is $f(e_3)$ ($f(e_2)$ is not depicted).

in the figure, sending P to P_t and Q to Q_t . If we collapse e_3 , and we homotope P_t to P along a , this corresponds to the automorphism $e_1 \mapsto e_1 \bar{e}_2, e_2 \mapsto \bar{e}_2$.

The following direct calculation shows that if we set $\delta = 1 + 2\sqrt{2}$ and $x = 2\sqrt{2}$, the map f_t is optimal for any t and all the three edges are stretched by the same amount.

The edges e_1 and e_2 are in different gates at P and e_1 and e_3 are in different gates at Q . In order to check that f_t is optimal it suffices to check that every edge is stretched by the same amount.

$$\lambda_{e_1}(f_t) = \frac{x + 4}{x} \quad \lambda_{e_2}(f_t) = \frac{1 + \delta}{2} \quad \lambda_{e_3}(f_t) = \frac{4 + 2\delta}{1 + \delta}.$$

In particular they do not depend on t . If we set $x = 2\sqrt{2}$ and $\delta = 1 + 2\sqrt{2}$ we get

$$\lambda_{e_1}(f_t) = \frac{2\sqrt{2} + 4}{2\sqrt{2}} \quad \lambda_{e_2}(f_t) = \frac{2 + 2\sqrt{2}}{2} \quad \lambda_{e_3}(f_t) = \frac{6 + 4\sqrt{2}}{2 + 2\sqrt{2}}$$

which are all equal to $1 + \sqrt{2}$. □

However, given a straight map, we can choose an optimal map which is in some sense the closest possible. Given two \mathcal{O} -maps $f, g : X \rightarrow Y$ we define

$$d_\infty(f, g) = \max_{x \in X} d_Y(f(x), g(x)).$$

Theorem 3.15 (Optimization). *Let $X, Y \in \mathcal{O}(\Gamma)$ and let $f : X \rightarrow Y$ be a straight map. There is a map⁴ $\text{weakopt}(f) : X \rightarrow Y$ which is weakly optimal and such that*

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

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

Proof of Theorem 3.15. By arguing component by component, we may assume without loss of generality that Γ is connected, hence that we can work in $\mathcal{O}(\mathcal{G})$. For this proof it will be convenient to work with both graphs and trees. (Recall the tilde-underbar Notation 2.11: $\underline{X} = G \setminus X$, and similarly for vertices and edges). By Remark 3.11 a non-free vertex will never be considered one-gated.

Let us concentrate on the first claim.

Let $\lambda = \Lambda(X, Y)$. Since straight maps are uniquely determined by their value on vertices, we need only to define $\text{weakopt}(f)$ (and g) on vertices of X . By Remark 3.4 the image of non-free vertices is fixed. We define straight maps f_t for $t \in [0, \lambda_f - \lambda]$ by moving the images of all one-gated vertices of $X_{\max}(f_t)$, in the direction given by the gate, so that

$$\frac{d}{dt}\lambda(f_t) = -1.$$

Let us be more precise on this point. We define a flow which is piecewise linear, depending on the geometry of the tension graph at time t . The key remark to have in mind is that if an edge is not in $X_{\max}(f)$, then it remains in the complement of the tension graph for small perturbations of f . Therefore, we can restrict our attention to the tension graph.

Suppose we are at time t . We inductively define sets of vertices and edges as follows:

- V_0 is the set vertices of $X_{\max}(f_t)$ which are one-gated in $X_{\max}(f_t)$;
- E_0 is the set of edges of $X_{\max}(f_t)$ incident to vertices in V_0 . We agree that such edges contain the vertices in V_0 but not others. (If an edge has both vertices in V_0 then it contains both, otherwise it contains only one of its vertices.)

Having defined V_0, \dots, V_i and E_0, \dots, E_i , we define V_{i+1} and E_{i+1} as follows:

- V_{i+1} is the set of one-gated vertices of $X_{\max}(f_t) \setminus \cup_{j=0}^i E_j$;
- E_{i+1} is the set of edges of $X_{\max}(f_t) \setminus \cup_{j=0}^i E_j$ incident to vertices in V_{i+1} . (As above such edges contain vertices in V_{i+1} but not others.)

We notice that since $G \setminus X$ is a finite \mathcal{G} -graph, we have only finitely many sets V_i , say V_0, \dots, V_k (each one formed by finitely many G -orbits).

Lemma 3.16. *If f_t is not weakly optimal, then $X_{\max}(f_t) \setminus \cup_{i=0}^k E_i$ is a (possibly empty) collection of vertices, that we name terminal vertices.*

Proof. Note that no vertex in $X_{\max}(f_t) \setminus \cup_{i=0}^k E_i$ can be one-gated, hence any vertex in $X_{\max}(f_t) \setminus \cup_{i=0}^k E_i$ is either isolated or has at least two gates in $X_{\max}(f_t) \setminus \cup_{i=0}^k E_i$. Thus if

⁴We describe an algorithm to find the map $\text{weakopt}(f)$, but the algorithm will depend on certain choices, hence the map $\text{weakopt}(f)$ may be not unique in general.

there is an edge e in $X_{\max}(f_t) \setminus \cup_{i=0}^k E_i$, the component of $X_{\max}(f_t) \setminus \cup_{i=0}^k E_i$ containing e must also contain an immersed legal loop and so f_t is weakly optimal. \square

By convention we denote the set of terminal vertices by V_∞ .

Remark 3.17. Any $e \in E_i$ has by definition at least one endpoint in V_i , and the other endpoint is in some V_j with $j \geq i$.

Our flow is defined by equivariantly moving the images $f_t(v)$ of vertices in $X_{\max}(f_t)$. We need to define a direction and a speed $s(v) \geq 0$ for any $f_t(v)$.

For $i < \infty$ each vertex in V_i has a preferred gate: the one that survives in $X_{\max}(f_t) \setminus \cup_{j=0}^{i-1} E_j$ (i.e. the unique gate of the map $f_t|_{X_{\max}(f_t)}$ such that some edge of that gate is in $X_{\max}(f_t) \setminus \cup_{j=0}^{i-1} E_j$). That gate gives us the direction in which we move $f_t(v)$.

The idea is the following. Since a vertex in V_0 is one-gated, we can define the flow so as to reduce the Lipschitz constant for every edge in E_0 (shrinking the image of each E_0 edge). Similarly, every vertex in V_1 is one gated in $X_{\max}(f_t) \setminus E_0$, so we define the flow to reduce the Lipschitz constants of edges in E_1 and so on. We have only to set speeds properly.

Lemma 3.18. *There exists G -equivariant speeds $s(v) \geq 0$ such that if we move the images of any v at speed $s(v)$ in the direction of its preferred gate, then for any edge $e \in X_{\max}(t)$*

$$\frac{d}{dt} \lambda_e(f_t) \leq -1.$$

Moreover, for any i , and for any $v \in V_i$, either $s(v) = 0$ or there is an edge $e \in E_i$ incident to v such that

$$\frac{d}{dt} \lambda_e(f_t) = -1.$$

Proof. We start by choosing a total order on the set orbits of vertices of $X_{\max}(f_t)$ (i.e. on the set of vertices of $X_{\max}(f_t)$) with the only requirement that orbits of vertices in V_i are bigger than those in V_j whenever $i > j$. This define a partial order on vertices by declaring $w > v$ when $\underline{w} > \underline{v}$. Now, we define speeds recursively starting from the the biggest vertex and going down through the order.

The speed of terminal vertices is set to zero. Let v be a vertex of $X_{\max}(f_t)$ and suppose that we already defined the speed $s(w)$ for all $w > v$.

The vertex v belongs to some set V_i . For any edge $e \in E_i$ emanating from v let u_e be the other endpoint of e , and define a sign $\sigma_e(u_e) = \pm 1$ as follows: $\sigma_e(u_e) = -1$ if the germ of e at u_e is in the preferred gate of u_e , and $\sigma_e(u_e) = 1$ otherwise. (So, for example, $\sigma_e(u_e) = 1$ if u_e is terminal, and $\sigma_e(u_e) = -1$ if $\underline{v} = \underline{u_e}$, or if $u_e \in V_i$.)

With this notation, if we move $f(v)$ and $f(u_e)$ in the direction given by their gates, and at speeds $s(v)$ and $s(u_e)$ respectively, then the derivative of $\lambda_e(f_t)$ is given by

$$-\left(\frac{s(v) - \sigma_e(u_e)s(u_e)}{L_X(e)} \right)$$

If $u_e > v$ we already defined its speed. We set

$$s(v) = \max\left\{0, \max_{u_e > v} \{L_X(e) + \sigma_e(u_e)s(u_e)\}, \max_{u_e = v} \frac{L_X(e)}{2}\right\}$$

where the maxima are taken over all edges $e \in E_i$ emanating from v . Note that there may exist some such edge with $\underline{u}_e < \underline{v}$. (By Remark 3.17 in this case $u_e \in V_i$ (same i as v), $\sigma_e(u_e) = -1$ and the derivative of λ_e will be settled later, when defining the speed of u_e .)

With the speeds defined in this way, we are sure that for any edge e we have $d/dt\lambda_e(f_t) \leq -1$ and, if $s(v) \neq 0$, then the edges that realize the above maximum satisfy $d/dt\lambda_e(f_t) = -1$. \square

The first consequence of this lemma is that if we start moving then $\lambda(f_t)$ decreases. Locally in t , when we start moving, the tension graph may lose some edges. However, the above lemma ensures that any vertex v with $s(v) \neq 0$ is incident to an edge e which is maximally stretched and $d/dt\lambda_e = -1$. Hence such an edge remains in the tension graph when we start moving. Since $d/dt\lambda_e \leq -1$ for any edge in the tension graph, it follows that when we start moving, the tension graph stabilizes. So our flow is well defined in $[t, t + \epsilon]$ for some $\epsilon > 0$. If at a time $t_1 > t$ some edge that was not previously in $X_{\max}(f_t)$ becomes maximally stretched, then we recompute speeds and we start again. A priori we may have to recompute speeds infinitely many times $t < t_1 < t_2 < \dots$ but the control on $d/dt\lambda(f_t)$ ensures that $\sup t_i = T \leq \lambda_f - \lambda$. Since the speeds, $s(v)$, are uniformly bounded (one can take the number of edges in $G \setminus X$ multiplied by the maximum length of an edge, as an upper bound) the flow has a limit for $t \rightarrow T$. More precisely, for any monotone sequence $t_n \rightarrow T$ as above, and any vertex v , the sequence $f_{t_n}(v)$ must be a Cauchy sequence and hence convergent, since all our trees are complete. Thus we can define $f_T(v) = \lim_{n \rightarrow \infty} f_{t_n}(v)$ for each vertex. This is enough to define a straight map, and then we can restart our flow from T . Therefore the set of times $s \in [0, \lambda_f - \lambda]$ for which the flow is well-defined for $t \in [0, s]$ is closed and open and thus is the whole $[0, \lambda_f - \lambda]$.

With these speeds, we have $d/dt(\lambda(f_t)) = -1$. Therefore for $t = \lambda(f) - \lambda$, and not before, we have $\lambda(f_t) = \lambda$ hence f_t is weakly optimal. We define

$$\text{weakopt}(f) = f_{\lambda(f) - \lambda}.$$

We prove now the claimed estimate on $d_\infty(f, f_t)$. The d_∞ -distance between straight maps is bounded by the d_∞ -distance of their restriction to vertices.

We first estimate the speed at which the images of vertices move. Let S be the maximum speed of vertices, i.e. $S = \max_v |s(v)|$. Let v be a fastest vertex. Since it moves, it belongs to V_s for some $s < \infty$. Let $v = v_1, v_2, \dots, v_m$ be a maximal sequence of vertices such that:

- (1) $s(v_i) > 0$ for $i < m$;
- (2) there is an edge e_i between v_i and v_{i+1} such that $e_i \in E_a$ if $v_i \in V_a$;
- (3) $\sigma_{e_i}(v_{i+1}) = 1$ for $i + 1 < m$;
- (4) $d/dt(\lambda_{e_i}(f_t)) = -1$.

By the above lemma, we have that either $s(v_m) = 0$ or $\sigma_{e_{m-1}}(v_m) = -1$. Moreover, by (2) – (3) and Remark 3.17 we have that $v_i < v_{i+1}$ and therefore the edges \underline{e}_i are all distinct.

Let γ be the path obtained by concatenating the e_i 's. By (2) – (3), γ is a legal path in the tension graph. So let

$$L = \sum_i L_X(e_i) = L_X(\gamma) \quad L_t = \sum_i L_Y(f_t(e_i)) = L_Y(f_t(\gamma)).$$

Since the e_i 's are in the tension graph and by condition (4) we have

$$L_t = \lambda(f_t)L \quad \frac{d}{dt}L_t = -L$$

On the other hand $-\frac{d}{dt}L_t \geq S$ because by (3) the contributions of the speeds of v_i do not count for $i = 2, \dots, m-1$ and $f(v_m)$ either stay or moves towards $f(v_1)$. It follows that

$$S \leq L \leq \text{vol}(X).$$

It follows that for any vertex v we have

$$d_Y(f(v), f_t(v)) \leq \int_0^t \left| \frac{d}{ds} f_s(w) \right| ds \leq \int_0^t S = tS \leq t \text{vol}(X)$$

hence

$$d_\infty(\text{weakopt}(f), f) = d_\infty(f_{\lambda(f)-\lambda}, f) \leq (\lambda(f) - \lambda) \text{vol}(X).$$

We prove the last claim of Theorem 3.15. If φ is optimal then we are done. Otherwise, there is some one-gated vertex in X_{\max} . We start moving the one-gated vertices as described above, by an arbitrarily small amount. Let g be the map obtained, clearly we can make $d_\infty(g, \varphi)$ arbitrarily small. Since φ is optimal, we must have $\lambda(g) = \lambda(\varphi)$. It follows that there is a core sub graph of X_{\max} which survives the moving. In particular, every vertex of $X_{\max}(g)$ is at least two-gated, hence g is optimal. \square

Definition 3.19. We denote by $\text{opt}(f)$ any optimal map obtained from $\text{weakopt}(f)$ as described in the proof of Theorem 3.15.

We want to stress the fact that Theorem 3.15 holds in a general context for X, Y metric one-dimensional complexes where the notions of straight and optimal maps are generalized in the obvious way.

Proposition 3.20. *Let A, B be metric one-dimensional simplicial complexes and let $f : A \rightarrow B$ a straight map. Then there is a weakly optimal map $\text{weakopt}(f)$ which is homotopic to f relatively to ∂A , such that*

$$d_\infty(f, \text{weakopt}(f)) \leq \text{vol}(A)(\text{Lip}(f) - \text{Lip}(\text{weakopt}(f))).$$

Moreover, for any $\varepsilon > 0$, there is an optimal map $g : A \rightarrow B$ homotopic to f relatively to ∂A such that $d_\infty(g, \text{weakopt}(f)) < \varepsilon$.

The proof is basically the same as that of Theorem 3.15 and it is left to the reader. (We do not use this generalization in what follows, and simply register the result as it may be interesting to the reader.)

Let $X, Y \in \mathcal{O}(\Gamma)$ and let $f : X \rightarrow Y$ be an optimal map. Let v be a vertex of X having an f -illegal turn $\tau = (e_1, e_2)$. Since $f(e_1)$ and $f(e_2)$ share an initial segment, we can identify an initial segment of e_1 and e_2 . We obtain a new element $X' \in \mathcal{O}(\Gamma)$, with an induced map, still denoted by f , from X' to Y . This is a particular case of Stallings fold ([24]). We refer to [9] for further details.

Definition 3.21. We call the above operation a **simple fold directed by f** . By a folding path directed by f we mean a sequence $(X_0, f_0), \dots, (X_n, f_n)$, where $f_0 = f$, and (X_{i+1}, f_{i+1}) is obtained by a simple fold directed by f_i .

We finish this section by proving the existence of optimal maps with an additional property, that will be used in the sequel.

Definition 3.22 (Minimal optimal maps). Let $X, Y \in \mathcal{O}(\Gamma)$. An optimal map $f : X \rightarrow Y$ is *minimal* if its tension graph consists of the union of axes of maximally stretched elements it contains. In other words, if any edge $e \in X_{\max}$ is contained in the axis of some element in $\pi_1(X_{\max})$ which is maximally stretched by f .

Note that not all optimal maps are minimal, as the following illustrates.

Example 3.23. Let X be the graph consisting of two barbells joined by an edge, as in Figure 2. All edges have length one except the two lower loops that have length two.




FIGURE 2. A non-minimal optimal map. The dots $f(x)$ and $f(y)$ are not vertices, all other crossings are. The red line is the image of the left “bar-edge” of the top barbell.

Let $f : X \rightarrow X$ be the straight map that exchanges the the top and bottom barbells (preserving left and right) and maps x to the middle point of the lower left loop, and y to the middle point of the lower right loop (see the figure).

The restriction of f to the lower barbell is 1-Lipschitz (each loop is shrunk and the bar is the same length as its image), while the stretching factor of all top edges is two. Hence the tension graph X_{\max} is the top barbell. The map is optimal because all vertices of X_{\max} are two gated, but the “bar-edges” of the top barbell are not in the axis of any maximally stretched loop. This is because the only legal loops in X_{\max} are the two lateral loops of

the barbell. Clearly this map can be homotoped to a map with smaller tension graph. As the next theorem shows this is always the case for non-minimal optimal maps. \square

Theorem 3.24. *Let $X, Y \in \mathcal{O}(\Gamma)$ and let $f : X \rightarrow Y$ be an optimal map. If f locally minimizes the tension graph amongst all optimal maps $X \rightarrow Y$, then f is minimal. Moreover, given $g : X \rightarrow Y$ optimal, for any $\varepsilon > 0$ there is a minimal optimal map $f : X \rightarrow Y$ with $d_\infty(g, f) < \varepsilon$.*

Proof. The first claim clearly implies the second, because the tension graph is combinatorially finite, hence the set of possible tension graphs is finite and we can always locally minimize it.

We will prove the contrapositive, that if f is not minimal then we can decrease the tension graph by perturbations as small as we want. The spirit is similar to that of the proof of Theorem 3.15.

As above, connectedness plays no role and we can work in $\mathcal{O}(\mathcal{G})$ without loss of generality. We will work with graphs rather than trees. For the ease of the reader we omit the underlines, and we declare that X, Y are G -graphs. Also we choose an orientation on edges, using the classical bar-notation to indicate the inverse.

At the level of graphs, the non-minimality of f translates to the fact that there is an edge α in the tension graph which is not part of any legal loop in X_{\max} .

Let x be the terminal vertex of the oriented edge α . We say that a path starting at x is α -legal, if it is a legal path in the tension graph, whose initial edge, e , is not in the same gate as $\bar{\alpha}$. We say a loop at x is α -legal if, considered as paths, both the loop and its inverse are α -legal.

If the terminal vertex of α admits an α -legal loop and the initial point of α also admits an $\bar{\alpha}$ -legal loop, then we can form the concatenation of these loops with α to get a legal loop in the tension graph crossing α and contradicting our hypothesis. (Note that an α -legal loop need not be legal as a loop; that is, the lift of the loop to the tree need not be a legal line. We simply require that the loops can be concatenated in this way with α to form a legal loop.)

Hence, by reversing the orientation of α if necessary, we may assume that the endpoint x (rather than the initial point) admits no α -legal loops.

We will show that it is possible to move the f -image of x a small amount (and possibly some other vertices) so that we obtain an optimal map with smaller tension graph. Let ε be small enough so that if an edge is not in X_{\max} , then it remains outside the tension graph for any perturbation of f by less than ε .

From now on, we restrict ourselves to the tension graph. We say that a vertex v is legally seen from x if there is an α -legal path γ from x to v . Note that in this case v is free. Indeed, otherwise the path γ followed by its inverse can in fact be turned into an α -legal loop thanks to the action of the vertex group ($\gamma\bar{\gamma}$ has a legal ‘‘lift’’ to \tilde{X} defined by using the action of the stabilizer of \tilde{v}). Since v is free, we can move $f(v)$.

We want to choose a direction to move the images of vertices α -legally seen from x . First, the direction we choose for $f(x)$ is given by the gate of α . That is, we move $f(x)$ so as to reduce the length of $f(\alpha)$. For any vertex, v , α -legally seen from x , via a path γ , we move

$f(v)$ backwards via the last gate of γ . That is, we move $f(v)$ so as to retrace γ . Note that this direction depends only on v and not on the choice of γ . This is because, were there to be another α -legal path from x to v , γ' , then the concatenation $\gamma\overline{\gamma'}$ would define an α -legal loop at x unless the terminal edges of γ , γ' lie in the same gate. Hence directions are well defined.

Observe that if the initial point x_0 of α is legally seen from x , and γ is an α -legal path from x to x_0 , then the last edge of γ must be in the same gate as α , because otherwise the concatenation of γ and α would form a legal loop containing α . It follows that, whether x_0 can be legally seen from x or not, in either case the length of $f(\alpha)$ decreases when we move x and possibly x_0 .

Next we move by ε all the images of vertices legally seen from x , in the directions given above. Consider an edge, β (not equal to α or its inverse) in the tension graph. If neither vertex of β is α -legally seen from x , then the image of β is unchanged and it remains in the tension graph. Otherwise, suppose that the initial vertex of β is α -legally seen from x , via a path γ , whose terminal edge is $\overline{\eta}$. If η and β are in different gates, then the terminal vertex of β is also α -legally seen from x and both vertices are moved the same amount, such that the length of the image of β remains unchanged. If, conversely, η and β are in the same gate then either the length of the image of β is reduced (if the terminal vertex is not α -legally seen) or it remains unchanged (if it is. For instance if $\eta = \beta$.) Moreover, by our above observation, the length of the image of α must strictly decrease. In particular, α itself is no longer in the tension graph.

On the other hand, since the tension graph has no one-gated vertices, there is at least one α -legal path emanating from x , and so some part of the tension graph survives. Since f is optimal, our assumption on ε implies that the new map is optimal and it has a tension graph strictly smaller than f . \square

4. DISPLACEMENT FUNCTION AND TRAIN TRACK MAPS FOR AUTOMORPHISMS

For the rest of the section we fix G, \mathcal{G} and $\Gamma = \sqcup_i \Gamma_i$ as in Definitions 2.6 and 2.20. (Recall that CV_n is a particular case of $\mathcal{O}(\Gamma)$.) If not specified otherwise, $\phi = (\sigma, \phi_1, \dots, \phi_k)$ will be an element of $\text{Aut}(\Gamma)$ - recall Definition 2.37.

This section is devoted to the study of train track maps, and related objects, from a metric point of view. In particular, we prove that the points which are minimally displaced by ϕ are exactly those admitting a *partial train track map* for ϕ , see Definition 4.11. (In the irreducible case, this amounts to showing that points of minimal displacement are precisely train track maps, in the usual sense. We broaden the class of maps to allow for the reducible case as well.)

The spirit of our analysis is that of [3, 9]. We will recall the main facts proved in [9] for irreducible elements of $\text{Out}(G)$, and generalize such facts to the case of $\text{Out}(\Gamma)$. Connectedness does not really play a crucial role, and most of the arguments of [9] transfer without requiring embellishment. The main contribution of this section is to generalize from irreducible to reducible automorphisms.

Definition 4.1 (Maps representing ϕ). Let $X \in \mathcal{O}(\Gamma)$ and $\phi = (\sigma, \phi_1, \dots, \phi_k)$ be an automorphism. We say that a (straight) map $f : X \rightarrow X$ represents ϕ if f maps X_i to $X_{\sigma(i)}$, by a map we denote by f_i which is equivariant in the following sense:

For each $i = 1 \dots, k$ the map

$$f_i : X_i \rightarrow X_{\sigma(i)},$$

is equivariant with respect to the isomorphism $\phi_i : H_i \rightarrow H_{\sigma(i)}$. This means that for each $x \in X_i$ and each $h \in H_i$ we have,

$$f_i(h \cdot x) = \phi_i(h) \cdot f_i(x)$$

We also require that each f_i is a straight map. We say that f is optimal if each f_i is optimal.

If X is a Γ -graph, then a map $f : X \rightarrow X$ represents ϕ if it has a lift $\tilde{f} : \tilde{X} \rightarrow \tilde{X}$ representing ϕ .

Note that a map $f : X \rightarrow X$ representing ϕ can be viewed as an \mathcal{O} -map $f : X \rightarrow \phi X$.

Definition 4.2 (Displacements). For any $[\phi] \in \text{Out}(\Gamma)$ we define the function

$$\lambda_\phi : \mathcal{O}(\Gamma) \rightarrow \mathbb{R} \quad \text{by} \quad \lambda_\phi(X) = \Lambda(X, \phi X).$$

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

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

Definition 4.3 (Minimally displaced points). For any automorphism ϕ we define sets:

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

$$\text{LocMin}(\phi) = \{X \in \mathcal{O}(\Gamma) : \exists U \ni X \text{ open s.t. } \forall Y \in U \lambda(X) \leq \lambda(Y)\}$$

Remark 4.4 (Fold-invariance of $\text{Min}(\phi)$). A fold directed by a weakly optimal map does not increase λ . This is because the fold naturally induces a map with the same Lipschitz constant, and it is easy to see that a legal loop in the tension graph for the original map becomes a legal loop for the folded map in the tension graph. See [9] for more details.

In particular, $\text{Min}(\phi)$ is invariant by folds directed by weakly optimal maps.

Definition 4.5 (Reducibility). An automorphism ϕ is called *reducible* if there is an $X \in \mathcal{O}(\Gamma)$ and $f : X \rightarrow X$ representing ϕ having a proper, non-trivial, f -invariant Γ -sub-forest (See Definition 2.12).

We say ϕ is *irreducible* if it is not reducible.

Remark 4.6. In the connected case, if $G = F_n$ then this definition coincides with the usual definition of irreducibility. For irreducible automorphisms we have $\text{Min}(\phi) \neq \emptyset$, but the converse is not true in general. (See [9] for more details.)

Remark 4.7. If ϕ is irreducible, then any closed simplex has a min-point for λ . (See for instance [9, Section 8]. See also Proposition 5.6 below.) In [3, 9] automorphisms so that $\text{Min}(\phi) \neq \emptyset$ and $\lambda > 1$ are called hyperbolic.

Definition 4.8 (Train track between trees). Let \sim be a gate structure on a (not necessarily connected) tree X . A map $f : X \rightarrow X$ is a **train track map w.r.t. \sim** if

- (1) any vertex has at least two gates w.r.t. \sim ;
- (2) f maps edges to legal paths (in particular, f does not collapse edges);
- (3) for any vertex v , if $f(v)$ is a vertex, then f maps inequivalent germs at v to inequivalent germs at $f(v)$.

We already defined the gate structure \sim_f induced by a straight map (Definition 3.10).

Definition 4.9 (Gate structure $\langle \sim_{fk} \rangle$). Let X be a (not necessarily connected) tree, and let $f : X \rightarrow X$ be a map whose components are straight. We define the gate structure $\langle \sim_{fk} \rangle$ as the equivalence relation on germs generated by all \sim_{fk} , $k \in \mathbb{N}$.

Lemma 4.10. Let $\phi \in \text{Aut}(\Gamma)$, $X \in \mathcal{O}(\Gamma)$ and \sim be a gate structure on X . Let $f : X \rightarrow X$ be a straight map representing ϕ . If $f : X \rightarrow X$ is a train track map w.r.t. \sim , then relation \sim is stronger than (i.e. it contains) $\langle \sim_{fk} \rangle$. In particular if f is a train track map w.r.t. some \sim then it is a train track map w.r.t. $\langle \sim_{fk} \rangle$.

See [9, Section 8] for a proof (where it is proved in the connected case, but connectedness plays no role).

Now we give a definition of partial train track map representing an automorphism. Our definition is given at once for both reducible and irreducible automorphisms. In the irreducible case coincides with the standard one. For reducible automorphisms there already exist notions of relative and absolute train tracks (see [5]). Our notion is different from that of relative train tracks; absolute train tracks are train tracks in our setting but not vice versa.⁵

The main motivation for this new definition is that it well-behaves with respect to the displacement function, as we will see that it characterise minimally displaced points.

Definition 4.11 (Partial train track maps for automorphisms). Let $[\phi] \in \text{Out}(\Gamma)$. Let $X \in \mathcal{O}(\Gamma)$ and let $f : X \rightarrow X$ be a straight map representing ϕ . Then we say that f is a

- **partial train track map with one-step gates** if there is a (not necessarily proper) f -invariant Γ -sub-forest $A \subseteq X_{\max}(f)$ such that
 - (1) $f|_A$ is a train track map w.r.t. \sim_f , and
 - (2) A is homotopically non-trivial; that is, A contains the axis of a hyperbolic element.

⁵Our present definition of partial train track map coincides with the notion of *optimal* train track map given in [9] for irreducible automorphisms in the connected case.

- **partial train track map** if there is a (not necessarily proper) f -invariant, homotopically non-trivial, G -sub-forest $A \subseteq X_{\max}(f)$ such that $f|_A$ is a train track map w.r.t. $\langle \sim_{f^k} \rangle$.

At level of graphs, a map $X \rightarrow X$ is said a partial train track if it is the projection of a partial train track map $\tilde{X} \rightarrow \tilde{X}$ (that is, if there is a non-trivial invariant sub-graph of the tension graph so that f restricted to that graph is a train-track map). We stress the fact that no requirements are made outside the tension graph.

Here some more remarks are needed, since the *metric* theory of train tracks maps, first introduced in [5], does not have a completely standard treatment. That is, train tracks can be defined topologically, from a simplicial viewpoint, and a metric is subsequently introduced. Our point of view is to always have a metric, and deduce the topological properties from certain minimizing conditions. Additionally, it should be noted that the standard definition requires train track maps (or representatives in general) to send vertices to vertices, whereas we do not. While this condition is extremely useful, and can often be recovered, our arguments are based on continuous deformations where it is more natural to relax this condition. These are sometimes called *simplicial* train-tracks and are useful for computation purposes. This is not a big issue as the closure of any simplex containing a partial train track also contains a simplicial one. (See [9].)

In the case that ϕ is irreducible there is not much difference between topological and metric train track maps. Indeed if $f : X \rightarrow X$ is a topological train track map representing ϕ , then one can rescale the edge-lengths of X so that f is a train track map for Definition 4.11. And the same holds true if f has no proper invariant sub-graphs. This is because train track maps do not collapse edges, hence edge-lengths can be adjusted so that every edge is stretched by the same amount. In particular, the following two results are proved in [9] for irreducible automorphisms and Γ connected. The proofs for general automorphisms are essentially the same (details are left to the reader).

Lemma 4.12. *Let $[\phi] \in \text{Out}(\Gamma)$, $X \in \mathcal{O}(\Gamma)$, and $f : X \rightarrow X$ be a straight map representing ϕ . Then f is partial train track if and only if there is an embedded periodic line L in \tilde{X}_{\max} such that $f^k(L) \subseteq \tilde{X}_{\max}$ and $f^k|_L$ is injective for all $k \in \mathbb{N}$ (here $A = \cup_k f^k(L)$ is the invariant sub-forest). In particular if f is partial train track then*

- (1) f^k is a partial train track;
- (2) $\text{Lip}(f) = \Lambda(X, \phi X)$ (hence f is weakly optimal);
- (3) $\text{Lip}(f)^k = \text{Lip}(f^k) = \Lambda(X, \phi^k X)$.

Corollary 4.13. *Let $\phi \in \text{Aut}(\Gamma)$, $X \in \mathcal{O}(\Gamma)$, and $f : X \rightarrow X$ be a map representing ϕ . Suppose that there is an embedded periodic line L in \tilde{X} such that $f^k|_L$ is injective for all $k \in \mathbb{N}$. Suppose moreover that $\cup_k f^k(L) = \tilde{X}$. Then there is X' obtained by rescaling edge-lengths of X so that $\text{Str}(f) : X' \rightarrow X'$ is a train track map.*

In general, if $\cup_k f^k(L)$ is just an f -invariant subtree Y of X , we can adjust edge lengths so that every edge of Y is stretched the same, but we cannot guarantee a priori that $Y \subset X_{\max}$.

Definition 4.14 (Train track sets). For any $[\phi] \in \text{Out}(\Gamma)$ we define:

$$\text{TT}(\phi) = \{X \in \mathcal{O}(\Gamma) : \exists f : X \rightarrow X \text{ partial train track}\}$$

$$\text{TT}_0(\phi) = \{X \in \mathcal{O}(\Gamma) : \exists f : X \rightarrow X \text{ partial train track with one-step gates}\}$$

If we need to specify the map we write $(X, f) \in \text{TT}(\phi)$ or $(X, f) \in \text{TT}_0(\phi)$.⁶

Theorem 4.15. *Let $[\phi] \in \text{Out}(\Gamma)$. Then*

$$\overline{\text{TT}_0(\phi)} = \text{TT}(\phi) = \text{Min}(\phi) = \text{LocMin}(\phi)$$

where the closure is made with respect to the simplicial topology.

Proof. If ϕ is irreducible and Γ connected, the proof is given in [9] and goes through the following steps:

- (1) $\text{TT}_0(\phi) \subseteq \text{TT}(\phi) \subseteq \text{Min}(\phi) \subseteq \text{LocMin}(\phi)$.
- (2) If X locally minimizes λ_ϕ in Δ_X , and $f : X \rightarrow X$ is an optimal map representing ϕ then X_{\max} contains a homotopically non-trivial f -invariant sub-forest, A .
- (3) $\text{TT}_0(\phi)$ is dense in $\text{LocMin}(\phi)$.
- (4) $\text{TT}(\phi)$ is closed.

We now adapt the proof so that it works also for ϕ reducible and general Γ . Clearly $\text{Min}(\phi) \subseteq \text{LocMin}(\phi)$. By Lemma 4.10 $\text{TT}_0(\phi) \subseteq \text{TT}(\phi)$. We now show that $\text{TT}(\phi) \subseteq \text{Min}(\phi)$, arguing by contradiction. If $X \in \text{TT}(\phi)$ and $\lambda_\phi(X) > \lambda(\phi)$ then there is $Y \in \mathcal{O}(\Gamma)$ such that $\lambda_\phi(Y) < \lambda_\phi(X)$. By Lemma 4.12 $\Lambda(X, \phi^k X) = \lambda_\phi(X)^k$ but then

$$\begin{aligned} \lambda_\phi(X)^k &= \Lambda(X, \phi^k X) \leq \Lambda(X, Y)\Lambda(Y, \phi^k Y)\Lambda(\phi^k Y, \phi^k X) \\ &= \Lambda(X, Y)\Lambda(Y, \phi^k Y)\Lambda(\phi^k Y, \phi^k X) \leq \Lambda(X, Y)\Lambda(Y, X)\lambda_\phi(Y)^k \end{aligned}$$

thus $(\frac{\lambda_\phi(X)}{\lambda_\phi(Y)})^k$ is bounded for any k , which is impossible if $\frac{\lambda_\phi(X)}{\lambda_\phi(Y)} > 1$.

Thus we have

$$\text{TT}_0(\phi) \subseteq \text{TT}(\phi) \subseteq \text{Min}(\phi) \subseteq \text{LocMin}(\phi).$$

Lemma 4.16. *Suppose (X, f) locally minimizes λ_ϕ in Δ_X . Then there is a homotopically non-trivial $A \subseteq X_{\max}$ which is f -invariant.*

Proof. For every $\epsilon > 0$, consider the ϵ -neighbourhood of X in Δ_X . For each point in this neighbourhood, f induces a map on it via rescaling. We optimize that map by using Theorem 3.15, and consider the tension graph, A_ϵ with respect to that optimal map. By abuse of notation, we think of A_ϵ as a subgraph of X (since all we have done is rescale edges). A_ϵ is homotopically non-trivial because of Proposition 3.13.

Now for each (sufficiently small) ϵ , choose a particular X_ϵ in the ϵ -neighbourhood of X , such that

- X_ϵ minimizes λ_ϕ in the ϵ -neighbourhood of X in Δ_X (we allow that X_ϵ could be X and in particular, we have that $\lambda_\phi(X) = \lambda_\phi(X_\epsilon)$), and

⁶We remark that, since in the irreducible case our present definition of train track map corresponds to that of *optimal* train track map of [9], the two definitions of TT and TT_0 coincide with those given in [9].

- the tension graph, A_ϵ is smallest, with respect to inclusion, amongst all possible choices, subject to the previous condition.

In particular, since there are only finitely many subgraphs, by taking sufficiently small ϵ we may assume that $A := A_\epsilon$ does not depend on ϵ .

Let g_ϵ denote the optimal map on X_ϵ (obtained as above) and f_ϵ denote the map on X obtained by rescaling g_ϵ . (That is, g_ϵ is simply the optimization — via Theorem 3.15 — of f , when thought of as a map on X_ϵ , and f_ϵ is g_ϵ , thought of as a map on X .)

Then it is clear, by Theorem 3.15, that $\lim_{\epsilon \rightarrow 0} d_\infty(f_\epsilon, f) = 0$ and hence $f_\epsilon \rightarrow f$, uniformly.

If A_ϵ (thought of as a subgraph of X_ϵ) contains an edge e whose image (under g_ϵ) is not in A_ϵ , then by shrinking (the orbit of) such an edge, either we reduce $\lambda_\phi(X) = \lambda_\phi(X_\epsilon)$ — which is impossible — or we reduce the tension graph — which is also impossible. Thus $A = A_\epsilon$ is g_ϵ and hence f_ϵ invariant.

But now the fact that $f_\epsilon \rightarrow f$, implies that A is f -invariant. Moreover, by Theorem 3.15, for sufficiently small ϵ , A_ϵ — hence A — will be a subgraph of X_{\max} since $d_\infty(g_\epsilon, f) \rightarrow 0$. \square

Lemma 4.17. $\text{LocMin}(\phi) \subseteq \overline{\text{TT}_0(\phi)}$. *More precisely, let $X \in \mathcal{O}(\Gamma)$ and fix $f : X \rightarrow X$ an optimal map representing ϕ . Suppose X has an open neighbourhood U such that for any $Y \in U$ obtained from X by a sequence of simple folds directed by f , we have $\lambda_\phi(X) \leq \lambda_\phi(Y)$. Then there is a sequence $Y_n \in U$, all contained in the same simplex, with $Y_n \rightarrow X$ and $\tilde{Y}_n \in \text{TT}_0$, each equipped with a partial train track map, f_n such that $f_n \rightarrow f$ uniformly.*

Proof. The proof is basically the same as in [9]. When Y obtained from X by folds directed by f , then we let f_Y denote the induced optimal map. First we remark that if Y is obtained from X by folds directed by f then $\lambda_\phi(Y) \leq \lambda_\phi(X)$ and by minimality of X we have $\lambda_\phi(Y) = \lambda_\phi(X)$. We consider the gate structure induced by f_Y . We call a vertex of Y_{\max} *foldable* if it has at least two edges of Y_{\max} in the same gate.

Locally, by using arbitrarily small folds in X_{\max} , directed by f , we find $Y \in U$ such that

- (1) $\lambda_\phi(Y) = \lambda_\phi(X)$;
- (2) the simplex $\Delta = \Delta_Y$ maximizes the dimension among simplices reachable from X via folds directed by f ;
- (3) Y minimizes Y_{\max} among points of Δ satisfying (1);
- (4) Y maximizes the number of orbit of foldable vertices of Y_{\max} among points of Δ satisfying (1), (3).

Let $A \subseteq Y_{\max}$ be a homotopically non-trivial f_Y -invariant sub-forest given by Lemma 4.16. We claim that $f_Y|_A$ is a train track map with one-step gates. Indeed, otherwise there is either an edge e or a legal turn τ in A having illegal image. Let v be the vertex of τ .

- If $f_Y(e)$ contains an illegal turn η then by folding (the orbit of) it a little, we would reduce the tension graph, contradicting (3). (Note that $\eta \subset Y_{\max}$ because $A \subseteq Y_{\max}$ is f_Y -invariant, thus by folding η we do not change simplex of $\mathcal{O}(\Gamma)$ thanks to (2).)
- If $f_Y(\tau)$ is an illegal turn η then we fold it a little. Either Y_{\max} becomes one-gated at v , and in this case the optimization process reduces the tension graph, contradicting (3), or v was not foldable at Y and becomes foldable, thus contradicting (4).

Finally, note that given such an Y , the sequence Y_n can be chosen in Δ_Y . \square

In particular, since $\text{Min}(\phi)$ is clearly closed, we now have:

$$\text{LocMin}(\phi) \subseteq \overline{\text{TT}_0(\phi)} \subseteq \overline{\text{TT}(\phi)} \subseteq \overline{\text{Min}(\phi)} = \text{Min}(\phi) \subseteq \text{LocMin}(\phi)$$

hence all inclusions are equalities.

Lemma 4.18. $\overline{\text{TT}(\phi)} = \text{TT}(\phi)$.

Proof. Let $X \in \overline{\text{TT}(\phi)} = \text{Min}(\phi)$. Let $f : X \rightarrow X$ be an optimal map representing ϕ . By Lemma 4.17 there is $Y_n \rightarrow X$ and $f_n \rightarrow f$ so that $(Y_n, f_n) \in \text{TT}_0(\phi)$. By Lemma 4.12 there is an embedded periodic line L_n in $(Y_n)_{\max}$ such that $f_n^k(L_n) \subset (Y_n)_{\max}$ is embedded for all $k \in \mathbb{N}$.

We will argue that, up to taking subsequences, there is a single embedded periodic line L which is f_n -legal for all n . (More precisely, there is a single hyperbolic element whose axis in each Y_n is both legal and contained within the tension graph of Y_n . We can think of this as a single topological line, since the Y_n all belong to the same simplex.)

First of all, we may assume that each f_n has the same (topological) tension graph. Next we claim that, for some n , the line L_n is f_m -legal for infinitely many m .

Suppose this is not the case. Then for each n , there is a turn crossed by L_n which is f_m -illegal for infinitely many m ; this is because each L_n only crosses finitely many orbits of turns, so one of these orbits must be illegal for infinitely many m (recall that gate structures are equivariant) since otherwise L_n would be f_m -legal for infinitely many m .

Starting with Y_1 , we can now define a subsequence Y_n as follows: Let τ_1 be a turn in L_1 which is illegal for infinitely many m . Now choose the first Y_2 where τ_1 is f_2 -illegal and let τ_2 be a turn crossed by L_2 such that there are infinitely many m where both τ_1 and τ_2 are f_m -illegal. (If there were no such turn, then amongst the infinitely many Y_m where τ_1 is illegal, we would get infinitely many m in which L_2 is f_m -legal and we would be done.) We may continue in this fashion to obtain (after renumbering) a sequence Y_n such that τ_n is a turn crossed by L_n (hence is f_n -legal) and all the turns $\tau_1, \dots, \tau_{n-1}$ are f_n -illegal.

By the equivariance of the f_n , the τ_n must all be in distinct orbits. If there were infinitely many of the τ_n we should be able to find a non-free vertex, v , and two edges e_1, e_2 so that the turn defined by e_1, e_2 is in the same orbit as some τ_r (for some r) and the turn defined by ge_1, he_2 is in the orbit of some τ_s (for some s), where g, h are group elements stabilising v and $g \neq h$. (That is, if we project to the quotient graph, then τ_r and τ_s look the same. The fact that they are in different orbits now implies that they are based at a non-free vertex as described.)

But now for some (almost all) Y_m , both τ_r and τ_s are f_m -illegal, which contradicts the fact that f_m is equivariant and the action on the edge set is free. Therefore the process of choosing our subsequence must terminate at some point, at which time we will have produced a single embedded periodic line L which is f_m -legal for infinitely many m . Without loss of generality, we may assume that L is f_n -legal for all n .

Since $f_n \rightarrow f$ and the maps are all straight, $L \subset X_{\max}$ and $f^k(L) \subset X_{\max}$. Moreover, if f^k were not injective on L for some k , then we could find $\varepsilon > 0$ and points p, q with $d_X(p, q) = \varepsilon$ and $f^k(p) = f^k(q)$. Now the fact that $f_n \rightarrow f$ would contradict the fact that

$f_n^k|_L$ is a homothety of ratio $\lambda(\phi)$. Thus $f^k|_L$ is embedded for any k , f is a partial train track map and so $X \in \text{TT}(\phi)$. \square

This completes the proof of Theorem 4.15. \square

We end this section by proving a lemma which is basically a rephrasing of Lemma 4.17 with a language which will be more usable. (For example we will use it in forthcoming part II of the present paper.)

Definition 4.19 (Exit points). Let $[\phi] \in \text{Out}(\Gamma)$. 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 exists an optimal map $f : X \rightarrow X$, representing ϕ , a point $X_E \in U$, and a folding path (Definition 3.21) directed by f , $X = X_0, X_1, \dots, X_m = X_E$ in U , such that Δ_{X_i} is finitary face of $\Delta_{X_{i+1}}$, Δ_X is a proper face of Δ_{X_E} , and such that

$$\lambda_\phi(X_E) < \lambda_\phi(X)$$

(strict inequality).

Remark. The idea is that an exit point is one which allows one to decrease the displacement by arbitrarily small folds. The complication is that a simple fold may not be sufficient, so we allow folding paths.

The principal application of this is as below; when X does not admit a partial train track representing ϕ , but is minimally displaced within its simplex, then it will be an exit point. This idea of folding to decrease the displacement goes back to [5].

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

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

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

(See Figure 3.)

In particular X is an exit point of Δ_X .




FIGURE 3. Graphical statement of Lemma 4.20

Proof. Let's prove the first claim. Since $X \notin \text{TT}(\phi)$, by Theorem 4.15 there is a neighbourhood of X in Δ_X which is contained in the complement of $\text{TT}_0(\phi)$. Without loss of generality we may assume that $U \cap \Delta_X$ is contained in such neighbourhood.

Let $f : X \rightarrow X$ be an optimal map representing ϕ . If there is a non-trivalent foldable vertex in X_{\max} then we set $Z = X$ and we are done. Otherwise, consider $Z \in U \cap \Delta_X$ obtained from X by a fold directed by f (we still denote by $f : Z \rightarrow Z$ the map induced by f). We have $\lambda(Z) \leq \lambda(X)$. Since $\lambda(X)$ is a local minimum in Δ_X , we must have $\lambda(Z) = \lambda(X)$. Let $\underline{A} \subset \underline{Z}_{\max}$ be an f -invariant sub-graph given by Lemma 4.16. Since $Z \notin \text{TT}_0(\phi)$, the restriction $f|_{\underline{A}}$ is not a train-track with one-step gates. That is, $f : Z \rightarrow Z$ is not a partial train track map with one-step gates.

It follows that by using folds directed by optimal maps we can either

- a) reduce the tension graph; or
- b) increase the number of foldable vertices; or
- c) create a non-trivalent foldable vertex.

So far Z is generic. We choose $Z \in U \cap \Delta_X$ so that, in order:

- (1) it locally minimizes the tension graph;
- (2) it locally maximizes the number of foldable vertices among points satisfying (1).

For such a Z the only possibility that remains in the above list of alternatives is *c*), which therefore admits a fold directed by an optimal map into a simplex of strictly larger dimension.

We continue this process inductively, and we get the result due to the fact that the simplicial dimension of $\mathcal{O}(\Gamma)$ is bounded. □

5. BEHAVIOUR OF λ AT BORDIFICATION POINTS

For the rest of the section we fix G, \mathcal{G} and $\Gamma = \sqcup_i \Gamma_i$ as in Definitions 2.6 and 2.20. We also fix $\phi \in \text{Aut}(\Gamma)$ and if there is no ambiguity we understand that

$$\lambda = \lambda_\phi.$$

In this section we discuss the behaviour of λ at boundary points of outer space, that is to say, when we reach points in $\partial_\infty \mathcal{O}(\Gamma)$. As above, we remind the reader that the results of this section hold true in particular for CV_n and its simplicial bordification.

We will see that the function λ is not continuous and we will provide conditions that assure continuity along particular sequences. We will also focus on the behaviour of λ on horoballs. In this section we will often work with Γ -graphs. We recall that we are denoting by \underline{X} the Γ -graph corresponding to $X \in \mathcal{O}(\Gamma)$.

Points near the boundary at infinity have some sub-graph that is almost collapsed. This is usually referred to as the ‘‘thin’’ part of outer space. We introduce now more quantified notions of ‘‘thinness’’.

Definition 5.1 (ε -thinness). Let $\varepsilon > 0$. A point $X \in \mathcal{O}(\Gamma)$ is ε -thin if there is a nontrivial loop γ in \underline{X} such that $L_X(\gamma) < \varepsilon \text{vol}(X)$.

Definition 5.2 ((M, ε) -collapsed points). Let $M, \varepsilon > 0$. A point $X \in \mathcal{O}(\Gamma)$ is (M, ε) -collapsed if there is a non-trivial loop γ in \underline{X} such that $L_X(\gamma) < \varepsilon \operatorname{vol}(X)$ and for any other loop η such that $L_X(\eta) \geq \varepsilon \operatorname{vol}(X)$ we have $L_X(\eta) > M \operatorname{vol}(X)$.

Definition 5.3 (ε -thin part). Let $\varepsilon > 0$. For any $X \in \mathcal{O}(\Gamma)$ we define X_ε the ε -thin part of X as the sub-forest formed by the union of the axes of elements γ with $L_X(\gamma) < \varepsilon \operatorname{vol}(X)$. (Note that $\underline{X}_\varepsilon$ is a core graph.)

Definition 5.4 (ϕ -invariance). Let $X \in \mathcal{O}(\Gamma)$. A Γ -sub-forest $A \subset X$ is called ϕ -invariant if there is a straight map $f : X \rightarrow X$ representing ϕ such that $f(A) \subseteq A$. The quotient, \underline{A} is called a ϕ -invariant subgraph of \underline{X} .

We now state some easy facts, the first of which can be found in [3].

Proposition 5.5. *For any $C > \lambda(\phi)$ there is $\varepsilon > 0$ such that for any $X \in \mathcal{O}(\Gamma)$, if $\lambda_\phi(X) < C$ and $X_\varepsilon \neq \emptyset$ then \underline{X} contains a non-trivial⁷ ϕ -invariant subgraph.*

For a proof in the case Γ is connected see [9, Section 8] (connectedness plays in fact no role).

However, we will need a slightly more precise statement, in order to be able to determine a particular invariant subgraph.

Proposition 5.6. *Let $C \geq 1$ and $M > 0$. Let D be the maximal number of orbits of edges for any tree in $\mathcal{O}(\Gamma)$.*

Let $\varepsilon = 1/2 \min\{M/CD, 1/D\}$. Then, for $X \in \mathcal{O}(\Gamma)$, if $\lambda_\phi(X) < C$ and X is (M, ε) -collapsed, then X_ε is not the whole X and it is ϕ -invariant.

Proof. By definition any edge in X_ε is shorter than $\varepsilon \operatorname{vol}(X)$. Thus we have $\operatorname{vol}(X_\varepsilon) < \varepsilon \operatorname{vol}(X)D$. In particular, since $\varepsilon D < 1$ then $X_\varepsilon \neq X$ (and thus there exists a loop η with $L_X(\eta) > \varepsilon \operatorname{vol}(X)$, whence $L_X(\eta) > M \operatorname{vol}(X)$), since X is (M, ε) -collapsed.

Let $f : X \rightarrow X$ be an optimal straight map representing ϕ . By picking a maximal tree in \underline{X} , we may find a generating set of the fundamental group of (each component of) $\underline{X}_\varepsilon$ whose elements have length at most $2 \operatorname{vol}(X_\varepsilon)$. For any such generator, γ , we have that $L_X(f(\gamma))/L_X(\gamma) \leq C$ and hence, $L_X(f(\gamma)) \leq CL_X(\gamma) \leq 2C \operatorname{vol}(X_\varepsilon) < 2CD\varepsilon \operatorname{vol}(X) \leq M \operatorname{vol}(X)$. But since X is (M, ε) -collapsed, we get that $L_X(f(\gamma)) < \varepsilon \operatorname{vol}(X)$. Hence $f(\gamma)$ is homotopic to a loop in X_ε .

Varying γ we deduce that X_ε is ϕ -invariant. \square

Proposition 5.7. *Let $X \in \mathcal{O}_{gr}(\Gamma)$ and $\phi \in \operatorname{Aut}(\Gamma)$. Suppose that $A \subset X$ is a ϕ -invariant core graph. Then $\lambda_{\phi|_A}(A) \leq \lambda_\phi(X)$.*

Proof. Let $f : X \rightarrow X$ be a straight map representing ϕ . Since A is ϕ -invariant, $f(A) \subset A$ up to homotopy. By passing to the universal covering we see that $f|_A : A \rightarrow X$ retracts to a map $f_A : A \rightarrow A$ representing ϕ with $\operatorname{Lip}(f_A) \leq \operatorname{Lip}(f)$, hence $\lambda_{\phi|_A}(A) \leq \operatorname{Lip}(f_A) \leq \operatorname{Lip}(f) = \lambda_\phi(X)$. \square

⁷In the sense of Definition 4.5.

Theorem 5.8 (Lower semicontinuity of λ). *Fix $\phi \in \text{Aut}(\Gamma)$ and $X \in \mathcal{O}_{\text{gr}}(\Gamma)$. 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 $X_i \rightarrow X_\infty \in \partial_\infty \Delta_X$ which is obtained from X by collapsing a sub-graph $A \subset X$. Then ϕ induces an element of $\text{Aut}(X/A)$, still denoted by ϕ .*

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

Proof. Let M be the ‘‘systole’’ of X_∞ , that is to say the shortest length of simple non-trivial loops in X_∞ . For any $M/\text{vol}(X) > \varepsilon > 0$, eventually on i , X_i is $(M/2 \text{vol}(X), \varepsilon)$ -collapsed and $(X_i)_\varepsilon = \text{core}(A)$. By Proposition 5.6 A is ϕ -invariant, thus $\phi \in \text{Aut}(X/A)$.

For any loop γ the lengths $L_{X_i}(\gamma)$ and $L_{X_i}(\phi(\gamma))$ converge to $L_{X_\infty}(\gamma)$ and $L_{X_\infty}(\phi(\gamma))$ respectively. Therefore, if γ is a candidate in X_∞ that realizes $\lambda_\phi(X_\infty)$, we have that $\lambda_\phi(X_i) \geq L_{X_i}(\phi(\gamma))/L_{X_i}(\gamma) \rightarrow \lambda_\phi(X_\infty)$ whence the lower semicontinuity of λ .

On the other hand, by Theorems 3.24 and 3.15, for any i there is a minimal optimal map $f_i : X_i \rightarrow X_i$ representing ϕ . Let γ_i be a candidate that realizes $\lambda_\phi(X_i)$, i.e. a f_i -legal candidate in $(X_i)_{\text{max}}$. Since X is combinatorically finite, we may assume w.l.o.g. that $\gamma_i = \gamma$ is the same loop for any i . We have

$$\lambda_\phi(X_i) = \frac{L_{X_i}(\phi(\gamma))}{L_{X_i}(\gamma)} \rightarrow \frac{L_{X_\infty}(\phi(\gamma))}{L_{X_\infty}(\gamma)}$$

Thus if $L_{X_\infty}(\gamma) \neq 0$ we have $\lambda_\phi(X_\infty) = \liminf \lambda_\phi(X_i)$. It follows that if there is a jump in λ at X_∞ , then any legal candidate is contained in A . Since f_i is minimal this implies that $\text{core}(A)$ contains the whole tension graph. \square

Remark 5.9. A comment on Theorem 5.8 is required. To avoid cumbersome notation, we have decided to denote by ϕ both the element of $\text{Aut}(X)$ and the one induced in $\text{Aut}(X/A)$. So when we write $\lambda_\phi(X_\infty)$ we mean $\Lambda(X_\infty, \phi X_\infty)$ as elements in $\mathcal{O}(X/A)$. In particular, $\lambda_\phi = \inf_X \lambda_\phi(X)$ can be different if computed in $\mathcal{O}(X)$ or in $\mathcal{O}(X/A)$. When this will be crucial we will specify in which space we take the infimum.

Moreover, if $\phi|_A$ is the restriction of ϕ to A , then $\lambda_{\phi|_A}$ is calculated in the space $\mathcal{O}(A)$. While the simplex Δ_{X_∞} is a simplicial face of Δ_X , $\Delta_A \in \mathcal{O}(A)$ does not have the same meaning. One could argue that Δ_A is the simplex ‘‘opposite’’ to Δ_{X_∞} in Δ , but ϕ does not necessarily produce an element of $\text{Aut}(X/(X \setminus A))$ as the complement of A may be not invariant.

Clearly, if $A \subset X$ is ϕ -invariant then $\lambda_\phi(X/A) < \infty$. On the other hand, if A is not ϕ -invariant, its collapse makes λ explode. Thus we can extend the function λ as follows.

Definition 5.10. Let $X_\infty \in \partial_\infty \mathcal{O}_{\text{gr}}(\Gamma)$. We say that $\lambda_\phi(X_\infty) = \infty$ if X_∞ is obtained from a Γ -graph X by collapsing a sub-graph $A \subset X$ which is not ϕ -invariant. (Note that X_∞ since is not in $\mathcal{O}(\Gamma)$, then A must have some non-trivial component).

In general, the function λ is not uniformly continuous with respect to the Euclidean metric, even in region where it is bounded, and so we cannot extend it to the simplicial closure of simplices. However we see now that the behaviour of λ is controlled on segments.

We recall the description of horoballs given in 2.7. Suppose that X_∞ is obtained from a Γ -graph X by collapsing a ϕ -invariant core sub-graph $A = \cup_i A_i$. Let k_i be the number of germs of edges incident to A_i in $X \setminus A$. Then $\text{Hor}(X_\infty)$ is a product of outer spaces with marked points $\mathcal{O}(A_i, k_i)$.

Notation 5.11. We denote $\pi : \text{Hor}(X_\infty) \rightarrow \mathbb{P}\mathcal{O}(A)$ the projection that forgets marked points.

Note that we chosen X_∞ to not be projectivized and $\mathbb{P}\mathcal{O}(A)$ to be projectivized. For any $Y \in \mathbb{P}\mathcal{O}(A)$ if $Z \in \pi^{-1}(Y)$, then there is a scaled copy of Y in Z . We denote by $\text{vol}_Z(Y)$ the volume of Y in Z . With this notation in place, we can now prove a key regeneration lemma.

Lemma 5.12 (Regeneration of optimal maps). *Fix $\phi \in \text{Aut}(\Gamma)$ 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 $f_A : A \rightarrow A$ representing $\phi|_A$, and for any $\varepsilon > 0$ there is $X_\varepsilon \in \Delta_X$ such that*

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

More precisely, for any $Y \in \mathbb{P}\mathcal{O}_{gr}(A)$ and map $f_Y : Y \rightarrow Y$ representing $\phi|_A$, for any map $f : X_\infty \rightarrow X_\infty$ representing ϕ , for any $\hat{X} \in \pi^{-1}(Y)$, and for any $\varepsilon > 0$; there is $0 < \delta = \delta(f, f_Y, X_\infty, \Delta_{\hat{X}})$, such that for any $Z \in \Delta_{\hat{X}} \cap \pi^{-1}(Y)$, if $\text{vol}_Z(Y) < \delta$ there is a straight map $f_Z : Z \rightarrow Z$ representing ϕ such that $f_Z = f_Y$ on Y and

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

(hence the optimal map $\text{opt}(f_Z)$ satisfies the same inequality⁸).

Proof. We denote by $\sigma : X \rightarrow X_\infty$ the map that collapses A . If A_i is a component of A , we denote by v_i the non-free vertex $\sigma(A_i)$. Let k_i be the valence of v_i in X_∞ . For any v_i let $E_i^1, \dots, E_i^{k_i}$ be the half-edges incident to v_i in X_∞ .

Let Y_i be the components of $Y \in \mathbb{P}\mathcal{O}_{gr}(A)$. Points in $\Delta_{\hat{X}} \cap \pi^{-1}(Y)$ are built by inserting a scaled copy of each Y_i at the v_i as follows. (Now we need to pass to the universal coverings.)

For every half-edge E_i^j of X_∞ we choose a lift in \tilde{X}_∞ . The tree \tilde{X} is given by attaching \tilde{E}_i^j to a point \tilde{y}_i^j of \tilde{Y}_i , and then equivariantly attaching any other lift of the E_i^j . At the level of graphs this is equivalent to choosing $y_i^j \in Y_i$. Two different choices at the level of universal coverings differ, at the level of graphs, by closed paths in Y_i and based at y_i^j . The choice of the simplex $\Delta_{\hat{X}}$ fixes such ambiguity. Moreover for any two graphs in $\pi^{-1}(Y) \cap \Delta_{\hat{X}}$ the points y_i^j are attached to the same edge of Y_i . Let $Z \in \pi^{-1}(Y) \cap \Delta_{\hat{X}}$.

Given $f_Y : Y \rightarrow Y$, consider its lift to \tilde{Y} and set $\tilde{z}_i^j = \tilde{f}_Y(\tilde{y}_i^j)$. There is a unique embedded arc $\tilde{\gamma}_i^j$ from \tilde{z}_i^j to \tilde{y}_i^j . Let L_i be the number of edges crossed $\tilde{\gamma}_i^j$; so L_i is an integer by including partial edges crossed in the count. L_i depends only on f_Y and the choices of \tilde{y}_i^j , hence it depends only on f_Y and $\Delta_{\hat{X}}$.

⁸We notice that while $f_Z = f_Y$ on Y , this may no longer be true for $\text{opt}(f_Z)$

Now, given $f : X_\infty \rightarrow X_\infty$, there exists a continuous map $g : Z \rightarrow Z$ representing ϕ , which agrees with f_Y on Y and which is obtained by a perturbation of f on edges of X_∞ . Namely on E_i^j when we consider the path $f(E_i^j)$ we need to lift this from a path in X_∞ to a path in Z by “filling in the missing parts”. That is, we can consider $f(E_i^j)$ to be a sequence of edges in Z whose completion to a path requires insertion of suitable paths which lie entirely in (the lifts of) Y , along with the path γ_i^j which tells us where to “attach” $f(E_i^j)$.

An accurate and detailed discussion on the properties of such a map will be carried on in [10].

For the present purpose it is sufficient to note that there is a constant C such that g can be obtained so that $\text{Lip}(\text{Str}(g)) \leq \max\{\text{Lip}(f) + C \text{vol}_Z(Y), \text{Lip}(f_Y)\}$. Moreover the constant C depends only on the L_i 's, the paths added to “fill in the missing parts”, and the edge-lengths of X_∞ . Hence it depends only on $f_Y, \Delta_{\widehat{X}}, X_\infty$. (We are using $\text{vol}_Z(Y)$ as a crude estimate for the maximum length of an edge in Y .)

The result follows by setting $\delta < \varepsilon/C$ and $f_Z = \text{Str}(g)$. \square

Definition 5.13. Fix $\phi \in \text{Aut}(\Gamma)$. Let $X_\infty \in \partial_\infty \Delta \subset \partial_\infty \mathcal{O}(\Gamma)$. We say that X_∞ has *not jumped in Δ* if there is a sequence of points $X_i \in \Delta$ such that $X_i \rightarrow X_\infty$ and $\lambda_\phi(X_\infty) = \lim_i \lambda_\phi(X_i)$. We say that $X_\infty \in \partial_\infty \mathcal{O}(\Gamma)$ *has not jumped* if there is a simplex Δ intersecting $\text{Hor}(X_\infty)$ such that X_∞ has not jumped in Δ .

The above definition is for points in $\partial_\infty \mathcal{O}(\Gamma)$. By convention, we say that X has not jumped for any $X \in \mathcal{O}(\Gamma)$.

Notice that even if X_∞ has not jumped, there may exist a simplex Δ intersecting $\text{Hor}(X_\infty)$ such that X_∞ has jumped in Δ . This is because if A is the collapsed part and $\phi|_A$ does not have polynomial growth, then we can choose a point in $\mathcal{O}(A)$ with arbitrarily high $\lambda_{\phi|_A}$. Moreover, even if X_∞ has not jumped in Δ it may happen that X_∞ is not a continuity point of λ . For example if the collapsed part A has a sub-graph B which is not invariant, then the collapse of B forces λ to increase due to Proposition 5.6, and thus we can approach X_∞ with arbitrarily high λ .

Also, note that if $\lambda > 1$ at some point X , then λ is in fact unbounded on Δ_X . This is because if X contains a loop which is not ϕ -invariant, then by collapsing that loop we force λ to explode. On the other hand, if any loop is ϕ -invariant then by Theorem 3.7 we get $\lambda = 1$.

Theorem 5.14. Let $\phi \in \text{Aut}(\Gamma)$. Let $X \in \mathcal{O}_{gr}(\Gamma)$ containing an invariant sub-graph A . Let $X_\infty = X/A$ and $C = \text{core}(A)$. 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 .

Proof. The first claim is a direct consequence of Proposition 5.7. Let $X_i \in \Delta_X$ with $X_i \rightarrow X_\infty$ without jump. Then

$$\lambda_\phi(\Delta_X) \leq \lambda_\phi(X_i) \rightarrow \lambda_\phi(X_\infty) \quad \text{hence} \quad (1) \Rightarrow (2).$$

If $\lambda_\phi(X_\infty) \geq \lambda_\phi(\Delta_X)$, then first claim implies (3), so (2) \Rightarrow (3).

Finally, suppose $\lambda_\phi(X_\infty) \geq \lambda_{\phi|_C}(\Delta_C)$. For any $\varepsilon > 0$ there is $C_\varepsilon \in \Delta_C$ and a straight map $f_{C_\varepsilon} : C_\varepsilon \rightarrow C_\varepsilon$ representing $\phi|_C$ such that $\text{Lip}(f_{C_\varepsilon}) < \lambda_{\phi|_C}(\Delta_C) + \varepsilon$. By Lemma 5.12 there is a point $X_\varepsilon \in X$ and a map $f_\varepsilon : X_\varepsilon \rightarrow X_\varepsilon$ representing ϕ such that $X_\varepsilon \rightarrow X_\infty$ as $\varepsilon \rightarrow 0$ and $\text{Lip}(f_\varepsilon) \leq \lambda_\phi(X_\infty) + \varepsilon$. This, plus lower semicontinuity (Theorem 5.8), implies $\lambda_\phi(X_\varepsilon) \rightarrow \lambda_\phi(X_\infty)$. So (3) \Rightarrow (1).

The final statement is now a consequence of the fact that (2) \Rightarrow (1). \square

Lemma 5.15 (Constant before jumping). *Let $\phi \in \text{Aut}(\Gamma)$. Let $X \in \mathcal{O}_{gr}(\Gamma)$ be a point with a ϕ -invariant sub-graph A . Let $X_\infty = X/A$ and let $C = \text{core}(A)$. Let*

$$X_t = (1-t)X_\infty + tX$$

and let C_t be the metric version of C in X_t . If $\lambda_\phi(X_\infty) < \liminf \lambda_\phi(X_t)$ then for small enough $t > 0$, the function $\lambda_\phi(X_t)$ is locally constant on t ; more precisely we have

$$\lambda_\phi(X_t) = \lambda_{\phi|_C}(C_1).$$

In particular, this is the case if displacement has jumped in $\Delta = \Delta_X$ along the segment XX_∞ .

Proof. By Theorem 5.8 for t small enough there is an optimal map $f_t : X_t \rightarrow X_t$ whose tension graph is contained in C_t . Since C_t is ϕ -invariant, $f_t(C_t) \subset C_t$ up to homotopy. Since the vertices of $(X_t)_{\max}$ are at least two gated, $f((X_t)_{\max}) \subset C_t$. Therefore $\lambda_{\phi|_C}(C_t) = \text{Lip}(f_t)$ and $\lambda_\phi(X_t) = \text{Lip}(f_t) = \lambda_{\phi|_C}(C_t) = \lambda_{\phi|_C}(C_1)$ (where the last equality follows from the fact that $[C_t] = [C_1] \in \mathbb{P}\mathcal{O}_{gr}(C)$).

The last claim follows because by Theorem 5.14, and since X_∞ has jumped in Δ , we have

$$\lambda_\phi(X_\infty) < \lambda_{\phi|_C}(\Delta_C) \leq \lambda_\phi(\Delta) \leq \lambda_\phi(X_t)$$

hence $\lambda_\phi(X_\infty) < \liminf_t \lambda_\phi(X_t)$. \square

Lemma 5.16. *Let $\phi \in \text{Aut}(\Gamma)$. Let $X \in \Delta \subseteq \mathcal{O}_{gr}(\Gamma)$ be a point with a ϕ -invariant sub-graph A . Let $X_\infty = X/A$ and let $C = \text{core}(A)$. Suppose that the displacement jumps at X_∞ along all segments of Δ . Then*

- (I) $\forall Y \in \Delta_C \subseteq \mathcal{O}(C)$ there is $X^Y \in \Delta$ such that $\lambda_{\phi|_C}(Y) = \lambda_\phi(X^Y) > \lambda_\phi(X_\infty)$;
- (II) $\lambda_{\phi|_C}(\Delta_C) = \lambda_\phi(\Delta) \geq \lambda(X_\infty)$;
- (III) if X_∞ does not jump⁹ in Δ then it is a minpoint of Δ , i.e. $\lambda_\phi(X_\infty) = \lambda_\phi(\Delta)$.

⁹See Example 5.19 for an explicit case of a non-jumping min-point that jumps along segments.

Proof. (I). For any $Y \in \Delta_C$, let A^Y be a metric version of A so that $\text{core}(A^Y) = Y$ and let X_∞^Y be a graph obtained by inserting a copy of A^Y in the collapsed part of X_∞ . Since X_∞ jumps along segments, by Lemma 5.15 there is a point X^Y in the segment $X_\infty^Y X_\infty$ so that $\lambda_\phi(X^Y) = \lambda_{\phi|_C}(Y)$. Inequality $\lambda_\phi(X^Y) > \lambda_\phi(X_\infty)$ follows from lower semicontinuity and jumping.

(II). Point (I) implies $\lambda_{\phi|_C}(\Delta_C) \geq \lambda_\phi(\Delta)$ and $\lambda_{\phi|_C}(\Delta_C) \geq \lambda(X_\infty)$; Proposition 5.7 gives $\lambda_{\phi|_C}(\Delta_C) \leq \lambda_\phi(\Delta)$.

(III). By Theorem 5.14, if X_∞ does not jump, then $\lambda_\phi(X_\infty) \geq \lambda_\phi(\Delta)$, and point (II) concludes. \square

Theorem 5.17. *Let $\phi \in \text{Aut}(\Gamma)$. 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 5.14).*

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 P does not jump in Δ by Theorem 5.14);
- 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).

Proof. We start by proving the existence of an X_{\min} .

Suppose first that at every point of $\overline{\Delta}^\infty$, there is a segment in Δ to that point such that λ_ϕ is continuous along the segment. Then the statement is clear, since we can choose a minimizing sequence in Δ , whose displacements tend to $\lambda_\phi(\Delta)$. This sequence has a limit point, X_{\min} whose displacement is bounded above by $\lambda_\phi(\Delta)$ by Theorem 5.8. But the continuity along the segment implies that $\lambda_\phi(X_{\min}) \geq \lambda_\phi(\Delta)$, and so $\lambda_\phi(X_{\min}) = \lambda_\phi(\Delta)$.

Thus we may assume that there is some point $X_\infty \in \overline{\Delta}^\infty$ whose displacement jumps along all segments in Δ .

Lemma 5.18. *For any such point X_∞ there exists a point in $\text{Hor}(X_\infty) \cap \overline{\Delta}^\infty$ that realises $\lambda_\phi(\Delta)$. (Possibly X_∞ itself realises the displacement.)*

Proof. We argue by induction on the rank of Γ . In rank 1 there is nothing to prove. Since λ_ϕ jumps at X_∞ along segments, then $\lambda_\phi(X_\infty) < \infty$. Let X_∞ be obtained by some X by collapsing a ϕ -invariant sub-graph A , and let $C = \text{core}(A)$.

By induction there is $Y \in \overline{\Delta_C}^\infty$ such that $\lambda_{\phi|_C}(Y) = \lambda_{\phi|_C}(\Delta_C)$. Let A^Y be a metric graph so that $\text{core}(A^Y) = Y$ (we are replacing the core part $C = \text{core}(A)$ with Y). Let

X_∞^{tY} be the graph obtained by inserting a volume- t copy of A^Y in the collapsed part of X_∞ . Note that $X_\infty^{tY} \in \text{Hor}(X_\infty) \cap \overline{\Delta}^\infty$. By Proposition 5.7

$$(1) \quad \lambda_\phi(X_\infty^{tY}) \geq \lambda_{\phi|_C}(Y) = \lambda_{\phi|_C}(\Delta_C)$$

and by Lemma 5.16, point (II),

$$(2) \quad \lambda_{\phi|_C}(\Delta_C) = \lambda_\phi(\Delta) \geq \lambda_\phi(X_\infty).$$

Thus, by Theorem 5.14, for any t , X_∞^{tY} does not jump in Δ . If the displacement jumps at X_∞ along the segment $X_\infty^{1Y}X_\infty$, then¹⁰ by Lemma 5.15 for small enough t we have $\lambda_\phi(X_\infty^{tY}) = \lambda_{\phi|_C}(Y) = \lambda_\phi(\Delta)$. Otherwise, by (1) and (2) above we have $\lambda_\phi(X_\infty) \leq \lambda_\phi(\Delta) \leq \lambda_\phi(X_\infty^{tY})$, and this plus non-jumping, force $\lambda_\phi(X_\infty) = \lambda_\phi(\Delta)$. In any case, we found a point in $\text{Hor}(X_\infty)$ that realises $\lambda_\phi(\Delta)$. \square

Lemma 5.18 proves the first claim of Theorem 5.17. Now choose X_{\min} to be maximal, as in the statement of Theorem 5.17 (always under the assumption that X_{\min} is a minimizing point), and we shall verify the list of properties.

If X_{\min} does not minimise the displacement in its simplex, then there is a point $X' \in \Delta_{X_{\min}}$ such that $\lambda_\phi(X') < \lambda_\phi(X_{\min}) = \lambda_\phi(\Delta)$. In particular X' jumps in Δ . Lemma 5.18 tells us that there is a point $X'' \in \text{Hor}(X') \cap \overline{\Delta}^\infty$ such that $\lambda_\phi(X'') = \lambda_\phi(\Delta)$ but this would contradict our maximality assumption.

In general, for any $P \in \overline{\Delta}^\infty$ with $\Delta_{X_{\min}} \subseteq \overline{\Delta}_P^\infty \subseteq \overline{\Delta}^\infty$, if P jumps in Δ , then by Lemma 5.18 there is $P' \in \text{Hor}(P) \cap \overline{\Delta}^\infty$ such that $\lambda_\phi(P') = \lambda_\phi(\Delta)$. By our maximality assumption, this cannot happen. So P does not jump in Δ and by Theorem 5.14 we have $\lambda_\phi(P) \geq \lambda(\Delta)$.

We check now the last property. If there is a segment in Δ to X_{\min} along which λ_ϕ is continuous, we are done by taking $W = X_{\min}$ and Z sufficiently close to X_{\min} . Hence we may assume that λ_ϕ jumps along all segments to X_{\min} .

In this case, we do the construction of Lemma 5.18 with $X_\infty = X_{\min}$. That is, we set W_t to be the graph X_∞^{tY} . The maximality assumption ensures both that λ_ϕ is continuous along the segment W_1X_{\min} and, that λ_ϕ is continuous along any segment from a point in Δ to some W_t . By choosing points sufficiently close to each other, we may ensure that $\lambda_\phi(Z), \lambda_\phi(W_t) \leq \lambda(\Delta) + \epsilon$.

\square

Example 5.19. [A non-jumping point which jumps along segments] Let $F_2 = \langle a, b \rangle$ and $\phi \in \text{Aut}(F_2)$ be any iwip (so $\lambda_\phi > 1$). For $n \geq 2$, let $F_{2n+2} = \langle a_0, b_0, a_1, b_1, \dots, a_n, b_n \rangle$. For any i , ϕ induces $\phi_i \in \text{Aut}(\langle a_i, b_i \rangle)$ by identifying $\langle a_i, b_i \rangle$ with $\langle a, b \rangle$. For any $i > 0$ choose a non-trivial $w_i \in \langle a_{i-1}, b_{i-1} \rangle$ and define $\psi \in \text{Aut}(F_{2n+2})$ by setting $\psi|_{\langle a_0, b_0 \rangle} = \phi_0$, and for $i > 0$, $\psi(a_i) = \phi_i(a_i)w_i$ and $\psi(b_i) = \phi_i(b_i)w_i$.

In order to understand the displacement of the simplex, we can just give each a_i, b_i the *projective* length coming from the minimum displacement of ϕ_i , but leave ourselves free

¹⁰Even if X_∞ is supposed to jump along all segments in Δ , it may not jump along some segment in the boundary, and X_∞^{tY} may belong to the boundary.

to choose the volume of the pair $\{a_i, b_i\}$, $i \leq n - 1$. This freedom allows us to impose the extra condition that w_i be as short as we like, thus showing that $\lambda_\psi \leq \lambda_\phi + \epsilon$, for any $\epsilon > 0$. Hence, $\lambda_\phi = \lambda_\psi$.

Let R be the rose whose petals are labelled a_i, b_i and let $\Delta = \Delta_R$. For any $X \in \Delta$, the displacement of ψ is strictly bigger than λ_ϕ and the minimum is attained at the graph X_∞ corresponding to the collapse of a_i, b_i , $i = 0, \dots, n - 1$, with length of petals a_n, b_n given by a train track for ϕ_n . Nonetheless, X_∞ jumps along all segments; in more detail, the stretching factor for the loop a_1 (for example) is strictly bigger than λ_ϕ for any $X \in \Delta$. Now if we consider the segment from X to X_∞ , the stretching factor for a_1 is constant along this segment, except for the discontinuity at X_∞ , because on this segment the thin part is shrunk uniformly to zero.

Hence the displacement of $\psi|_{\langle a_0, b_0, \dots, a_{n-1}, b_{n-1} \rangle}$ equals λ_ϕ and it is attained at a boundary point. Point W of Theorem 5.17 corresponds to the collapse of petals $a_0, b_0, \dots, a_{n-2}, b_{n-2}$ from a graph Z_W of Δ whose petals a_n, b_n are stretched by ψ more than any other.

6. CONVEXITY PROPERTIES OF THE DISPLACEMENT FUNCTION

We recall that we are using the terminology ‘‘simplex’’ in a wide sense, as Δ_X is a standard simplex if we work in $\mathbb{P}\mathcal{O}(\Gamma)$ and the cone over it if we work in $\mathcal{O}(\Gamma)$. (Remember we use Definition 2.20 for Γ .)

The displacement function λ is scale invariant on $\mathcal{O}(\Gamma)$ so it descends to a function on $\mathbb{P}\mathcal{O}(\Gamma)$. In order to control the value of λ on segments in terms of its value on vertices, we would like to say that λ is convex on segments. A minor issue appears with projectivization; if Δ is a simplex of $\mathcal{O}(\Gamma)$, then its euclidean segments are well defined, and their projections on $\mathbb{P}(\mathcal{O}(\Gamma))$ are euclidean segments in the image of Δ . However, the linear parametrization is not a projective invariant (given X, Y , the points $(X + Y)/2$ and $(5X + Y)/2$ are in different projective classes).

It follows that convexity of a scale invariant function is not well-defined. In fact if σ is a segment in Δ , $\pi : \Delta \rightarrow \mathbb{P}\Delta$ is the projection, and f is a convex function on σ , then $f \circ \pi^{-1}$ may be not convex. It is convex only up to reparametrization of the segment $\pi(\sigma)$. Such functions are called quasi-convex, and this notion will be enough for our purposes.

Definition 6.1. A function $f : [A, B] \rightarrow \mathbb{R}$ is called *quasi-convex* if for all $[a, b] \subseteq [A, B]$

$$\forall t \in [a, b] \quad f(t) \leq \max\{f(a), f(b)\}.$$

Note that quasi-convexity is scale invariant.

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

Proof. Let X be a Γ -graph such that $\Delta = \Delta_X$. We use the Euclidean coordinates of Δ labelled with edges of X , namely a point P in Δ is given by a vector whose e^{th} entry is the length of edge e in P . In the same way, to any reduced loop η in X we associate its occurrence vector, whose e^{th} entry is the number of times that η passes through the edge

e. We will denote by η both the loop and its occurrence vector. With this notation, the length function is bilinear:

$$L_X(\gamma) = \langle X, \gamma \rangle$$

(where $\langle \cdot, \cdot \rangle$ denotes the standard scalar product on \mathbb{R}^k .)

Let σ be a segment in Δ with endpoints A, B . Let γ be a candidate. We consider both γ and $\phi\gamma$ as loops in X . Up to switching A and B , we may assume that

$$\frac{\langle A, \phi\gamma \rangle}{\langle A, \gamma \rangle} \geq \frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle}.$$

Such a condition is scale invariant, and since λ is scale invariant, up to rescaling B we may assume that $\langle B, \gamma \rangle > \langle A, \gamma \rangle$. We now parametrize σ in $[0, 1]$:

$$\sigma(t) = A_t = Bt + (1-t)A$$

We are interested in the function:

$$F_\gamma(t) = \frac{\langle A_t, \phi\gamma \rangle}{\langle A_t, \gamma \rangle} = \frac{\langle Bt + (1-t)A, \phi\gamma \rangle}{\langle Bt + (1-t)A, \gamma \rangle} = \frac{\langle A, \phi\gamma \rangle + t\langle B - A, \phi\gamma \rangle}{\langle A, \gamma \rangle + t\langle B - A, \gamma \rangle}$$

A direct calculation shows that the second derivative of a function of the type $f(t) = (a + tb)/(c + td)$ is given by $2(ad - bc)d/(c + td)^3$.

So the sign of $F_\gamma''(t)$ is given by

$$(\langle A, \phi\gamma \rangle \langle B, \gamma \rangle - \langle B, \phi\gamma \rangle \langle A, \gamma \rangle) (\langle B - A, \gamma \rangle)$$

which is non-negative by our assumption on A, B . Hence $F_\gamma(t)$ is (weakly)-convex and therefore quasi-convex:

$$F_\gamma(t) \leq \max\{F_\gamma(A), F_\gamma(B)\}.$$

Now, by the Sausage Lemma 3.7 we have:

$$\begin{aligned} \lambda_\phi(A_t) &= \max_\gamma F_\gamma(t) \leq \max\{\max_\gamma F_\gamma(A), \max_\gamma F_\gamma(B)\} \\ &= \max\{\lambda_\phi(A), \lambda_\phi(B)\}. \end{aligned}$$

Finally, since there are finitely many *lengths* of candidates, there is a candidate γ_o such that for t sufficiently small we have $\lambda_\phi(A_t) = F_{\gamma_o}(t)$. By convexity, if F_{γ_o} is not strictly monotone near A , then it must be locally constant, and thus $F_{\gamma_o}''(t) = 0$. Hence

$$\lambda_\phi(A) = \lambda_\phi(A_0) = \frac{\langle A, \phi\gamma_o \rangle}{\langle A, \gamma_o \rangle} = \frac{\langle B, \phi\gamma_o \rangle}{\langle B, \gamma_o \rangle} \leq \lambda_\phi(B).$$

□

Lemma 6.3. *Let $[\phi] \in \text{Out}(\Gamma)$, 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 .

Proof. Let X be a graph of groups so that $\Delta_X \subseteq \overline{\Delta}^\infty$ contains the interior of the segment \overline{AB} . By Lemma 6.2, the function λ is quasi-convex on the interior of \overline{AB} as a segment in $\mathcal{O}(X)$. Let $\{A_i\}$ and $\{B_i\}$ sequences in Δ such that $A = \lim A_i$ and $B = \lim B_i$ with $\lim \lambda(A_i) = \lambda(A)$ and $\lim \lambda(B_i) = \lambda(B)$. Such sequences exist because of the non-jumping hypothesis. For all points P in the segment \overline{AB} , there is a sequence of points P_i in the segment $\overline{A_i B_i}$ such that $P_i \rightarrow P$. By Lemma 6.2 we know

$$\lambda(P_i) \leq \max\{\lambda(A_i), \lambda(B_i)\},$$

and by lower semicontinuity (Theorem 5.8) of λ and the non-jumping assumption, such an inequality passes to the limit. In particular, if $\lambda(A) \geq \lambda(B)$, then $\lambda(P) \leq \lambda(A)$ for any $P \in \overline{AB}$.

Now suppose that $P^j \rightarrow A$ is a sequence in the segment \overline{AB} . Then by lower semicontinuity Theorem 5.8 applied to the space $\mathcal{O}(X)$ on the segment, \overline{AB} , we have

$$\lambda(A) \geq \lim_j \lambda(P_j) \geq \lambda(A).$$

□

We end this section with an estimate of the derivative of functions like the $F_\gamma(t)$ defined as in Lemma 6.2, which will be used in the sequel. As above, we use the formalism $\langle X, \gamma \rangle = L_X(\gamma)$.

Lemma 6.4. *Let X be a Γ -graph and let $\Delta = \Delta_X$ be its simplex in $\mathcal{O}(\Gamma)$. Let $A, B \in \overline{\Delta}^\infty$. Let γ be a loop in X which is not collapsed neither in A nor in B and set*

$$C = \max\left\{\frac{L_A(\gamma)}{L_B(\gamma)}, \frac{L_B(\gamma)}{L_A(\gamma)}\right\}$$

Let ϕ be any automorphism of Γ . Suppose that $\frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle} \geq \frac{\langle A, \phi\gamma \rangle}{\langle A, \gamma \rangle}$. Let $A_t = tB + (1-t)A$ be the linear parametrization of the segment AB in Δ and define $F_\gamma(t) = \frac{\langle A_t, \phi\gamma \rangle}{\langle A_t, \gamma \rangle}$. Then

$$0 \leq F'_\gamma(t) \leq C \frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle}$$

In particular, for any point P in the segment AB we have

$$\lambda_\phi(P) \geq \frac{\langle P, \phi\gamma \rangle}{\langle P, \gamma \rangle} \geq \frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle} - C \lambda_\phi(B) \frac{\|P - B\|}{\|A - B\|}$$

where $\|X - Y\|$ denotes the standard Euclidean metric on Δ .

Before the proof, a brief comment on the statement is desirable. First, note that the constant C does not depend on ϕ . Moreover, by taking the supremum where γ runs over all candidates given by the Sausage Lemma 3.7, then C does not even depend on γ . Finally if γ is a candidate that realizes $\lambda_\phi(B)$, then we get a bound of the steepness of F_γ which does not depend on ϕ nor on γ but just on $\lambda_\phi(B)$ and $\|A - B\|$.

Proof. We have

$$F_\gamma(t) = \frac{\langle A_t, \phi\gamma \rangle}{\langle A_t, \gamma \rangle} = \frac{\langle Bt + (1-t)A, \phi\gamma \rangle}{\langle Bt + (1-t)A, \gamma \rangle} = \frac{\langle A, \phi\gamma \rangle + t\langle B - A, \phi\gamma \rangle}{\langle A, \gamma \rangle + t\langle B - A, \gamma \rangle}$$

and a direct calculation show that

$$(3) \quad F'_\gamma(t) = \frac{\langle B, \gamma \rangle \langle A, \gamma \rangle}{(\langle A_t, \gamma \rangle)^2} \left(\frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle} - \frac{\langle A, \phi\gamma \rangle}{\langle A, \gamma \rangle} \right)$$

The first consequence of this equation is that the sign of F'_γ does not depend on t , and since $\frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle} \geq \frac{\langle A, \phi\gamma \rangle}{\langle A, \gamma \rangle}$, then $F'_\gamma \geq 0$. Moreover, since $\langle A_t, \gamma \rangle$ is linear on t , we have $\frac{\langle B, \gamma \rangle \langle A, \gamma \rangle}{(\langle A_t, \gamma \rangle)^2} \leq C$. Therefore we get

$$F'_\gamma(t) \leq C \frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle}$$

and the first claim is proved. For the second claim, note that the parameter t is nothing but $\|A - A_t\|/\|A - B\|$ and thus

$$F_\gamma(1) - F_\gamma(t) \leq (1-t)C \frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle} = \frac{\|B - A_t\|}{\|B - A\|} C \frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle}.$$

If $P = A_t$, we have $F_\gamma(1) = \frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle}$ and $F_\gamma(t) = \frac{\langle P, \phi\gamma \rangle}{\langle P, \gamma \rangle}$. By taking in account $\lambda_\phi(B) \geq \frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle}$ and $\lambda_\phi(P) \geq \frac{\langle P, \phi\gamma \rangle}{\langle P, \gamma \rangle}$ we get the result. \square

7. EXISTENCE OF MINIMAL DISPLACED POINTS AND TRAIN TRACKS AT THE BORDIFICATION

The first question that naturally arises in the study of the displacement function of automorphisms is about the existence of min-points. The existence of points that minimize the displacement is proved in [9, Theorem 8.4] for irreducible automorphisms. The philosophy of the proof works in the general case, but we are forced to pass to the boundary at infinity — whence taking in account possible jumps. The notion of “train track at infinity” will be introduced in order for deal with such situations. A second issue that appears in general case is that, since the bordification of $\mathcal{O}(\Gamma)$ is not locally compact, one cannot use compactness for claiming that minimizing sequences have accumulation points. We overcome that difficulties first by using a Sausage Lemma trick as in [9], and then by proving that the set of all possible simplex-displacements form a well-ordered subset of \mathbb{R} .

As a product of this machinery we have also other interesting results, such as the fact that the collection of partial train tracks detect all invariant free factors.

We use the terminology of Definitions 2.6 and 2.20 for G, \mathcal{G} and Γ .

Lemma 7.1 (Sausage lemma trick). *For any Γ , for any $X \in \overline{\mathcal{O}(\Gamma)}^\infty$ the set $\{\lambda_\phi(X) : [\phi] \in \text{Out}(\Gamma)\}$ is discrete¹¹. In other words, given X , all possible displacements of X with respect to all automorphisms (hence markings) run over a discrete set (plus possibly ∞).*

¹¹We include the possibility that $\lambda_\phi(X) = \infty$; e.g. if X has a collapsed part which is not ϕ -invariant.

Proof. This proof is similar to that of [9, Theorem 8.4], we include it for completeness. By the Sausage Lemma 3.7, $\lambda_\phi(X) = \Lambda(X, \phi X)$ is computed by the ratio of translation lengths of candidates (we include the possibility that $\lambda_\phi(X) = \infty$, e.g. if X has a collapsed part which is not ϕ -invariant). The possible values of $L_X(\phi\gamma)$ (with $\underline{\gamma}$ any loop in \underline{X}) form a discrete set just because X has finitely many orbits of edges. Candidates are in general infinitely many in number, but there are only finitely many lengths arising from them. Thus the possible values of $\Lambda(X, \phi X)$ runs over a discrete subset of \mathbb{R} . \square

Theorem 7.2. *For any Γ the global simplex-displacement spectrum*

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

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

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

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

Proof. Recall that we defined $\lambda_\phi(\Delta)$ as $\inf_{X \in \Delta} \lambda_\phi(X)$. For this proof we work with co-volume one graphs (so we are in $\mathcal{O}_{\text{gr}_1}(\Gamma)$). In any simplex we use the standard Euclidean norm, denoted by $\|\cdot\|$.

We argue by induction on the rank of Γ (See Definition 2.22). Clearly if the rank of Γ is one there is nothing to prove. We now assume the claim true for any Γ' of rank smaller than Γ .

We will show that any monotonically decreasing sequence in $\text{spec}(\Gamma)$ has a (non trivial) sub-sequence which is constant, whence the original sequence is eventually constant itself. This implies that $\text{spec}(\Gamma)$ is well-ordered. For the second claim, since $\text{spec}(\phi)$ is a subset of a well-ordered set, it is well-ordered.

We follow the line of reasoning of [9, Theorem 8.4]. Let $\lambda_i \in \text{spec}(\Gamma)$ be a monotonically decreasing sequence. Note that displacements are non-negative so λ_i converges to some number $L \in \mathbb{R}$. For any i we chose ϕ_i and a simplex Δ_i such that $\lambda_i = \lambda_{\phi_i}(\Delta_i)$ (those exist by definition of $\text{spec}(\Gamma)$). By Theorem 5.17 there exists $X_i \in \overline{\Delta_i}^\infty$ such that $\lambda_{\phi_i}(X_i) = \lambda_{\phi_i}(\Delta_i)$ and $\lambda_{\phi_i}(X_i) = \lambda_{\phi_i}(\Delta_{X_i})$. In particular, up to replacing Δ_i with Δ_{X_i} we may assume $\Delta_i = \Delta_{X_i}$. Since the displacement is scale invariant, point X_i can be chosen in $\overline{\mathcal{O}_{\text{gr}_1}(\Gamma)}^\infty$. Up to possibly passing to sub-sequences we may assume that there is $[\psi_i] \in \text{Out}(\Gamma)$ such that $\psi_i X_i$ belongs to a fixed simplex Δ . Therefore, by replacing ϕ_i with $\psi_i \phi_i \psi_i^{-1}$ we may assume that the X_i all belong to the same simplex Δ .¹² So we have

$$\lambda_{\phi_i}(X_i) = \lambda_{\phi_i}(\Delta_{X_i}) = \lambda_i \searrow L \qquad \Delta_{X_i} = \Delta_i = \Delta.$$

Up to sub-sequences, X_i converges to a point X_∞ in the simplicial closure of Δ . We show now that $\lambda_{\phi_i}(X_\infty) < +\infty$ eventually on i .

¹²Note that X_i may be a boundary point of $\mathcal{O}(\Gamma)$ and that we have made no assumption about jumps, so X_i may jump.

If $X_\infty \in \overline{\Delta}$ (the finitary closure), then it follows from the fact that $\lambda_{\phi_i}(X_i) < +\infty$. Otherwise $X_\infty \in \partial_\infty \overline{\Delta}$. In this case, let M be the length of the shortest loop in X_∞ . For any $\varepsilon > 0$, X_i is eventually (M, ε) -collapsed (because X_∞ is not in a finitary face of Δ and $X_i \rightarrow X_\infty$). For ε small enough, and eventually on i , the ε -thin part $(X_i)_\varepsilon$ is the core-subgraph of X_i which is collapsed in order to reach the deformation space $\mathcal{O}(X_\infty)$ where X_∞ lives.

Since $\lambda_i \rightarrow L$, in particular $\lambda_i < L + 1$ eventually on i . Thus, if ε is small enough to satisfy the hypothesis of Proposition 5.6 (with $C = L + 1$) then $(X_i)_\varepsilon$ is ϕ_i -invariant. This implies that λ_{ϕ_i} is not infinite on $\mathcal{O}(X_\infty)$, and in particular $\lambda_{\phi_i}(X_\infty) < +\infty$ as claimed.

Since X_i is a min-point for the function λ_{ϕ_i} on $\Delta_{X_i} = \Delta$, by Lemma 6.2 the function λ_{ϕ_i} either is constant on the segment $X_i X_\infty$ or it is not locally constant near X_∞ . By Lemma 5.15 in the latter case X_∞ has not jumped w.r.t. λ_{ϕ_i} along the segment $X_i X_\infty$.

Therefore we have the following three cases, and up to subsequences we may assume that we are in the same case for any i :

- (1) λ_{ϕ_i} is constant and continuous on $X_i X_\infty$;
- (2) λ_{ϕ_i} is constant on the interior of $X_i X_\infty$ and there is a jump at X_∞ , hence $\lambda_{\phi_i}(X_\infty) < \lambda_{\phi_i}(X_i)$ by lower semicontinuity Theorem 5.8;
- (3) λ_{ϕ_i} is monotone increasing on the segment $X_i X_\infty$, and continuous at X_∞ .

In the first case $\lambda_i = \lambda_{\phi_i}(X_\infty)$ is a (bounded) converging sequence of displacements of the single point X_∞ . By Lemma 7.1, it must be eventually constant.

In the second case we use the inductive hypothesis. Since the displacement is continuous on $\mathcal{O}(X)$, and since we have a jump at X_∞ , in this case X_∞ is a point at infinite of $\mathcal{O}(X)$. Let $C_i = (X_i)_\varepsilon$ the ε -thin part of X_i . We choose ε as above so that C_i is the ϕ_i -invariant core sub-graph of X_i which is collapsed to reach the deformation space $\mathcal{O}(X_\infty)$. Up to sub-sequences we may assume that C_i is topologically the same graph for any i . Since λ_{ϕ_i} jumps at X_∞ along the segment $X_i X_\infty$, by Lemma 5.15 we have

$$\lambda_i = \lambda_{\phi_i}(Y) = \lambda_{\phi_i|_{C_i}}(C_i)$$

for any Y in the interior of the segment $X_i X_\infty$.

If C_i would not locally minimise the function $\lambda_{\phi_i|_{C_i}}$ on its simplex Δ_{C_i} , then we could perturb a little C_i and strictly decrease the displacement $\lambda_{\phi_i}(X_i)$ contradicting the fact that $\lambda_{\phi_i}(X_i) = \lambda_{\phi_i}(\Delta_{X_i})$. So C_i locally minimises the displacement on its simplex. By quasi-convexity Lemma 6.2, in any simplex local minima are minima. Therefore

$$\lambda_i = \lambda_{\phi_i|_{C_i}}(C_i) = \lambda_{\phi_i|_{C_i}}(\Delta_{C_i}) \in \text{spec}(C_i).$$

By induction $\text{spec}(C_i)$ is well-ordered, hence the monotonically decreasing sequence λ_i must be eventually constant.

All that remains is case (3). In this case

$$\lambda_{\phi_i}(X_i) < \lambda_{\phi_i}(X_\infty).$$

Let $R > 0$ be such that for any face Δ' of Δ such that $X_\infty \notin \overline{\Delta'}^\infty$, the ball $B(X_\infty, 2R)$ is disjoint from Δ' . In other words, if $P \in B(X_\infty, 2R)$ and it is obtained from X by collapsing a sub-graph P_0 , then P_0 is collapsed also in X_∞ . Eventually on i , $X_i \in B(X_\infty, R)$. Let Y_i be the point on the Euclidean half-line from X_∞ toward X_i , at distance exactly R from X_∞ .

The stretching factor $\Lambda(X_\infty, \phi_i X_\infty)$ is realised by some candidates (Theorem 3.7), and since λ_{ϕ_i} is strictly decreasing from X_∞ to X_i , among such candidates there is at least one, say γ_i , whose stretching factor $\frac{L_X(\phi_i(\gamma_i))}{L_X(\gamma_i)}$ locally decreases near X_∞ as a function of X in the segment from X_∞ to X_i . Moreover, the function $\frac{L_X(\phi_i(\gamma_i))}{L_X(\gamma_i)}$ is the ratio of two functions that are linear on X (it is the same function $F_\gamma(t)$ of the proof of Lemma 6.2); then its derivative on any Euclidean segment has constant sign, and in particular it is strictly decreasing from X_∞ to Y_i .

By Lemma 6.4 applied with $A = Y_i$, $B = X_\infty$ (and $P = X_i$) we have

$$\lambda_{\phi_i}(X_i) \geq \lambda_{\phi_i}(X_\infty) \left(1 - C \frac{\|X_i - X_\infty\|}{R} \right).$$

where $C = \max\{\frac{L_{Y_i}(\gamma_i)}{L_{X_\infty}(\gamma_i)}, \frac{L_{X_\infty}(\gamma_i)}{L_{Y_i}(\gamma_i)}\}$. Let $\varepsilon_i = \|X_i - X_\infty\|$. Since there are finitely many lengths of candidates and by our choice of R , the constant C is uniformly bounded independently on i . Since $X_i \rightarrow X_\infty$ we have $\varepsilon_i = \|X_i - X_\infty\| \rightarrow 0$ and thus

$$(4) \quad \lambda_{\phi_i}(X_\infty)(1 - C\varepsilon_i) \leq \lambda_{\phi_i}(X_i) \leq \lambda_{\phi_i}(X_\infty).$$

Since $\lambda_i \searrow L$, left-hand side inequality tells us that $\lambda_{\phi_i}(X_\infty)$ is uniformly bounded. Therefore, by Theorem 7.1, up to sub-sequences we may assume that $\lambda_{\phi_i}(X_\infty)$ is a constant not depending on i . Moreover, (4) implies that $|\lambda_{\phi_i}(X_\infty) - \lambda_i| \rightarrow 0$, hence $\lambda_{\phi_i}(X_\infty) = L$.

Now (4) implies $\lambda_i \leq L$, and since $\lambda_i \searrow L$ and it is bounded above by L , it must be constant equal to L . \square

We suspect that $\text{spec}(\phi)$ and $\text{spec}(\Gamma)$ are not only well-ordered but in fact discrete. However, Theorem 7.2 will be enough for our purposes.

Theorem 7.3 (Existence of minpoints). *Let Γ be as in Notation 2.20. Let $[\phi]$ be any element in $\text{Out}(\Gamma)$. Then there exists $X \in \overline{\mathcal{O}(\Gamma)}^\infty$ that has not jumped and such that*

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

Proof. Let $X_i \in \mathcal{O}(\Gamma)$ be a minimizing sequence for λ_ϕ . Without loss of generality we may assume that the sequence $\lambda_\phi(\Delta_{X_i})$ is monotone decreasing and Theorem 7.2 implies that it is eventually constant. Therefore X_i can be chosen in a fixed simplex Δ . Theorem 5.17 concludes. \square

An interesting corollary of Theorem 7.3 is that we can characterize (global) jumps extending equivalence “(3) \Leftrightarrow (1)” of Theorem 5.14 from a local to a global statement.

Theorem 7.4. *Let $[\phi] \in \text{Out}(\Gamma)$. Let $X \in \mathcal{O}_{gr}(\Gamma)$ and let $X_\infty \in \partial_\infty \Delta_X$ be obtained from X by collapsing a ϕ -invariant core graph A . Then X_∞ has not jumped if and only if*

$$\lambda(\phi|_A) \leq \lambda_\phi(X_\infty).$$

Proof. Suppose that X_∞ has not jumped (Definition 5.13). Then there is a simplex Δ where X_∞ has not jumped, and the claim follows from Theorem 5.14 because $\lambda_{\phi|_A} \leq \lambda_{\phi|_A}(\Delta_A)$.

On the other hand, suppose $\lambda(\phi|_A) \leq \lambda_\phi(X_\infty)$. By Theorem 7.3 there is a simplex in $\mathcal{O}(A)$ containing a minimizing sequence for $\phi|_A$. Let A_ε be an element in that simplex so that $\lambda_{\phi|_A}(A_\varepsilon) < \lambda(\phi|_A) + \varepsilon$, and let $f_A : A_\varepsilon \rightarrow A_\varepsilon$ be an optimal map representing $\phi|_A$. Note that A_ε and A may be not homeomorphic. Let \widehat{X} be a Γ -graph obtained by inserting a copy of A_ε in X_∞ . (We notice that since A_ε may be not homeomorphic to A , we can have $\Delta_{\widehat{X}} \neq \Delta_X$. We also notice that such $\Delta_{\widehat{X}}$ is not unique as we have plenty of freedom of attaching the edges of X_∞ to A_ε .) By Lemma 5.12, for any $\varepsilon > 0$ there is an element $X_\varepsilon \in \Delta_{\widehat{X}}$ and an optimal map $f_\varepsilon : X_\varepsilon \rightarrow X_\varepsilon$ representing ϕ so that $X_\varepsilon \rightarrow X_\infty$ and $\text{Lip}(f_\varepsilon) \leq \lambda_\phi(X_\infty) + \varepsilon$, hence $\lambda_\phi(X_\varepsilon) \leq \lambda_\phi(X_\infty) + \varepsilon$. Thus X_∞ has no jump in $\Delta_{\widehat{X}}$, and therefore has not jumped. \square

Warning: Differences between min-points at infinity and partial train tracks.

By Theorem 4.15 we know that minimally displaced points and partial train tracks coincide. But some care is needed here, as that theorem is stated for points of $\mathcal{O}(\Gamma)$, and not for points at infinity. In fact, given $\phi \in \text{Aut}(\Gamma)$, $X \in \mathcal{O}(\Gamma)$, and $A \subset X$ a ϕ -invariant subgraph, a priori it may happen that $\lambda(\phi)$ is different if we consider ϕ as an element of $\text{Aut}(X)$ or of $\text{Aut}(X/A)$. That is to say, we may have $X_\infty = X/A$ such that $\lambda_\phi(X_\infty) = \lambda(\phi)$ but X_∞ is not a train track point in $\mathcal{O}(X/A)$.

For instance, consider the case where $X = A \cup B$, with both A and B invariant. Suppose that $\lambda(\phi) = \lambda(\phi|_A) > \lambda(\phi|_B)$. Now suppose that $\lambda(\phi) = \lambda_\phi(X) = \lambda_{\phi|_A}(A) = \lambda_{\phi|_B}(B)$. Collapse A . Then the resulting point X_∞ is a min point for ϕ in $\overline{\mathcal{O}(\Gamma)}^\infty$ which has not jumped, but since $\lambda(\phi|_B) < \lambda(\phi)$, it is not a min point for ϕ on $\mathcal{O}(X/A)$.

We want to avoid such a pathology. Here we need to make a distinction between $\lambda(\phi)$ computed in different spaces, so we will specify the space over which we take the infimum.

Lemma 7.5 (Min-points vs relative min-points). *Let $[\phi] \in \text{Out}(\Gamma)$. Let $X_\infty \in \overline{\mathcal{O}(\Gamma)}^\infty$ be such that:*

- *There is $X \in \mathcal{O}(\Gamma)$ such that X_∞ is obtained from X by collapsing a (possibly empty) core sub-graph A in X , and such that X_∞ has not jumped in Δ_X ;*
- $\lambda_\phi(X_\infty) = \inf_{Y \in \mathcal{O}(\Gamma)} \lambda_\phi(Y)$.

Suppose moreover that X_∞ maximizes the dimension of Δ_{X_∞} among the set of elements in $\overline{\mathcal{O}(\Gamma)}^\infty$ satisfying such conditions (such a set is not empty by Theorem 7.3). Then

$$\lambda_\phi(X_\infty) = \inf_{Y \in \mathcal{O}(X/A)} \lambda_\phi(Y).$$

(Hence it is in $\text{TT}(\phi) \subset \mathcal{O}(X/A) = \mathcal{O}(X_\infty)$.)

Proof. If A is empty this is an instance of Theorem 4.15. In general, again by Theorem 4.15, it suffices to show that X_∞ is a partial train track point for ϕ in $\mathcal{O}(X_\infty)$. Suppose the

contrary. In this case, near X_∞ there is a point $X'_\infty \in \mathcal{O}(X_\infty)$ such that $\lambda_\phi(X'_\infty) < \lambda_\phi(X_\infty)$. Indeed, if X_∞ is not a local min point in $\Delta_{X_\infty} \subset \mathcal{O}(X_\infty)$, then we can find X'_∞ just near X_∞ in Δ_{X_∞} . Otherwise, by Lemma 4.20 there is a point X'_∞ obtained from X_∞ by folds directed by optimal maps (and such that $\dim(\Delta_{X'_\infty}) > \dim(\Delta_{X_\infty})$) such that $\lambda_\phi(X'_\infty) < \lambda_\phi(X_\infty)$.

Let $\varepsilon = (\lambda_\phi(X_\infty) - \lambda_\phi(X'_\infty))/2$.

Since X_∞ has not jumped, by Theorem 7.4 we have $\lambda(\phi|_A) \leq \lambda_\phi(X_\infty)$. If $\lambda(\phi|_A) < \lambda_\phi(X_\infty)$, let $A' \in \mathcal{O}(A)$ be a point such that $\lambda_{\phi|_A}(A') < \lambda_\phi(X_\infty)$. Now Lemma 5.12 provides an element of $\mathcal{O}(\Gamma)$ which is displaced less or equal than $\max\{\lambda_{\phi|_A}(A'), \lambda_\phi(X'_\infty) + \varepsilon\}$, contradicting the fact that X_∞ is a minpoint for λ_ϕ in $\mathcal{O}(\Gamma)$. Therefore $\lambda(\phi|_A) = \lambda(X_\infty)$.

By Theorem 7.3 there is $A_\infty \in \overline{\mathcal{O}(A)}^\infty$ such that $\lambda_{\phi|_A}(A_\infty) = \lambda(\phi|_A)$ and which has not jumped in $\mathcal{O}(A)$. Thus A_∞ is obtained, without jumps, from a point $A'_\infty \in \mathcal{O}(A)$ by collapsing a (possibly empty) invariant core sub-graph B . So $A_\infty \in \mathcal{O}(A'_\infty/B)$.

Let Y be a Γ -graph obtained by inserting a copy of A'_∞ in X'_∞ . Let Y' be the graph obtained collapsing B . Y' belongs to the simplicial boundary of Δ_Y and, since A_∞ has no jump, then so does Y' . Now, observe that $Y' \in \mathcal{O}(Y/B)$ and A_∞ is a ϕ -invariant subgraph of Y' so that $Y'/A_\infty = X'_\infty$. Lemma 5.12 provides an element in $Y'_\infty \in \mathcal{O}(Y/B)$, in the same simplex of Y' , which is displaced no more than $\lambda_{\phi|_A}(A_\infty)$ (because $\lambda_\phi(X'_\infty) < \lambda_\phi(X_\infty) = \lambda_{\phi|_A}(A_\infty)$). Now, Y'_∞ is a new minpoint for λ_ϕ with $\dim(\Delta_{Y'_\infty}) > \dim(\Delta_{X_\infty})$ contradicting the maximality hypothesis on X_∞ . It follows that X_∞ is a train track point in $\mathcal{O}(X_\infty)$ as desired. \square

We have just seen that, even if non-jumping min-points are not necessarily partial train tracks, some of them are. Conversely, we see now non-jumping partial train tracks at the bordification are always min-points for λ_ϕ .

Lemma 7.6. *Let $\phi \in \text{Aut}(\Gamma)$ and let $X \in \mathcal{O}(\Gamma)$. If there is k so that there is a constant $A > 0$ such that for any $n \gg 1$*

$$Ak^n \leq \Lambda(X, \phi^n X)$$

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

Proof. This follows from the multiplicative triangular inequality. For any $Y \in \mathcal{O}(\Gamma)$ we have $\Lambda(Y, \phi^n Y) \leq \Lambda(Y, \phi Y)^n$. Define a constant $C = \Lambda(X, Y)\Lambda(Y, X)$ and notice that we also have $C = \Lambda(X, Y)\Lambda(\phi Y, \phi X)$. Then,

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

whence, for any n

$$\left(\frac{k}{\Lambda(Y, \phi Y)}\right)^n \leq \frac{C}{A}.$$

This implies $k \leq \Lambda(Y, \phi(Y))$. By choosing a minimizing sequence of points Y_i we get $k \leq \lambda(\phi)$. \square

Lemma 7.7. *Let $\phi \in \text{Aut}(\Gamma)$. Let $X_\infty \in \overline{\mathcal{O}_{gr}(\Gamma)}$ which has not jumped. Suppose that there is a loop $\gamma \in X_\infty$ and $k > 0$ such that $L_{X_\infty}(\phi^n)(\gamma) \geq k^n L_{X_\infty}(\gamma)$. Then*

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

In particular, if X_∞ is a partial train track for ϕ as an element of $\text{Aut}(X_\infty)$, then it is a min-point for ϕ as an element of $\text{Aut}(\Gamma)$.

Proof. Let $X \in \mathcal{O}_{\text{gr}}(\Gamma)$ so that X_∞ is obtained from X by collapsing a core sub-graph $A \subset X$. Let X_ε be a point of X where $\text{vol}(A) < \varepsilon$. Let γ be as in the hypothesis. For ε small enough we have $L_{X_\varepsilon}(\gamma) \leq 10L_{X_\infty}(\gamma)$, and therefore

$$\Lambda(X_\varepsilon, \phi^n X_\varepsilon) \geq \frac{L_{X_\varepsilon}(\phi^n \gamma)}{L_{X_\varepsilon}(\gamma)} \geq \frac{L_{X_\infty}(\phi^n \gamma)}{10L_{X_\infty}(\gamma)} \geq \frac{k^n L_{X_\infty}(\gamma)}{10L_{X_\infty}(\gamma)} = \frac{k^n}{10}.$$

By Lemma 7.6 we have $\lambda(\phi) \geq k$.

For the second claim it suffice to choose let γ a legal candidate that realizes $\Lambda(X_\infty, \phi X_\infty)$. So $L_{X_\infty}(\phi^n(\gamma)) = \lambda_\phi(X_\infty)^n L_{X_\infty}(\gamma)$.

Hence $\lambda(\phi) \geq \lambda_\phi(X_\infty)$ and since X_∞ has not jumped $\lambda(\phi) \leq \lambda_\phi(X_\infty)$. \square

We are now in position to complete extension of Theorem 5.14 from a local to a global statement (see Theorem 7.4).

Theorem 7.8. *Let $\phi \in \text{Aut}(\Gamma)$. 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

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

Proof. Let $\lambda = \lambda(\phi|_A)$. By Lemma 7.5 and Theorem 4.15, there is $\widehat{A} \in \overline{\mathcal{O}(A)}^\infty$ which is a min-point for $\phi|_A$, which has not jumped in $\mathcal{O}(A)$, and which is a partial train track for $\phi|_A$ as an element of $\text{Aut}(\widehat{A})$. Let f_A be a partial train track map $f_A : \widehat{A} \rightarrow \widehat{A}$ representing $\phi|_A$. Therefore, there is a periodic line γ in \widehat{A}_{max} with legal images in \widehat{A}_{max} and stretched exactly by λ . Let now $\widehat{X} \in \mathcal{O}(\Gamma)$ be obtained by inserting a copy of \widehat{A} in X_∞ . Since \widehat{A} has not jumped in $\mathcal{O}(A)$, then \widehat{X} has not jumped in $\mathcal{O}(\Gamma)$.

Let $f : \widehat{X} \rightarrow \widehat{X}$ be any straight map representing ϕ so that $f|_A = f_A$. Therefore $f_A^n(\gamma)$ is immersed for any n and the length of $f_A^n(\gamma)$ is λ^n times the length of γ . It follows that $L_{\widehat{X}}((\phi^n)\gamma) = \lambda^n(L_{\widehat{X}}(\gamma))$.

By Lemma 7.7 $\lambda(\phi|_A) = \lambda \leq \lambda(\phi)$, and the first claim is proved. Moreover, if $\lambda(\phi|_A) = \lambda_\phi(X_\infty)$, then

$$\lambda(\phi) \leq \lambda(X_\infty) = \lambda(\phi|_A) = \lambda \leq \lambda(\phi)$$

and therefore all inequalities are equalities. Finally, if X has not jumped then $\lambda(X) \geq \lambda(\phi)$ just because this inequality is true by definition for points in $\mathcal{O}(\Gamma)$ and clearly passes to limits of non-jumping sequences. The converse inequality follows from the second claim and Theorem 7.4. \square

Note that Theorem 7.8 implies that *a posteriori* we can remove the non-jumping requirement from Theorem 7.3 and Lemma 7.5.

Corollary 7.9 (Min-points don't jump). *Let ϕ be any element in $\text{Aut}(\Gamma)$. If $X \in \overline{\mathcal{O}(\Gamma)}^\infty$ is such that $\lambda_\phi(X) = \lambda(\phi)$, then it has not jumped.*

Proof. This is a direct consequence of Theorem 7.8. \square

We introduce the notion of partial train track at infinity.

Definition 7.10 (Partial Train track at infinity). Let $\phi \in \text{Aut}(\Gamma)$. The set $\text{TT}^\infty(\phi)$ is defined as the set of points $X \in \overline{\mathcal{O}(\Gamma)}^\infty$ such that X has not jumped, and X is a partial train track point for ϕ in $\mathcal{O}(X)$. (Hence $\lambda_\phi(X) = \lambda(\phi)$ by Lemma 7.7.)

Note that $\text{TT}(\phi) \subset \text{TT}^\infty(\phi)$. The main differences are that $\text{TT}(\phi)$ may be empty, while any ϕ has a partial train track in $\text{TT}^\infty(\phi)$. On the other hand, $\text{TT}(\phi)$ coincides with the set of minimally displaced points, while $\text{TT}^\infty(\phi)$ may be strictly contained in the set of minimally displaced points.

With this definition we can collect some of the above results in the following simple statement, which is a straightforward consequence of Theorems 4.15, 7.3 and Lemmas 7.5, 7.7.

Theorem 7.11 (Existence of partial train tracks at infinity). *For any $[\phi] \in \text{Out}(\Gamma)$, $\text{TT}^\infty(\phi) \neq \emptyset$. For any $X \in \text{TT}^\infty(\phi)$, $\lambda_\phi(X) = \lambda(\phi)$.*

We end this section by discussing some interesting consequences of the theory developed so far. In particular we show that if ϕ is reducible then there is a train track showing reducibility.

Theorem 7.12 (Detecting reducibility). *Let $\phi \in \text{Aut}(\Gamma)$ be reducible. Then there is $T \in \text{TT}^\infty(\phi)$ such that either $T \in \partial_\infty \mathcal{O}(\Gamma)$ or there is an optimal map $f_T : T \rightarrow T$ representing ϕ such that there is a proper sub-graph of T which is f_T -invariant.*

Proof. Since ϕ is reducible there is $X \in \mathcal{O}(\Gamma)$, a straight map $f : X \rightarrow X$ representing ϕ and a proper non trivial sub-graph $A \subset X$ such that $f(A) = A$. We can therefore collapse A and λ won't explode. By Theorem 7.11 there is a partial train track point Z for ϕ in $\overline{\mathcal{O}(X/A)}^\infty$ and a partial train track point Y for $\phi|_A$ in $\overline{\mathcal{O}(A)}^\infty$. If $\lambda_{\phi|_A}(Y) \leq \lambda_\phi(Z)$, then $Z \in \text{TT}^\infty(\phi) \cap \partial_\infty \mathcal{O}(\Gamma)$ and we are done. Otherwise, since Z has not jumped (as a point of $\partial_\infty \mathcal{O}(X/A)$), we can regenerate it to a point $Z' \in \mathcal{O}(X/A)$ with $\lambda_\phi(Z') < \lambda_{\phi|_A}(Y)$. We now apply regeneration Lemma 5.12 to Y and Z' . If $Y \in \partial_\infty \mathcal{O}(A)$, then we get a partial train track for ϕ in $\partial_\infty \mathcal{O}(\Gamma)$. If $Y \in \mathcal{O}(A)$ we get a partial train track for ϕ in $\mathcal{O}(\Gamma)$ admitting Y as an invariant sub-graph. \square

In fact, the proof of Theorem 7.12 proves more: that the set of partial train tracks detect any (maximal) invariant free factor system. Roughly, if A is an invariant free factor system, then there is $B \supseteq A$ (possibly $B = A$) and a partial train track point which shows B as the fundamental group of an invariant sub-graph. The precise statement is the following.

Theorem 7.13 (Strong reformulation of Corollary 7.12). *Let $\phi \in \text{Aut}(\Gamma)$. 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 $W \in \text{Hor}_{\mathcal{O}(\Gamma)}(Z)$ such that the simplex Δ_W contains a minimizing sequence for λ . Moreover if Y is the graph used to regenerate W from Z , then the minimizing sequence can be chosen with straight maps f_i such that $f_i(Y) = Y$ and $\text{Lip}(f_i) \rightarrow \lambda(\phi)$.*

Proof. Follows from the proof of Theorem 7.12 (and Lemma 5.12). \square

Finally, as in the case of irreducible automorphisms, the existence of partial train tracks gives the following fact.

Corollary 7.14. *For any $\phi \in \text{Aut}(\Gamma)$ we have $\lambda(\phi^n) = \lambda(\phi)^n$.*

Proof. This follows from Theorem 7.11 and Lemma 4.12. \square

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 $\text{Out}(F_n)$. 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] Marc Culler and Karen Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), no. 1, 91–119. MR 830040
- [7] Stefano Francaviglia and Armando Martino, *Metric properties of outer space*, Publ. Mat. **55** (2011), no. 2, 433–473. MR 2839451
- [8] ———, *The isometry group of outer space*, Adv. Math. **231** (2012), no. 3-4, 1940–1973. MR 2964629
- [9] Stefano Francaviglia and Armando Martino, *Stretching factors, metrics and train tracks for free products*, Illinois J. Math. **59** (2015), no. 4, 859–899.
- [10] Stefano Francaviglia and Armando Martino, *Displacements of automorphisms of free groups II: Connectedness of level sets*, in preparation.
- [11] 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
- [12] ———, *On fixed points of certain automorphisms of free groups*, Proc. London Math. Soc. (3) **48** (1984), no. 1, 72–90. MR 721773
- [13] ———, *Fixed points of automorphisms of free groups*, Adv. in Math. **64** (1987), no. 1, 51–85. MR 879856
- [14] 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
- [15] 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
- [16] ———, *The free splitting complex of a free group, II: Loxodromic outer automorphisms*, Trans. Amer. Math. Soc. **372** (2019), no. 6, 4053–4105. MR 4009387
- [17] Allen Hatcher, *Homological stability for automorphism groups of free groups*, Comment. Math. Helv. **70** (1995), no. 1, 39–62. MR 1314940
- [18] 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.
- [19] Ilya Kapovich, *Detecting fully irreducible automorphisms: a polynomial time algorithm*, <https://arxiv.org/abs/1609.03820>.
- [20] Ilya Kapovich, *Algorithmic detectability of iwip automorphisms*, Bull. Lond. Math. Soc. **46** (2014), no. 2, 279–290. MR 3194747

- [21] 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
- [22] Martin Lustig, *Conjugacy and centralizers for iwip automorphisms of free groups*, *Geometric group theory*, Trends Math., Birkhäuser, Basel, 2007, pp. 197–224. MR 2395795
- [23] Sebastian Meinert, *The Lipschitz metric on deformation spaces of G -trees*, *Algebr. Geom. Topol.* **15** (2015), no. 2, 987–1029. MR 3342683
- [24] John R. Stallings, *Topology of finite graphs*, *Invent. Math.* **71** (1983), no. 3, 551–565. MR 695906
- [25] 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