ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

The mimimally displaced set of an irreducible automorphism is locally finite

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:
Stefano Francaviglia, A.M. (2020). The mimimally displaced set of an irreducible automorphism is locally finite. GLASNIK MATEMATICKI. SERIJA 3, 55(2), 301-336 [10.3336/gm.55.2.09].

Availability:
This version is available at: https://hdl.handle.net/11585/784942 since: 2023-01-24

Published:
DOI: http://doi.org/10.3336/gm.55.2.09

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (https://cris.unibo.it/).
When citing, please refer to the published version.

(Article begins on next page)

12 October 2024

This is the final peer-reviewed accepted manuscript of:

Francaviglia, S., Martino, A., & Syrigos, D. (2020). The minimally displaced set of an irreducible automorphism is locally finite. Glasnik Matematicki, 55(2), 301-336

The final published version is available online at ***https://dx.doi.org/10.3336/gm.55.2.09***

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (https://cris.unibo.it/)

When citing, please refer to the published version.

THE MINIMALLY DISPLACED SET OF AN IRREDUCIBLE AUTOMORPHISM IS LOCALLY FINITE

Stefano Francaviglia, Armando Martino and Dionysios Syrigos
University of Bologna, Italy, University of Southampton, UK and University of Southampton, UK

Abstract. We prove that the minimally displaced set of a relatively irreducible automorphism of a free splitting, situated in a deformation space, is uniformly locally finite. The minimally displaced set coincides with the train track points for an irreducible automorphism. We develop the theory in a general setting of deformation spaces of free products, having in mind the study of the action of reducible automorphisms of a free group on the simplicial bordification of Outer Space. For instance, a reducible automorphism will have invariant free factors, act on the corresponding stratum of the bordification, and in that deformation space it may be irreducible (sometimes this is referred as relative irreducibility).

1. Introduction

Overview. In this paper we study deformation spaces of marked metric graphs of groups.

Since its first appearance on the scene ([8]), the celebrated Culler-Vogtmann Outer Space became a classical subject of research. It turned out to be a very useful tool for understanding properties of automorphisms of free groups (see for instance [4, 14, 15, 16, 20, 26, 27]). A typical object in the Outer Space of F_n is a marked graph with fundamental group of rank n, and locally Euclidean coordinates are defined by turning graphs into metric graph by an assignment of positive edge-lengths. Outer Space is not compact and there are basically two ways of going to infinity: making the marking diverge or collapsing a collection of sub-graphs of a given element X of Outer Space. The second operation has a local flavour and it is similar to the operation of pinching a curve of a surface. Attaching these "collapsed" points, leads one to define the simplicial bordification of the deformation space. If one starts

2000 *Mathematics Subject Classification.* 20E06, 20E36, 20E08.

with Culler-Vogtmann space, the result is the free splitting complex, which is related to the free factor complex (see for instance [3, 5, 6, 18, 19, 21, 25]).

Collapsing comes naturally into play when one is analysing a reducible automorphism of F_n induced by simplicial map $f : X \to X$ which exhibits an invariant collection of sub-graphs of X.

Once collapsed, X is turned into a graph of groups corresponding to a free splitting of F_n. On the other hand, the collapsed part is not necessarily connected. This phenomenon has led researchers to investigate more general deformations spaces. Namely deformation spaces of (not necessarily connected) graph of groups, possibly with marked points or "hairs" (see for instance [17, 11, 12, 13, 28, 7]).

One of the main tools used to study the action of automorphisms on deformation spaces is the theory of Stalling folds ([30]) and the so-called Lispchitz metric ([9, 10, 28]). In particular, given an automorphism ϕ, one can study the displacement function λ_ϕ defined as $\lambda_\phi(X) = \Lambda(X, \phi X)$ (here Λ denotes the maximal stretching factor from X to ϕX, whose logarithm is the asymmetric Lipschitz metric). Of particular interest is the set $\mathrm{Min}(\phi)$ of minimally displaced points. When ϕ is irreducible, this coincides with the set of points supporting train-track maps ([12]) and its structure is particularly useful for example in building algorithm for decision problems. It is used in [13] for a metric approach to the conjugacy problem for irreducible automorphism of free groups (solved originally in [26]) and the reducibility problem of free groups (solved originally in [23, 24]).

Main results of the paper. If one is interested in effective procedures, one of the main problem is that general deformation spaces have a simplicial structure that is not locally finite. So if one starts from a simplex and wishes to enumerate neighbouring simplices, there is no chance to make this procedure effective.

In this paper we prove that the minset $\mathrm{Min}(\phi)$ for irreducible automorphisms of exponential growth is locally finite; namely given a simplex intersecting $\mathrm{Min}(\phi)$, one can give a finite of its neighbours so that any simplex *not* in that list, does not intersect $\mathrm{Min}(\phi)$. This is the content of our Theorem 6.4. Moreover, it is also uniformly locally finite, Corollary 6.10.

Theorem (Theorems 6.4 and 6.10). *Let G be a group equipped with a free splitting, $\mathcal{G}$. Let ϕ be an automorphism of G which preserves the splitting and is irreducible with $\lambda(\phi) > 1$. Then* $\mathrm{Min}(\phi)$ *- also seen as the points which support train track maps for ϕ - is uniformly locally finite both as a subset of the deformation space $\mathcal{O}(\mathcal{G})$ and its volume 1 subspace, $\mathcal{O}_1(\mathcal{G})$.*

Remark 1.1. We note that the number $\lambda(\phi)$ in the Theorem above is the (minimal) displacement of ϕ relative to the splitting $\mathcal{G}$.

For instance, if one takes a relative train track representative for an automorphism of F_n, then one gets a free splitting of $G = F_n$ by taking the largest invariant subgraph (the union of all the strata except the top one). The resulting automorphism is irreducible in the corresponding relative space, and the number $\lambda(\phi)$ is the Perron-Frobenius eigenvalue of the top stratum.

The application we have in mind for this kind of result is an effective study of the minset for a reducible automorphism.

In general this minset is empty, but starting with a reducible automorphism ϕ of a free group, one can collapse an invariant free factor to obtain a new deformation space on which the automorphism acts. If ϕ is relatively irreducible in that space, then its minset is locally finite. Otherwise, one can keep collapsing free factors until it is relatively irreducible.

It is easy to see that the minset for a reducible automorphism is not locally finite in general, however for any automorphism and any simplex with a given displacement, there are only finitely many possible simple folds which produce simplices of strictly smaller displacement - Corollary 4.15.

The idea behind Theorem 6.4 is the following. For a minimally displaced point X, it is known that folding an illegal turn of an optimal map $f : X \to X$ representing ϕ produces a path in $\mathrm{Min}(\phi)$, called folding path. (See for instance [11, 12]). But it is also clear that there are legal turns that can be folded without exiting $\mathrm{Min}(\phi)$, for instance, this may happen at illegal turns for ϕ^{-1}. The strategy is to understand which legal turns can be folded, and we are able to produce a finite list such that if a turn τ is not in that list, then by folding τ one exits the minimally displaced set. In our terminology, folding a *critical* turn could allow one to remain in the minset, whereas folding a *regular* turn forces one to leave it. One then understands arbitrary neighbouring simplices, by looking at which of them can be reached by a (uniformly bounded) number of critical folds - these are the only ones that may be minimally displaced. However, one complication is that it is possible that a critical fold could *increase* the displacement, and a subsequent critical fold decrease it so that one re-enters the minimally displaced set. Nevertheless, our result produces a finite list containing all neighbouring simplices that are minimally displaced.

Remark 1.2. We have written the paper for deformation spaces of free splittings of G, namely connected graph of groups with trivial edge-groups. However, every result of the paper remains true for deformations spaces of non-connected graph of groups, as developed for instance in [12, 13]. This is because connectedness plays no role in our proofs. (In those papers, non-connectedness was crucial since the main argument was an inductive one.) Nonetheless, we decided to stick to the connected case for the benefit of the reader.

Structure of the paper. We have decided to write the paper in a reverse order; we start immediately with the core of the paper, postponing the section of general definitions to the end. This is because the definitions and terminology we use are quite standard, and the reader used to the subject can start reading directly.

2. Preliminaries

We recall some notation here and, we refer to Section 7 for more details.

Convention 2.1. Deformations spaces - here, of free splittings, even though the concept exists more generally - can be viewed either as spaces of trees, or graphs. We adopt here the graphs-viewpoint, but one can easily pass from one viewpoint to the other by taking universal covers and G-quotients (see below for more details).

Throughout the whole paper, $\mathcal{G}$ will denote a fixed free splitting of the group G. That is, we write $G = G_1 * \ldots * G_k * F_n$, but this need not be the Grushko decomposition of G. In fact, in the examples we have in mind, G is a free group, and the free factors G_i correspond to a collection of invariant free factors under some automorphism of G. $\mathcal{O}(\mathcal{G})$ will denote the deformation space of a free splitting of a group G.

The typical object $X \in \mathcal{O}(\mathcal{G})$ is therefore a marked metric graph of groups, with trivial edge groups, and whose valence one or two vertices have non-trivial vertex group. (One can also think of X as a G-tree with trivial edge stabilisers, where the vertex stabilisers are precisely the conjugates of the G_i. Elements of G which fix some vertex are called elliptic, and the others are hyperbolic.) Note that every $X \in \mathcal{O}(\mathcal{G})$ has the same elliptic elements (and this characterises the points in the space). For a vertex $v \in X$ we denote G_v its vertex group. If G_v is trivial, then v is said to be *free*.

These spaces - $\mathcal{O}(\mathcal{G})$ - naturally occur in the bordification of classical Culler-Vogtmann Outer Space, on collapsing invariant subgraphs.

We denote by $\mathrm{Aut}(\mathcal{G})$ the group of automorphisms of G which preserve the splitting; that is, each G_i in the splitting is sent to a conjugate of another (possibly the same) G_j. That is, $\mathrm{Aut}(\mathcal{G})$ the group of automorphisms of G which preserve the elliptic elements. Similarly, $\mathrm{Out}(\mathcal{G}) = \mathrm{Aut}(\mathcal{G})/\mathrm{Inn}(G)$.

2.1. *Graph of Groups and G-trees.* We recall some basic notions of Bass-Serre theory. The main references for this section are [1] and [29]. Given a graph Γ, we denote by $V(\Gamma)$ the set of vertices of Γ and by $E(\Gamma)$ the set of (oriented) edges of Γ. If e is an (oriented) edge, we denote by $\iota(e)$, the initial vertex of e and by $\tau(e)$, the terminal vertex of e.

Definition 2.2 (Graph of Groups). *A graph of groups X consists of a connected graph Γ together with groups G_v for every vertex $v \in V(\Gamma)$ and edge groups $G_e = G_{\bar{e}}$, and monomorphisms $\alpha_e : G_e \to G_{\tau(e)}$ for every (oriented) edge $e \in E(\Gamma)$.*

In this paper, we work with free products of groups, which by Bass-Serre theory, arise as the fundamental groups of graph of groups with trivial edge groups. For the remainder of the section, we suppose that all the graph of groups have trivial edge groups. In this case, we can see a graph of groups X, as a pair which consists of a graph Γ and a collection of groups G_v, one for each v vertex of Γ; we simply write $X = (\Gamma, \{G_v\}_{v\in V(\Gamma)})$. We refer to Γ, as the topological space (or the graph) associated to X.

Let $X = (\Gamma, \{G_v\}_{v\in V(\Gamma)})$ be a graph of groups. An *edge- path* P (of combinatorial length k) in X is a sequence of the form $(g_1, e_1, g_2, e_2, \dots, g_k, e_k, g_{k+1})$ where the e_i's are oriented edges of Γ such that the terminal point of each e_i is the initial point of e_{i+1}, and each g_i is a group element from the vertex group based at the initial point of e_i; in this case, we simply write $P = g_1e_1 \dots g_ke_kg_{k+1}$. In the present work, we also use "paths" which are not edge-paths, as their endpoints may not be vertices. The previous definition is extended from edge-paths to *paths*, by allowing the first and the last segments e_1, e_k of P to be partial edges (note that if $\iota(e_1)$ is not a vertex, then we set g_1=1 and, similarly, if $\tau(e_k)$ is not a vertex then we set $g_{k+1} = 1$). A path is called *loop*, if the endpoint of its last edge coincides with the initial point of its first edge. We say that a path P, as above, is *reduced*, if whenever $e_i = \bar{e}_{i+1}$, $i = 1, \dots, k-1$, then g_i is not the trivial group element. Furthermore, we say that a loop $P = g_1e_1g_2e_2 \dots g_ke_kg_{k+1}$ is *cyclically reduced* if it is reduced and if $e_k = \bar{e}_1$ then $g_{k+1}g_1 \neq 1$.

We can always represent the conjugacy class of a loop in the graph of groups in the form $g_1e_1g_2e_2 \dots g_ke_k$; in this case, being cyclically reduced means that whenever $e_i = \bar{e}_{i+1}$, then g_i is not the trivial group element, with the subscripts taken modulo k.

We say that two paths are *equivalent* if one can be obtained from the other by a sequence of insertions or deletions of inverse-pairs. We then denote by $\pi(X)$, the set of equivalence classes of edge paths.

Then $\pi(X)$ has the structure of a groupoid, whose operation is the concatenation of paths. Note that $\pi(X)$ is not a group in general, for if $p, q \in \pi(X)$, then the concatenation $p * q = pq$ is only defined exactly when the terminal vertex of p is the initial vertex of q. In our setting, every path p is equal (in $\pi(X)$) to a unique reduced path (with the same endpoints), which we denote by $[p]$.

(Alternatively, we can take $\pi(X)$ to be the set of reduced edge-paths, under the operation of concatenation followed by reduction.)

If a, b are two points of Γ (not necessarily vertices), we can define the set $P[a, b]$, of paths in X, from a to b. If a, b are vertices, then $P[a, b]$ can be seen as a subset of $\pi(X)$. In the special, case where $a = b$, the set of loops based at a, $P[a, a]$, form a subgroup of $\pi(X)$ (as any two loops based at a, can be concatenated, inducing a loop based at a).

Definition 2.3 (Fundamental group). *Let $X = (\Gamma, \{G_v\}_{v\in V(\Gamma)})$ be a graph of groups. If u is a vertex of Γ, then the fundamental group $\pi_1(X, u) = P[u, u]$ of Γ is the set of loops in Γ based at u (as elements of $\pi(X)$). In our case,*

$$\pi_1(X, u) = (*_{v\in V(\Gamma)} G_v) * \pi_1(\Gamma, u).$$

Note that the isomorphism type of $\pi_1(X, u)$ does not depend on the choice of u. If u, w are vertices of Γ, then $\pi_1(X, u), \pi_1(X, w)$ are conjugate in $\pi(X)$ (and we simply write $\pi_1(X)$, when we ignore the base point).

We can now define the universal cover (or Bass-Serre tree) of a graph of groups.

Definition 2.4 (Universal Cover). *Let $X = (\Gamma, \{G_v\}_{v\in V(\Gamma)})$ be a graph of groups. If v is a vertex of Γ, then we define the tree $T = \widetilde{(X, v)} = \widetilde{X}$, as follows:*

- *The set of vertices $V(T)$ are exactly the "cosets" pG_w, $w \in V(\Gamma)$, where p is a reduced edge- path from w to v.*
- *There is an (oriented) edge from $p_1G_{w_1}$ to $p_2G_{w_2}$, if and only if there are elements $g_i \in G_{w_i}, i = 1, 2$ and an edge $e \in E(\Gamma)$, so that $p_2 = p_1g_1eg_2$.*

Remark 2.5. (i) As every element of $G = \pi_1(X, v)$ is a loop based at v, there is natural simplicial (left) action of the fundamental group on the universal cover T.

(ii) If w is a vertex of Γ, then there is a unique (reduced) path p_w of Γ_V, from v to w. Then the stabiliser $Stab_T(pG_w) = pG_wp^{-1}$

(iii) The stabilisers of edges are, by construction, trivial.

(iv) One can define arbitrary points of T in exactly the same way, by allowing the paths p to be general paths (not just edge-paths) from v to w, with the convention that G_w is trivial if w is not a vertex.

Given a graph of groups with trivial edge groups, we described the action of its fundamental group on a tree with trivial edge stabilisers (and vertex stabilisers which coincide with the set of conjugacy classes of the vertex groups).

Conversely, if a group G acts on a tree with trivial edge stabilisers, we can construct a graph of groups with fundamental group G.

Definition 2.6 (Quotient graph of groups). *Let's suppose that G acts on a tree T with trivial edge stabilisers. Then we can define the quotient graph of groups X, by taking the quotient graph $\Gamma = G/T$ and as vertex groups the stabilisers of some orbit of each vertex (and trivial edge groups). In this case, $\pi_1(X)$ is (isomorphic to) G.*

Remark 2.7. By Bass-Serre theory, we know that these two constructions (graph of groups and the universal cover) are equivalent and we can go from the level of the tree to the level of the graph of groups and vice versa. (We

have done some choices of fundamental domains which are inessential, as they produce isomorphic structures, every time.)

Now we restrict to graph of groups in the same relative outer space, as in our Convention. More specifically, we assume that our graph of groups, will be $\mathcal{G}$-graph of groups (see Definition 7.3, for more details), for the fixed splitting $\mathcal{G}$ of our group G.

As in our description for the $\mathcal{G}$-trees (see 7), every $\mathcal{G}$-graph as a point (simplex) of $\mathcal{O}(\mathcal{G})$, is equipped with a marking.

Definition 2.8. *A $\mathcal{G}$-graph dual to a $\mathcal{G}$-tree. Namely, it is a finite connected G-graph of groups X, along with an isomorphism, $\Psi_X : G \to \pi_1(X)$ - a marking - such that:*

- *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, via Ψ_X, to $\mathcal{G}$. That is, Ψ_X restricts to a bijection from the conjugacy classes of the G_i to the vertex groups of X.*

The following definition is not the definition of the morphism that is given by Bass in [1], as he requires the map to be a graph morphism (sending edges to edges and vertices to vertices). However, in our case, his method can be easily adjusted for the more general maps that we define.

Definition 2.9 (Maps Between Graphs of Groups). *Let $X = (\Gamma, \{G_v\})$ and $Y = (\Gamma', \{H_w\})$ be two (marked) $\mathcal{G}$-graphs.*

A map F between X, Y consists of:

1. *Two maps f_V, f_E:*
   (a) *$f_V : V(\Gamma) \to \Gamma'$ and*
   (b) *For each edge e with $\iota(e) = v, \tau(e) = v'$, $f_E(e)$ is a path from $f_V(v)$ to $f_V(v')$.*
2. *Isomorphisms between the vertex groups: $\phi_v : G_v \to H_{f_V(v)}$ (by abusing the notation, we set $H_{f_V(v)} = 1$, if $f_V(v)$ is not a vertex).*

Note that after sub-dividing some edges of the co-domain, we can suppose that our maps are simplicial. For any two $\mathcal{G}$-graphs X, Y, a map $F : X \to Y$, induces a (natural) homomorphism $\Phi_F : \pi(X) \to \pi(Y)$ between the set of paths, which restricts to a homomorphism $\Phi_F : \pi_1(X) \to \pi_1(Y)$ between the fundamental groups (well defined up to composition with inner automorphisms). The maps that we define in the following definition, play the role, at level of $\mathcal{G}$-graphs, of G-equivariant maps between $\mathcal{G}$-trees.

Definition 2.10. *We say that a map $F : X \to Y$ is a $\mathcal{G}$-map, if the induced homomorphism Φ_F is the "change of markings", i.e. $\Phi_F \Psi_X = \Psi_Y$ (up to composition with inner automorphisms).*

Remark 2.11. (i) By Bass-Serre theory, it's clear now that the notion of $\mathcal{G}$-maps at the level of marked $\mathcal{G}$- graphs is equivalent to the notion of G-equivariant "straight" maps at the level of $\mathcal{G}$-trees - that is, equivariant, surjective, continuous maps that are determined by the images of vertices by extending linearly over edges. In fact, we can move from one to the other by simply considering lifts and projections, using Bass-Serre theory (even if the lifts and the projections are not unique, as they depend on the choice of some fundamental domain, all of them are equivalent).

(ii) Note that we cannot talk about continuity between graph of groups, but we could alternatively consider the equivalent notion of graph of spaces (we simply change the vertex groups G_i with a topological space Γ_i, with $\pi_1(\Gamma_i) = G_i$). In that case, the maps between two graph of spaces would be continuous.

2.2. *Basic Notation.*

Notation 2.12. *We will use the following standard notation:*

- $\mathcal{O}_1(\mathcal{G})$ *the volume one subspace of* $\mathcal{O}(\mathcal{G})$.
- Δ *denotes an open simplex of* $\mathcal{O}(\mathcal{G})$.
- Δ_X *denotes the simplex with underlying graph of groups* X.
- $\bar{e}$ *denotes the inverse of oriented edge* e. *Same notation for paths.*
- $\gamma \cdot \eta$ *denotes concatenation of paths.*
- $L_X(\gamma)$ *denotes the reduced length in* X *of a loop* γ, *if* X *is seen as a* G*-graph.*
- *Folding a turn* $\{a, b\}$ *by an amount of* t *means identify initial segments of* a *and* b *of length* t. *This is always well-defined for small enough* t.
- *Given an automorphism* $\phi \in \mathrm{Out}(\mathcal{G})$, $\lambda_\phi : \mathcal{O}(\mathcal{G}) \to \mathbb{R}$ *denotes the displacement function* $\lambda_\phi(X) = \Lambda(X, \phi X)$ *(this is well defined as the inner automorphisms act trivially). For a simplex* Δ *we set* $\lambda_\phi(\Delta) = \inf_{X \in \Delta} \lambda_\phi(X)$; *we set* $\lambda(\phi) = \inf_{X \in \mathcal{O}(\mathcal{G})} \lambda_\phi(X)$.
- *An* $\mathcal{O}$*-map between elements of* $\mathcal{O}(\mathcal{G})$ *is a map that realises the difference of markings. A straight map between elements of* $\mathcal{O}(\mathcal{G})$ *is an* $\mathcal{O}$*-map with constant speed on edges. (See the definitions section on page 30).*

Since we decided to adopt the graphs-viewpoint, some words of explanation are needed about turns. A turn at a non-free vertex v of X is given by the equivalence class of unoriented pair $\{g_1e_1, g_2e_2\}$, where e_1, e_2 are (germs of) oriented edges with the same initial vertex, v; g_1, g_2 are elements in the vertex-group G_v, and the equivalence relation is given by the diagonal action of G_v. If we think X as a G-tree, a turn at a non-free vertex v, is a G_v-orbit of an unoriented pair of edges $\{g_1e_1, g_2e_2\}$ where $g_1, g_2 \in G_v$ and e_1, e_2 are oriented edges of X emanating at v. In other words, the projection of a turn

in a tree is a turn in the quotient graph of groups, and, conversely, the lift of a turn in the graph of groups is a turn in the universal cover. In any case, we denote the turn given by the class of $\{a, b\}$ simply $[a, b]$. (Note that $[a, b] = [b, a] = [ga, gb]$)

For the convenience of the reader, in this section we give the definitions for both points of view, by thinking X both as $\mathcal{G}$-graph of groups and as $\mathcal{G}$-tree. More specifically, everything in this section is written, the graphs-viewpoint and we translate into the language of trees, if needed. In the rest of the paper, we use these two notions interchangeably, as everything can be easily translated between these two point of views.

Definition 2.13. *Let $X \in \mathcal{O}(\mathcal{G})$. A turn $[x, x]$ is said* **trivial**. *A turn $\tau = [a, gb]$ at a vertex v of X is called* **degenerate**, *if $a = b$ (resp., if a, b are in the same G_v-orbit); it is called non-degenerate otherwise. If e is an edge starting* and *ending at v (resp., if both endpoints of e are in the same G-orbit, as v), it determines a non-degenerate turn at v.*

Definition 2.14. *Given a straight map $f : X \to Y$, at the graph level, we say that f maps the turn $[a, gb]$ to the turn $[c, hd]$, if the initial paths of (combinatorial) length 1 in the graph of groups (resp., if the initial edge or germ of edge), of $f(a)$ and $f(gb)$ are c and hd (in some order).*

We will sometimes abuse notation and say that f maps the turn $[a, gb]$ to the turn $[f(a), f(gb)]$, even though we really mean this to be the initial sub-paths of combinatorial length 1 of the (in general) paths given.

We say that $[a, gb]$ is f-legal, if f maps $[a, gb]$ to a non-trivial turn. If f maps either of a or b to a vertex, then we say the turn is illegal.

If, moreover, $X = Y$, then we say that $[a, gb]$ is $\langle \sim_{f^k} \rangle$-legal if f^k maps $[a, gb]$ to a non-trivial turn for all integers $k \geq 1$.

Lemma 2.15. *Let $X \in \mathcal{O}(\mathcal{G})$ and $\tau = [a, gb]$ be a non-degenerate turn at a vertex v. Then (equivariantly) folding a and gb gives a new element in $\mathcal{O}(\mathcal{G})$.*

Proof. The proof is straightforward and left to the reader. □

Remark 2.16. Note that if e is an edge emanating from v, a non- free vertex of $X \in \mathcal{O}(\mathcal{G})$, and if $g \in G_v$ is such that $< g > \neq G_v$, then by folding a degenerate turn $[e, ge]$ we obtain a tree with non-trivial edge groups(resp., stabilisers). Namely, the new edge e emanating from v has edge group $< g >$ (resp., stabiliser).

Therefore, in practice, a non-degenerate turn is the same as a "foldable" turn.

Remark 2.17. Suppose that $X \in \mathcal{O}(\mathcal{G})$ and $f : X \to X$ is a straight $\mathcal{O}$-map. By this we mean that there are two G markings on X (resp., G-actions), and f is $\mathcal{G}$-map (resp., G-equivariant) with respect to the two different markings (resp., actions) (both of which lie in the same deformation space, and hence have the same elliptic elements).

We will describe the change of action in more detail (at graph level, we have a very similar description, by changing the marking, instead of changing the action). Let's give these names; we will denote the first action by $\cdot$ and the second action by $\star$. Then we will always have that there is an element, $\phi \in \mathrm{Aut}(\mathcal{G})$ such that,

$$f(g \cdot x) = g \star f(x) = \phi(g) \cdot f(x).$$

Then, on iterating f, we get,

$$f^r(g \cdot x) = \phi^r(g) \cdot f^r(x).$$

Note that if e is an edge emanating from a vertex v, then for every $\mathcal{G}$-map f (resp., G-equivariant map), the degenerate turn $[e, ge]$ is f-legal for any $g \neq Id \in G_v$, as long as f does not map e to a vertex.

Lemma 2.18. *Let $X, Y \in \mathcal{O}(\mathcal{G})$ and $f : X \to Y$ be a straight $\mathcal{O}$-map. Let v be a vertex of X and G_v its stabiliser. Let $\tau = [a, gb]$ be a turn at v, such that neither $f(a)$ nor $f(b)$ is a single vertex. If τ is f-illegal, then for any $g' \neq g \in G_v$ the turn $[a, g'b]$ is f-legal. If $X = Y$ and τ is $\langle \sim_{f^k} \rangle$-illegal, then for any $g' \neq g \in G_v$ the turn $[a, g'b]$ is $\langle \sim_{f^k} \rangle$-legal.*

Proof. Let's prove the second claim first. Since τ is $\langle \sim_{f^k} \rangle$-illegal, then there is some power $r \geq 1$ so that $f^r(a)$ and $\phi^r(g) f^r(b)$ are the same germ - we are using the automorphism ϕ as in Remark 2.17. It follows that $[f^r(a), \phi^r(g') f^r(b)]$ is degenerate and legal. Since f is an $\mathcal{O}$-map, it follows that $f^{r+l}(\tau) = [f^{r+l}(a), \phi^{r+l}(g') f^{r+l}(b)]$ is degenerate and legal for any $l \geq 0$. Since f-images of illegal turns are f^n-illegal for any n, then $[f^m(a), \phi^m(g') f^m(b)]$ is legal also for $m \leq r$.

First claim now follows by exactly the same argument with $r = 1$ and $l = 0$. ∎

Definition 2.19. *Let $X \in \mathcal{O}(\mathcal{G})$ and let τ be a turn of X. For any loop γ in X (resp., a path with endpoints which lie in the same G-orbit) we denote by*

$$\#(\gamma, \tau)$$

the number of times that the **cyclically reduced** *representative of γ crosses τ. We recall that τ is not an oriented object, so we do not take in account crossing directions.*

The following lemma is almost tautological, but important for our purposes.

Lemma 2.20. *Let $X, Y \in \mathcal{O}(G)$ and $f : X \to Y$ any straight map. If γ is a f-legal path in X, then it is reduced.*

Proof. If γ is not reduced, then it contains a sequence $\bar{e}e$, hence a turn of the kind $[x, x]$. That turn cannot be f-legal. ∎

Definition 2.21. *Let $X \in \mathcal{O}(\mathcal{G})$, and τ be a turn of X. We say that τ is* **(non-) free** *if it is based at a (non-) free vertex. We say that τ is* **infinite non-free** *if it is based at a vertex with infinite vertex group (resp., infinite stabiliser). We say that τ is* **finite non-free** *if it is non-free and based at a vertex with finite vertex group (resp., finite stabiliser). Given an invariant sub-graph $Y \subseteq X$ (resp., G- sub-forest), we say that τ is in Y if both germs of edges of τ belong to Y.*

3. Unfolding projections and local surgeries on paths

Suppose Δ is a simplex in $\mathcal{O}(\mathcal{G})$, with underlying graph of groups X. Let τ be a non-degenerate turn in X. We denote by Δ_τ the simplex obtained by (equivariantly) folding τ. If τ is free and trivalent then Δ_τ trivially equals Δ. Otherwise, Δ is a codimension-one face of Δ_τ. In the latter case there is a natural projection $\Delta_\tau \to \Delta$ corresponding to the collapse of the newly created edge. Rather, we will use the **unfolding projection**, which is defined as follows.

Given $Y \in \Delta_\tau$, we will define lengths of edges of X so that isometrically folding τ eventually produces Y. Let e_1, e_2 be the edges defining τ (possibly $e_1 = e_2$ is τ arises at an edge-loop) and let e be the extra edge added in Δ_τ after folding τ.

Firstly, every edge of X, different to e_1, e_2, will have the same length as its length in Y. Then, for $i = 1, 2$ we set the length of e_i to be $L_Y(e_i) + L_Y(e)$ if $e_1 \neq e_2$, and $L_Y(e_i) + 2L_Y(e)$ if $e_1 = e_2$. We denote the resulting metric graph, which is an element of Δ, by $\mathrm{unf}_\tau(Y)$ and we say that it is obtained by **unfolding** τ. The map

$$\mathrm{unf}_\tau : \Delta_\tau \to \Delta$$

is our unfolding projection.

Lemma 3.1. *The map unf_τ is surjective. Moreover, Folding τ by an amount of $L_Y(e)$ produces a simplicial segment from $\mathrm{unf}_\tau(Y)$ to Y.*

Proof. The proof immediately follows from the construction. ∎

We describe now local surgeries on paths. As above, let Δ be a simplex of $\mathcal{O}(\mathcal{G})$ with underlying graph X. Since we adopt the graphs-viewpoint, then we may view G as the fundamental group of the graph of groups given by X.

If $g_1 e_1 g_2 e_2 \dots g_k e_k$ is a loop in the graph of groups, then it crosses k turns (including multiplicity); each sub-path of the form $e_i g_{i+1} e_{i+1}$ determines a turn, $[\bar{e}_i, g_{i+1} e_{i+1}]$, where the indices are taken modulo k. (The specific metric on X is not relevant for this discussion, merely the fact that we have a way of representing elements/conjugacy classes as loops in the underlying graph of groups for X.)

Thus a path (or loop) is reduced (cyclically reduced) if the turns it crosses (cyclically crosses) are all non-trivial.

With this description, we may modify any given path by replacing one of the g_i with some other element g in the same vertex group. The turns crossed by this new path are exactly the same as the original, except for one turn $\tau = [\bar{e}_{i-1}, g_i e_i]$ which is replaced with $[\bar{e}_{i-1}, g e_i]$. We denote the modified turn and loop respectively

$$\tau_g \qquad \text{and} \qquad \gamma_{\tau,g}.$$

Putting everything in formulas we have:

LEMMA 3.2 (Turn-surgery of paths). *Let Δ be a simplex of $\mathcal{O}(\mathcal{G})$, $\gamma = g_1 e_1 \dots g_k e_k$ a cyclically reduced loop realised in the underlying graph of groups and $\tau = [\bar{e}_{i-1}, g_i e_i]$ be a turn crossed by γ. Let v be the initial vertex of e_i. Then, if τ is non-degenerate, for any $g \neq g_i \in G_v$ the loop $\gamma_{\tau,g}$ is cyclically reduced and satisfies*

$$\#(\gamma_{\tau,g}, \tau') = \begin{cases} \#(\gamma, \tau') & \text{if } \tau' \neq \tau, \tau_g \\ \#(\gamma, \tau') - 1 & \text{if } \tau' = \tau \\ \#(\gamma, \tau') + 1 & \text{if } \tau' = \tau_g \end{cases}$$

Moreover, if τ is degenerate (hence $e_{i-1} = \bar{e}_i$), than the same is true if in addition we choose $g \neq id$.

PROOF. Since γ is cyclically reduced and τ is not degenerate, then $\gamma_{\tau,g}$ is reduced. The same holds true if τ is degenerate and $g \neq id$. The claim now easily follows by counting the number of times that a turn appears along $\gamma_{\tau,g}$. □

We introduce also a second surgery on paths. Let $\gamma = g_1 e_1 \dots g_k e_k$ denote a loop as above. Let $e = e_i$ be an oriented edge crossed by γ at least twice and let j be the next index so that $e_j = e$. We can therefore form the loop $g_j e_i \dots g_{j-1} e_{j-1}$ (note that the formed loop starts with the group element g_j instead of g_i, as in this case any turn which is crossed by the this new loop, was seen as a turn crossed by γ). We refer to such procedure as *edge-surgery*, and denote the resulting loop by

$$\gamma_e.$$

Note that every turn (cyclically) crossed by γ_e is also a turn crossed by γ, so if γ is cyclically reduced, then γ_e is cyclically reduced as well. By construction, γ_e crosses the oriented edge e only once. Still, it may cross $\bar{e}$ and other edges multiple times.

LEMMA 3.3 (Edge-reduction of loops). *Let Δ be a simplex of $\mathcal{O}(\mathcal{G})$, $\gamma = g_1 e_1 \dots g_k e_k$ a cyclically reduced loop realised in the underlying graph of groups. Then for every e_i there is a cyclically reduced loop γ', obtained by recursive edge-surgeries on γ, such that first, γ' crosses e_i, and second,*

γ' crosses every oriented edge at most once. (Possibly $\gamma' = \gamma$ if γ had those properties).

Proof. For a loop η set $n(\eta)$ the total number of repetitions (counted with multiplicity) of oriented edges. So η crosses any oriented edge at most once if and only if $n(\eta) = 0$. If $n(\gamma) > 0$, then there is $g_j e_j \dots g_i e_i \dots g_l e_l$ a sub-path of γ containing e_i and so that $e_j = e_l$ (indices are taken cyclically). The loop γ_{e_j} contains e_i and $n(\gamma_{e_j}) \leq n(\gamma) - 1$. We conclude by arguing inductively as $n(\gamma)$ is strictly decreasing under edge-surgeries. ∎

There is a version of the previous lemma for turns:

Lemma 3.4. *Let Δ be a simplex of $\mathcal{O}(\mathcal{G})$, $\gamma = g_1 e_1 \dots g_k e_k$ a cyclically reduced loop realised in the underlying graph of groups. If $\tau = [e, ge']$ is a non-trivial turn, which is crossed by γ, then we can find some cyclically reduced loop γ', obtained by recursive edge surgeries on γ, such that first, γ' crosses τ and second, γ' crosses every oriented edge at most once.*

Proof. Let $\gamma = g_1 e_1 \dots g_k e_k$ be a cyclically reduced loop, as above. Without loss of generality, we can assume that γ is of the form $\gamma = ge' \dots \bar{e}$ and there are no other occurrences of e' or $\bar{e}$ in γ, as otherwise we can preform edge-surgeries to change γ to a cyclically reduced loop satisfying this property and which still crosses τ (cyclically).

Now suppose that there is some oriented edge E which is crossed by γ at least twice. In this case, if e_i and e_j are the first and the last occurrences of E in γ, respectively, then we replace γ with the cyclically reduced loop $\gamma_1 = ge' \dots g_{i-1} e_{i-1} g_j e_j \dots g_k e_k$ which still crosses τ and, in addition, crosses E once. By arguing inductively on the number of repetitions, we can find a γ' with the requested properties. ∎

4. Critical and Regular turns

Firstly we explain our strategy. Given $X \in \mathcal{O}(\mathcal{G})$ which is minimally displaced by an automorphism ϕ, we want to control the number of ways we can fold a turn of X, without exiting $\mathrm{Min}(\phi)$. If a straight map $f : X \to X$ representing ϕ sends an edge of a maximally stretched loop γ across a turn τ, then by folding τ we decrease the length of $f(\gamma)$. "Morally", this is the only way we can decrease stretching factors of loops, and if we fold a loop *not* in the image of an edge, we increase the displacement. "Morally" does not mean "literally", and in fact one has to (focus on legal loops in tension graph, and) analyse what happens to the images of turns. Our plan is to select a finite number of turns that will be enough to control the displacement. These will be our "simplex critical turns" that we introduce at the end of this section. The upshot of this process will be that the folding of simplex regular (i.e. non-critical) turns, strictly increases the displacement. We note that our set of critical turns won't be optimal, in the sense that we may a priori increase

the displacement also by folding a critical turn; for instance we include all free turns for convenience.

It would be interesting to have a nice characterisation of exactly those turns whose folding do not increase the displacement. The next lemma is the key observation we begin with.

LEMMA 4.1. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$. Let Δ be a simplex of $\mathcal{O}(\mathcal{G})$ and f be an optimal map representing ϕ on a point X of Δ. Let τ be a non-degenerate turn, and let Δ_τ be the simplex obtained by folding τ. Let X^t denote the point of Δ_τ obtained from X by folding τ by an amount t.*

If there is an f-legal loop γ in the tension graph of f (see Definition 7.10) such that

$$(\heartsuit) \qquad \#(f(\gamma),\tau) \leq \lambda_\phi(X)\#(\gamma,\tau) \qquad \textit{(resp. with strict inequality)}$$

then

$$\lambda_\phi(X^t) \geq \lambda_\phi(X) \qquad \textit{(resp. with strict inequality).}$$

PROOF. For any legal loop γ in the tension graph of X, in X^t we have

$$\lambda_\phi(X^t) = \sup_g \frac{L_{X^t}(f(g))}{L_{X^t}(g)} \geq \frac{L_{X^t}(f(\gamma))}{L_{X^t}(\gamma)} = \frac{\lambda_\phi(X)L_X(\gamma) - 2t\#(f(\gamma),\tau)}{L_X(\gamma) - 2t\#(\gamma,\tau)}.$$

and $\lambda_\phi(X^t) \geq \lambda_\phi(X)$ is guaranteed (with strict inequality) provided that

$$\lambda_\phi(X)L_X(\gamma) - 2t\#(f(\gamma),\tau) \geq \lambda_\phi(X)(L_X(\gamma) - 2t\#(\gamma,\tau))$$

(resp. with strict inequality), and that last inequality clearly reduces to ($\heartsuit$). □

What we will do from now on is showing that, except for finitely many turns, we can guarantee the existence of a loop γ satisfying the hypothesis of Lemma 4.1.

We now make a choice of a single non-trivial element $h_v \in G_v$, for each non-trivial G_v. Some of our subsequent constructions will be dependent on this choice, but we will never need to revise this choice so we will not need to refer to the specific elements. We denote the collection of such chosen elements by H:

$$H = \{h_v : v \text{ is a non-free vertex}\}.$$

DEFINITION 4.2. *For any simplex Δ of $\mathcal{O}(\mathcal{G})$ define a set of loops, A_Δ as follows: a cyclically reduced loop $g_1e_1 \dots g_ke_k$ in the underlying graph of Δ is in A_Δ if and only if*

1. *it crosses every (un-oriented) edge at most 4 times, and*
2. *every non-trivial g_i belongs to H.*

REMARK. *The reason for constructing A_Δ is that it is finite, and gives us a local coordinate system of loops which will be sufficient for calculating displacements and the Lipschitz metric, locally.*

Definition 4.3. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$. For a simplex, Δ, in the underlying graph of Δ we say that a turn τ is* **candidate regular** *if it is infinite non-free and $\#(\phi(\gamma), \tau) = 0$ for all loops in A_Δ; a turn is* **candidate critical** *if it is not regular (so τ is critical if either its vertex group is finite or if it appears in $\phi(A_\Delta)$). We denote the set of candidate critical turns of Δ by $\mathcal{C}_C(\Delta)$. (We remark that even if we do not refer to ϕ in the notation, the set $\mathcal{C}_C(\Delta)$ depends on ϕ).*

Lemma 4.4. *Let $X, Y \in \mathcal{O}(\mathcal{G})$ and $f : X \to Y$ a straight map. Let $\Delta = \Delta_X$. Suppose ξ is either an edge or a free turn of X, which is crossed by an f-legal loop γ_0. Then ξ is also crossed by a f-legal loop $\gamma \in A_\Delta$ which additionally crosses any oriented edge at most once. If γ_0 is in the tension graph, then so is γ.*

Moreover, under the same hypotheses, if additionally $X = Y$ and γ_0 is $\langle \sim_{f^k} \rangle$-legal, then γ may also be chosen to be $\langle \sim_{f^k} \rangle$-legal.

Proof. Let γ_0 be a legal loop crossing ξ. By Lemma 2.20 γ_0 is cyclically reduced. By Lemmas 3.3 or 3.4, as appropriate, we can reduce γ_0, via edge-surgeries, to a loop γ_1, still crossing ξ, and which crosses any oriented edge at most once. In particular, γ_1 satisfies condition (1) for belonging to A_Δ.

Since γ_1 is obtained from γ_0 by edge-surgeries, the turns (cyclically) crossed by γ_1 are also crossed by γ_0. Hence if γ_0 is f-legal (respectively $\langle \sim_{f^k} \rangle$-legal), then so is γ_1.

We now perform turn surgeries on γ_1 to produce a loop in A_Δ. Condition (1) of Definition 4.2 is already satisfied, so we only need to concern ourselves with condition (2), which is about the non-free turns crossed by the loop.

Suppose that γ contains a sub-path at a non-free vertex, v, $e_{i-1} g_i e_i$, crossing the corresponding non-free turn, $[\bar{e}_{i-1}, g_i e_i]$. Let $h \in H$ be the corresponding group element of G_v. Then by Lemma 2.18, at least one of $[\bar{e}_{i-1}, e_i]$ and $[\bar{e}_{i-1}, h e_i]$ is f-legal (respectively $\langle \sim_{f^k} \rangle$-legal).

Therefore, by making appropriate choices at each non-free turn crossed by γ_1, we can perform a sequence of turn surgeries to produce an f-legal loop γ (respectively $\langle \sim_{f^k} \rangle$-legal) which is in A_Δ and still crosses ξ (since ξ is unaffected by turn surgeries).

Moreover, γ contains only edges that were originally edges of γ_0, so if γ_0 is in the tension graph, so is γ. ∎

Remark 4.5. Note that in the previous result, we prove that the path γ crosses each oriented edge at most one, hence each un-oriented edge at most twice, even though the requirement for being in A_Δ is that it crosses each un-oriented edge at most 4 times. The reason is that we use two such loops in the following Lemma; we contruct the loops in Lemma 4.6 by using two loops from Lemma 4.4.

LEMMA 4.6. *Let $X, Y \in \mathcal{O}(\mathcal{G})$ and $f : X \to Y$ a straight map. Let $\Delta = \Delta_X$. Let $\tau = [a, gb]$ be a non-free f-legal turn so that both edges a, b are crossed by f-legal loops. Then, there exists a f-legal loop γ which crosses τ. If the loops for a and b are in the tension graph, then so is γ.*

Moreover, we could take $\gamma = \gamma'_{\tau,g_1}$, for some $g_1 \in G_v$ where $\gamma' \in A_\Delta$. Finally, the same is true for $\langle \sim_{f^k} \rangle$-legality in the case where $X = Y$.

PROOF. We orient a, b so that v is the common starting point. By Lemma 4.4 there exist legal loops γ^a and γ^b, crossing a and b respectively, each crossing any oriented edge at most once, so that $\overline{\gamma^a}$ and γ^b and are in A_Δ; we choose, γ_a, γ_b to start with a and b respectively. The loop $\gamma = \overline{\gamma^a} g \gamma_b$ crosses τ by construction, and it crosses any un-oriented edge at most 4 times. Let ω be the non-free turn determined at the concatenation of the end γ^b and the beginning $\overline{\gamma^a}$. By construction γ is legal except possibly at ω. Hence, by Lemma 2.18, up to possibly replacing γ with γ_{ω,h_v} or $\gamma_{\omega,id}$ we may assume that γ is legal. By Lemma 2.20 γ is cyclically reduced.

Moreover, both $\gamma_{\tau,id}, \gamma_{\tau,h_v}$ satisfy condition (2) for belonging to A_Δ and at least one of them is legal by Lemma 2.18. Clearly if both γ^a and γ^b are in the tension graph of f, then so is γ. □

REMARK 4.7. If $f : X \to Y$ is a minimal optimal map, then any edge in the tension graph is crossed by a f-legal loop in the tension graph, and so satisfies hypothesis of Lemma 4.4, and any non-free legal turn in the tension graph satisfies the hypothesis of Lemma 4.6, in the tension graph. This is just by the definition of *minimal* optimal map (see Definition 7.15). Moreover, we recall also that if ϕ is irreducible and $f : X \to X$ is an optimal map representing ϕ on a minimally displaced point, then the tension graph of f is the whole X (see Lemma 7.25).

REMARK 4.8. Note that for Lemma 4.6, the hypothesis that the turns are non-free is essential, as the lemma fails for free turns.

Let ϕ be the automorphism of $F_2 =< a, b >$, which sends a to aba and b to ba. Then the iwip automorphism ϕ admits a natural train track representative - which we also call ϕ - on the rose R, where we identify each petal of R with an element of the free basis $\{a, b\}$. Moreover, the turn $\tau = [a, b]$ is $\langle \sim_{\phi^k} \rangle$-legal, as for every positive integer k, $\phi^k(a), \phi^k(b)$ start with a,b, respectively.

However, note that a legal loop cannot contain the cyclic subwords ab^{-1} or ba^{-1}. Therefore the only legal loops are either positive or negative words in a and b. In particular, the free turn τ is ϕ-legal, but it cannot be extended to a ϕ-legal loop.

LEMMA 4.9. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ be an irreducible element, let $X \in \mathrm{Min}(\phi)$ and $f : X \to X$ a train track map representing ϕ. Then, every edge of X is crossed by a $\langle \sim_{f^k} \rangle$-legal loop. In particular, Lemma 4.4 holds true for any edge of X, and Lemma 4.6 for any non-free $\langle \sim_{f^k} \rangle$-legal turn.*

PROOF. Any train track map is also train track with respect to $\langle\sim_{f^k}\rangle$-legality; namely, it maps $\langle\sim_{f^k}\rangle$-legal paths to $\langle\sim_{f^k}\rangle$-legal paths ([11, Corollary 8.12]).

Since ϕ is irreducible and, the tension graph of f is the whole of X and any vertex is at least two-gated with respect to $\langle\sim_{f^k}\rangle$. Therefore, there exists a $\langle\sim_{f^k}\rangle$-legal loop, γ_0, in X. The iterated images $f^n(\gamma_0)$ form a sub-graph of X which is f-invariant. By irreducibility, that sub-graph must be the whole X. In particular any edge e is in the loop $f^n(\gamma_0)$ for some n, and that loop is $\langle\sim_{f^k}\rangle$-legal because f is a train-track map. ∎

The following is just a list of immediate corollaries of previous lemmas.

LEMMA 4.10. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ and Δ a simplex in $\mathcal{O}(\mathcal{G})$. Then $\mathcal{C}_C(\Delta)$ contains (at least) all turns of the following kinds:*

- *(i) Free and finite non-free turns;*
- *(ii) f-images of finite non-free turns, where f is any straight $\mathcal{O}$-map landing on X;*
- *(iii) turns in the f-image of an edge crossed by some f-legal loop, where $f : X \to X$ is any straight map representing ϕ. In particular those include:*
  - *(a) edges in the tension graph of f, when $f : X \to X$ is a minimal optimal map representing ϕ;*
  - *(b) any edge, provided the tension graph of f is the whole X,* e.g. *if ϕ is irreducible and f is an optimal map representing ϕ on the minimally displaced point X;*
- *(iv) turns in the f-image of a free turn crossed by some f-legal loop, where $f : X \to X$ is any straight map representing ϕ.*

PROOF. (i) is by definition. For (ii), note that, since f is an $\mathcal{O}$-map, the f-image of a finite non-free vertex is again a finite non-free vertex. Cases (iii) and (iv) follow immediately from Lemmas 4.4 and 4.7. In particular, case $(iii) - (a)$ follow from Lemma 4.4 by Remark 4.7; case $(iii) - (b)$ from Lemma 4.4 by Remark 4.7; case (iv) from Lemma 4.4. ∎

PROPOSITION 4.11. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$, Δ a simplex in $\mathcal{O}(\mathcal{G})$, and $f : X \to X$ be a straight map representing ϕ at a point $X \in \Delta$. Let $\tau_1, \dots, \tau_k$ be candidate-regular turns. Then for any $j = 1, \dots, k$, any f-legal loop γ_0 crossing τ_j can be modified via turn-surgeries (at infinite non-free turns) to an f-legal loop γ so that*

- *(i) $\#(\gamma, \tau_j) = 1$ and,*
- *(ii) $\sum_{i=1}^{k} \#(\gamma, \tau_i) = 1$ and,*
- *(iii) $\sum_{i=1}^{k} \#(f(\gamma), \tau_i) \leq 1$ and,*
- *(iv) $\sum_{i=1}^{k} \#(f(\gamma), \tau_i) = 0$ unless f maps τ_j to some τ_l.*

Moreover, if in addition γ_0 is $\langle\sim_{f^k}\rangle$-legal, then γ can be chosen to be $\langle\sim_{f^k}\rangle$-legal.

Proof. Without loss of generality we may assume that γ_0 crosses τ_1. We modify γ_0 by using turn-surgeries (Lemma 3.2) in order to get a new loop γ that satisfies the extra properties. We will apply surgeries only on turns at non-free vertices with infinite stabilisers, so there will be infinitely many choices every time.

Concretely, let γ_0 be represented as $g_1e_1\dots g_ne_n$ (with cyclic indices modulo n) so that $\tau_1 = [\bar{e}_1, g_2e_2]$. For any infinite non-free turn $\tau = [\bar{e}_i, g_{i+1}e_{i+1}]$ with $i \neq 1$, we choose an element a in the corresponding vertex group so that $\tau_a = [\bar{e}_i, ae_{i+1}]$ satisfies

(1) τ_a is not one of the τ_i;
(2) τ_a is $\langle\sim_{f^k}\rangle$-legal;
(3) $f(\tau_a)$ is not one of the τ_i;

Such an element exists because τ is infinite non-free, there are finitely many τ_i, and by Lemma 2.18 all but one choice for a produces a $\langle\sim_{f^k}\rangle$-legal turn. We define γ as the result of the turn-surgeries at all such infinite non-free vertices, by using the chosen group elements.

Condition (2) assures that γ is legal, and $\langle\sim_{f^k}\rangle$-legal if γ_0 where so. Since we did not touch τ_1, condition (1) gives us point (i) and (ii). As for (iii), let's analyse the turns crossed by $f(\gamma)$. They come in several types:

(a) a turn crossed by the f-image of an edge of γ,
(b) the f-image of a free turn of γ,
(c) the f-image of a finite non-free turn of γ,
(d) the f-image of an infinite non-free turn of γ.

By Lemma 4.10, the first three are all candidate critical, so none of the τ_i appears in this way (note that in type (b), the free turns that appear, are crossed by the f-legal loop γ, so the hypothesis of 4.10 (iv) is satisfied). Crossings of kind (d) are avoided by condition (3), except possibly if $f(\tau_1)$ equals one of the τ_i's. Points (iii) and (iv) follow. □

Corollary 4.12. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$, Δ a simplex in $\mathcal{O}(\mathcal{G})$, and $f : X \to X$ be a minimal optimal map representing ϕ on a point $X \in \Delta$.*

Suppose that $\tau_1, \dots, \tau_k$ are candidate regular turns. If there is a turn τ_j which is f-legal and in the tension graph of f, then there exists an f-legal loop, γ, in the tension graph, and such that:

(i) $\#(\gamma, \tau_j) = 1$ *and,*
(ii) $\sum_{i=1}^{k} \#(\gamma, \tau_i) = 1$ *and,*
(iii) $\sum_{i=1}^{k} \#(f(\gamma), \tau_i) \leq 1$ *and,*
(iv) $\sum_{i=1}^{k} \#(f(\gamma), \tau_i) = 0$ *unless f maps τ_j to some τ_l.*

Proof. By Remark 4.7, Lemma 4.6 applies for the infinite non-free turn, τ_j. So there is a f-legal loop γ_0, in the tension graph, and crossing τ_j. Proposition 4.11 applies. Since γ is obtained from γ_0 via turn-surgeries, and since γ_0 is in the tension graph, so also γ is in the tension graph. ∎

Corollary 4.13. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ an irreducible element, Δ a simplex in $\mathcal{O}(\mathcal{G})$, and $f : X \to X$ a train-track map representing ϕ on a point $X \in \Delta$.*

Suppose that $\tau_1, \dots, \tau_k$ are candidate regular turns. If there is a turn τ_j which is $\langle \sim_{f^k} \rangle$-legal, then there exists a $\langle \sim_{f^k} \rangle$-legal loop, γ such that:

(i) $\#(\gamma, \tau_j) = 1$ *and,*
(ii) $\sum_{i=1}^{k} \#(\gamma, \tau_i) = 1$ *and,*
(iii) $\sum_{i=1}^{k} \#(f(\gamma), \tau_i) \leq 1$ *and,*
(iv) $\sum_{i=1}^{k} \#(f(\gamma), \tau_i) = 0$ *unless f maps τ_j to some τ_l.*

Proof. Lemma 4.9 applied for τ_j (which is necessarily infinite non-free, as it is regular) guarantees the existence of a $\langle \sim_{f^k} \rangle$-legal loop γ_0 crossing τ_j. Proposition 4.11 applies. ∎

Corollary 4.14. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$. Let Δ be a simplex of $\mathcal{O}(\mathcal{G})$ and f be a minimal optimal map representing ϕ on a point X of Δ. Let τ be a non-degenerate candidate regular turn, and let Δ_τ be the simplex obtained by folding τ. If X^t denotes the point of Δ_τ obtained from X by folding τ by an amount t, then*

$$\lambda_\phi(X^t) \geq \lambda_\phi(X).$$

Moreover, if τ is f-legal and in the tension graph, and if $\lambda(\phi) > 1$, then the inequality is strict.

Proof. By Remark 4.7 and Lemma 4.4, there exists an f-legal loop $\gamma \in A_\Delta$ in the tension graph, and any turn which is crossed by the image of $f(\gamma) = \phi(\gamma)$ is candidate critical just by definition of candidate critical. Since τ is regular

$$\#(f(\gamma), \tau) = 0,$$

the non-strict version of hypothesis ($\heartsuit$) of Lemma 4.1 is fulfilled, and first claim follows.

If in addition τ is legal and in the tension graph, we invoke Corollary 4.12 (with $k = 1$) to build a legal loop γ in the tension graph so that $\#(\gamma, \tau) > 0$ and

$$\#(f(\gamma), \tau) \leq \#(\gamma, \tau).$$

The non strict version of inequality ($\heartsuit$) follows because $\lambda_\phi(X) \geq 1$. Moreover, if $\lambda(\phi) > 1$, then $\lambda_\phi(X) \geq \lambda(\phi) > 1$ and also the strict version is proved. ∎

Corollary 4.15. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$. Let τ be a non-degenerate candidate regular turn with respect to a simplex Δ of $\mathcal{O}(\mathcal{G})$, and Δ_τ be the simplex obtained by folding τ. Then*

$$\lambda_\phi(\Delta_\tau) \geq \lambda_\phi(\Delta).$$

Proof. Any point $Y \in \Delta_\tau$ is obtained by folding $\mathrm{unf}_\tau(Y)$ (Lemma 3.1). Corollary 4.14 tells us $\lambda_\phi(Y) \geq \lambda_\phi(\mathrm{unf}_\tau(Y))$. The claim follows taking infima. □

Corollary 4.15 provides the kind of non-strict inequalities we are searching for. We now focus on turns whose folding guarantees the strict inequality $\lambda_\phi(\Delta_\tau) > \lambda_\phi(\Delta)$.

Lemma 4.16. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$. Let $X \in \mathcal{O}(\mathcal{G})$ and $f : X \to X$ be an optimal map representing ϕ. For any $X_0 \in \Delta_X$ let $f^0 : X_0 \to X_0$ denote the map f read in X_0.*

There is a neighbourhood U of X in Δ_X such that for any $X_0 \in U$ there is a minimal optimal map $f_0 : X_0 \to X_0$ such that

$$d_\infty(f^0, f_0) < \varepsilon$$

and, for any $x, y \in X$

$$|d_X(f(x), f(y)) - d_{X_0}(f^0(x), f^0(y)| < \varepsilon.$$

Proof. The function $\lambda_\phi(X)$ is continuous on X and, tautologically, the metric of X changes continuously on X. Therefore, for any $\varepsilon > 0$ there is a neighbourhood U of X in Δ_X such that

- $|\lambda_\phi(X_0) - \lambda_\phi(X)| < \varepsilon$,
- $d_\infty(\mathrm{Str}(f^0), f^0) < \varepsilon$,
- $|\operatorname{Lip}(\mathrm{Str}(f^0)) - \operatorname{Lip}(f)| < \varepsilon$.

By [12, Theorem 3.15] there exists a weakly optimal map $f_1 : X_0 \to X_0$ representing ϕ such that

$$d_\infty(f_1, \mathrm{Str}(f^0)) \leq \mathrm{vol}(X_0)(\operatorname{Lip}(\mathrm{Str}(f^0)) - \lambda_\phi(X_0))$$

and, by [12, Theorem 3.15 and Theorem 3.24] there exists a minimal optimal map $f_0 : X_0 \to X_0$ representing ϕ such that

$$d_\infty(f_1, f_0) < 2\varepsilon.$$

Putting together all such inequalities, and since ε is arbitrary, we get that for any $\varepsilon > 0$ there is U so that for all $X_0 \in U$ we have

$$d_\infty(f^0, f_0) < \varepsilon.$$

Moreover, it is clear that we can choose U in such a way that for any $x, y \in X$ we have

$$|d_X(f(x), f(y)) - d_{X_0}(f^0(x), f^0(y)| < \varepsilon.$$

□

LEMMA 4.17. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$. For any $X \in \mathcal{O}(\mathcal{G})$ and optimal map $f : X \to X$ representing ϕ, there is a neighbourhood U of X in Δ_X such that for any $X_0 \in U$ there is a minimal optimal map $f_0 : X_0 \to X_0$ such that if τ is a non-free turn in X which is f-legal, then τ is f_0-legal.*

PROOF. We apply Lemma 4.16. By equivariance, if v is the non-free vertex where τ is based, then $f(v) = f_0(v)$ is a non-free vertex. Estimates of Lemma 4.16 now easily imply that if τ is legal in X, it remains legal for small perturbations. ∎

LEMMA 4.18. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$, Δ be a simplex in $\mathcal{O}(\mathcal{G})$, and $X \in \Delta$ be a point which is minimally displaced by ϕ. Suppose that Δ' is a simplex with face Δ and that there is a point $Y \in \Delta'$ which is minimally displaced by ϕ. Then for any open neighborhood U of X in Δ' there is a point Z in U which is minimally displaced by ϕ.*

PROOF. This is an immediate application of the convexity properties of λ_ϕ (namely, by [12, Lemma 6.2]). More specifically, for X, Y as above, the linear segment $\overline{YX}$ eventually enters in U, by continuity. On the other hand, by convexity properties of λ_ϕ, any point of the segment $\overline{YX}$ is minimally displaced by ϕ, which gives us the required result. ∎

PROPOSITION 4.19. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ be irreducible and with $\lambda(\phi) > 1$. Let Δ be a simplex of $\mathcal{O}(\mathcal{G})$, and let $X \in \Delta$. Let $f : X \to X$ be a minimal optimal map representing ϕ. Let τ be a non-degenerate, candidate regular turn with respect to Δ, which is also f-legal and let Δ_τ be the simplex obtained by folding τ[1]. Then for all $Y \in \Delta_\tau$ we have*

$$\lambda_\phi(Y) > \lambda(\phi).$$

PROOF. From Corollary 4.15 we know $\lambda_\phi(\Delta_\tau) \geq \lambda_\phi(\Delta)$, and if $\lambda_\phi(\Delta) > \lambda(\phi)$ the claim follows. Thus we may assume $\lambda_\phi(\Delta) = \lambda(\phi)$.

For any $Z \in \Delta$ denote by Z^t the point of Δ_τ obtained from Z by folding τ by an amount of t (which is well defined, for any Z, for small enough t).

We will prove that there is an open neighbourhood U of X in Δ and $T > 0$ (which depends only on ϕ and X) so that for all $Z \in U$ and $t < T$, the point Z^t, is not minimally displaced. Then the result follows, as $U^T = \{Z^t : Z \in U, t < T\}$ is an open neighbourhood of X in Δ_τ, and by Lemma 4.18, if Δ_τ were to contain a minimally displaced point, we would be able to find a minimally displaced point in U^T, leading to a contradiction. Whence $\lambda_\phi(Y) > \lambda(\phi)$ for all $Y \in \Delta_\tau$.

We prove now our claim. Since τ is candidate regular, in particular it is non-free. By Lemma 4.17 there is a neighbourhood U of X in Δ so that for

[1] Note also that since τ is regular, it is in particular infinite non free, so Δ_τ is different from Δ.

any point $Z \in U$, there is a minimal optimal map $f_Z : Z \to Z$, such that τ is f_Z-legal. Clearly, for any such U there is $T > 0$ so that Z^t is well defined for all $Z \in U$ and $t < T$. By Corollary 4.14 $\lambda_\phi(Z^t) \geq \lambda_\phi(Z)$, and if Z is not minimally displaced, the result follows.

So, suppose that $Z \in \mathrm{Min}(\phi)$. In this case, since f_Z is an optimal map representing ϕ, and since ϕ is irreducible, then the tension graph of f_Z is the whole Z (Lemma7.25), and since $\lambda(\phi) > 1$, Corollary 4.14 applies in its strict inequality version. In any case, Z^t cannot be minimally displaced and our claim follows. ∎

Remark 4.20. If one is interested in a version of Proposition 4.19 for reducible automorphisms, one has just to add the hypothesis that τ is *stably* in the tension graph, that is to say, that τ is in the tension graph of any f_Z for Z close enough to X. This will be enough to apply Corollary 4.14 as we did in the proof for irreducible automorphisms.

Lemma 4.21. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ be an irreducible element with $\lambda(\phi) > 1$. Suppose that $X_1, X_2 \in \Delta$ are two points minimally displaced by ϕ and let f_1, f_2 be train track representative of ϕ on X_1 and X_2 respectively.*

Suppose that τ is candidate regular turn. Then τ is f_1-legal in X_1 if and only if it is f_2-legal in X_2.

Proof. Suppose for contradiction that τ is f_1-legal but f_2-illegal (in particular it is non-degenerate). Let Δ_τ be the simplex obtained by folding τ. Since τ is candidate regular, we apply Proposition 4.19 (we can because train tracks are minimal optimal map by Lemma 7.25), and we have that $\lambda_\phi(Y) > \lambda(\phi)$ for any $Y \in \Delta_\tau$. On the other hand, τ is f_2-illegal, and $\mathrm{Min}(\phi)$ is invariant under isometrically folding illegal turns (see [11, Theorem 8.23]), which means that there is a point $Y \in \Delta_\tau$ which is minimally displaced and that leads us to a contradiction. Clearly we can switch the roles of X_1 and X_2, and the proof is complete. ∎

Definition 4.22 (Simplex Critical Turns). *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$. For a simplex, Δ, we say that a turn τ is* **simplex critical** *if it is either candidate critical (see 4.3) or f-illegal for some train track representative of ϕ defined on some point of Δ (if any). A turn is called* **simplex regular** *if it is not critical. The set of simplex critical turns is denoted by $\mathcal{C}_\Delta$. Sometimes we will use the short notation* Δ-critical *to means* simplex critical in Δ.

Theorem 4.23. *If $[\phi] \in \mathrm{Out}(\mathcal{G})$ is irreducible with $\lambda(\phi) > 1$, then the set $\mathcal{C}_\Delta$ is finite. In fact, it is sufficient to add to $\mathcal{C}_C(\Delta)$ the illegal turns for a single train track map (if any) to obtain the whole of $\mathcal{C}_\Delta$.*

Proof. The set $\mathcal{C}_C(\Delta)$ is clearly finite by construction (see Definition 4.3), and if there is no minimal displaced point in Δ, we have nothing to prove. Otherwise, chose any train track representative of ϕ, f, defined on a point of

Δ. f-illegal turns are finitely many (Lemma 2.18) and Lemma 4.21 tells us that $\mathcal{C}_\Delta = \mathcal{C}_C(\Delta) \cup \{\tau : \tau \text{ is an } f\text{-illegal turn}\}$ and, in particular, it is finite. $\square$

Remark 4.24. It is worth mentioning that simplex regular turns can be effectively detected, having a train track map f in hand. Namely, suppose that a turn τ is

1. not a free turn nor a turn with finite vertex group; and
2. not the f-image of an edge; and
3. not the f-image of a free turn; and
4. not the f-image of a turn involving group elements in H; and
5. not f-illegal;

then τ is simplex regular. In particular, Proposition 4.19 tells us that if we have $X \in \mathrm{Min}(\phi)$ and we want to find all neighbours of X obtained from X by a single turn-fold, and which still are in $\mathrm{Min}(\phi)$, then we only need to check turns in the finite complement of the above effective list, namely turns that are either

1. free or with finite vertex group; or
2. in the f-image of an edge; or
3. the f-image of a free turn; or
4. the f-image of a turn involving group elements in H; or
5. f-illegal.

5. Folding and unfolding collapsed forests

Here we extend the unfolding construction of Section 3 to the general case of two simplices, one face of the other. We remind that we always understand that a straight map between elements of $\mathcal{O}(\mathcal{G})$ is an $\mathcal{O}$-map (i.e. G-equivariant at level of trees).

A straight map $p : X \to Y$, defines on X a simplicial structure σ_p, by pulling back that of Y. With respect to σ_p, the map p is tautologically simplicial. We define the **simplicial volume** of p, $\mathrm{svol}(p)$ as the number of edges of σ_p.

If in addiction p is **locally isometric on edges**, then it defines (some) folding paths $X = X_0, \dots, X_n = Y$ obtained by recursively identifying pairs of edges of σ_p having a common vertex and the same p-image. Together with the X_i there are quotient maps $q_i : X_{i-1} \to X_i$ given by the identification, and maps $p_i : X \to Y$ defined by $p_i(x) = p(q_i^{-1}(x))$. (Note that $p_0 = p$ and $p_n = id$). We refer to any folding path obtained as above as a **folding path directed by** p. We say that $X = X_0, \dots, X_n = Y$ has length n.

Lemma 5.1. *Let $X, Y \in \mathcal{O}(\mathcal{G})$ and $p : X \to Y$ be a straight map which is locally isometric on edges. Then any folding path directed by p has length at most* $\mathrm{svol}(p)$.

PROOF. At any step the number of edges decreases by one. ∎

LEMMA 5.2. *Let Δ, Δ' be simplices of $\mathcal{O}(\mathcal{G})$ so that Δ is a face of Δ'. For any $Y \in \Delta'$ there is a point,* $\mathrm{unf}(Y) \in \Delta$ *and a straight map, $p : \mathrm{unf}(Y) \to Y$, such that:*

1. *p is a local isometry on edges;*
2. *If v is a non-free vertex in Y, then $p^{-1}(v)$ is a single vertex;*
3. $\mathrm{svol}(p)$ *is at most $2D(\Delta')^2$, where $D(\Delta')$ is the number of edges of Δ'.*

Moreover, any folding path directed by p produces maps which still satisfy (1) *and* (2).

PROOF. The underlying graph X of Δ is obtained by the collapse in Y of a simplicial forest $\mathcal{F} = T_0 \sqcup \cdots \sqcup T_k$ each of whose tree T_i contains at most one non-free vertex. We define $\mathrm{unf}(Y)$ by isometrically unfolding each tree. More precisely, for any T_i we choose a root-vertex w_i with the requirement that w_i is the unique non-free vertex of T_i, if any. For any leaf y of the forest, say y is a leaf of T_i, there is a unique path γ_y connecting y to w_i in T_i. For notational convenience we define γ_y to be the constant path for any other vertex of Y.

The metric on X defining the point $\mathrm{unf}(Y)$ is given as follows. Any edge e of X has a preimage in Y which is also an edge. We declare

$$L_{\mathrm{unf}(Y)}(e) = L_Y(\gamma_a) + L_Y(e) + L_Y(\gamma_b)$$

where a, b are the endpoints of the preimage of e in Y. As an oriented edge, e is therefore the concatenation of three sub-segments

$$e = A \cdot E \cdot B$$

of lengths $L_Y(\gamma_a), L_Y(e)$, and $L_Y(\gamma_b)$ respectively. The map p is now defined by isometrically identifying A with γ_a, E with the copy of e in Y, and B with γ_b. The union of all A-segments and B-segments form a forest which can be viewed as an isometric unfolding of $\mathcal{F}$, and out of that forest, p is basically the identity by definition. Conditions (1) and (2) immediately follow. As for (3), it suffices to note that for any $y \in Y$, the cardinality of $p^{-1}(y)$ is bounded by the number of leaves of $\mathcal{F}$, which is bounded by $2D(\Delta')$. Therefore σ_p has at most $2D(\Delta)D(\Delta') \leq 2D(\Delta')^2$ edges. The last claim is easy to verify and we leave it to the reader. ∎

LEMMA 5.3. *Let $X, Y \in \mathcal{O}(\mathcal{G})$. Let $p : X \to Y$ be a straight map such that if v is a non-free vertex in Y, then $p^{-1}(v)$ is a single vertex in X (condition* (2) *in Lemma 5.2).*

Let α, β be paths in X, starting at the same vertex, and so that $\bar{\alpha}\beta$ is reduced. If $p(\alpha) = p(\beta)$, then the only non-free vertex crossed by each of them, if any, is their initial vertex (which is crossed only once).

PROOF. We argue by contradiction assuming that α crosses a non-free vertex other than its initial point (we consider multiple crossings of the same

vertex as distinct crossings). Up to possibly truncating α and β, we may assume that the last vertex w of α is non-free, and that α crosses no other non-free vertex except possibly its initial point. Then our assumption on p implies that the last vertex of β must be w. But in this case $\alpha\bar{\beta}$ would define a non trivial group-element which is collapsed by p, contradicting that X, Y are in the same deformation space. ∎

PROPOSITION 5.4. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ be an irreducible element with $\lambda(\phi) > 1$. Let $X, Y \in \mathcal{O}(\mathcal{G})$, $X \neq Y$, and suppose that there is $p : X \to Y$ a straight map such that*

1. *p is a local isometry on edges;*
2. *if v is a non-free vertex in Y, then $p^{-1}(v)$ is a single vertex.*

If any p-illegal turn is simplex regular (Definition 4.22), then $\lambda_\phi(Y) > \lambda(\phi)$.

PROOF. By Lemma 2.18, the set of p-illegal turn is a finite set $\{\tau_1, \dots, \tau_k\}$. Since $X \neq Y$, this set is non-empty.

Denote Δ the simplex of X. Let $f : X \to X$ be a minimal optimal map representing ϕ. Note that for any element of G, seen as a loop γ in X, the loop $p(f(\gamma))$ represents $\phi(p(\gamma))$ in Y.

Firstly, we deal with the case where $X \notin \mathrm{Min}(\phi)$. By Remark 4.7 and Lemma 4.4 there is an f-legal loop $\gamma \in A_\Delta$ in the tension graph of f. Its image, $f(\gamma)$, crosses only candidate critical turns by definition of candidate critical and, in particular, it doesn't cross any of the τ_i's. In other words, $f(\gamma)$ is p-legal and therefore, as p is a local isometry on every edge, $L_Y(p(f(\gamma))) = L_X(f(\gamma))$. On the other hand, again because p is a local isometry, $L_Y(p(\gamma)) \leq L_X(\gamma)$ which means that

$$\lambda_\phi(Y) \geq \frac{L_Y(p(f(\gamma)))}{L_Y(p(\gamma))} \geq \frac{L_X(f(\gamma))}{L_X(\gamma)} = \lambda_\phi(X) > \lambda(\phi).$$

Suppose now $X \in \mathrm{Min}(\phi)$. In this case we may assume that f is a train track map representing ϕ (in particular f is a minimal optimal map, see Section 7). All τ_i are f-legal because of simplex-regularity.

Let's first assume that there is some τ_j which is mapped by f to a turn distinct from any of the τ_i's. Then, by Corollary 4.12, there is an f-legal loop γ (which is in the tension graph) such that,

1. $\sum_{i=1}^k \#(\gamma, \tau_i) = 1$,
2. $\sum_{i=1}^k \#(f(\gamma), \tau_i) = 0$.

Hence, $L_Y(p(\gamma)) < L_X(\gamma)$ whereas, $L_Y(p(f(\gamma))) = L_X(f(\gamma))$. Hence,

$$\lambda_\phi(Y) \geq \frac{L_Y(p(f(\gamma)))}{L_Y(p(\gamma))} > \frac{L_X(f(\gamma))}{L_X(\gamma)} = \frac{\lambda_\phi(X) L_X(\gamma)}{L_X(\gamma)} = \lambda(\phi).$$

Otherwise, f must leave invariant the set of τ_i. We will now work with the $\langle\sim_{f^k}\rangle$-legal structure, in order to ensure that the image of a legal loop is again legal. Since all τ_i are legal and set of τ_i is invariant under the action of f, then they also are $\langle\sim_{f^k}\rangle$-legal.

Let Σ be the set of $\langle\sim_{f^k}\rangle$-legal loops γ in X that satisfy

$$\sum_{i=1}^{k} \#(\gamma, \tau_i) = 1.$$

By Corollary 4.13, the set Σ is not empty. Let

$$C = \sup\{L_X(\gamma) - L_Y(p(\gamma)) : \gamma \in \Sigma\}.$$

The Bounded Cancellation Lemma (see for instance [22, Proposition 3.12]) and discreteness show that C is a maximum, which means that C is realised by some loop γ_C. Moreover, since the τ_i's are p-illegal, $C > 0$. We claim that γ_C can be chosen so that $f(\gamma_C)$ also belongs to Σ. This will be enough as, since γ_C realises C, then $L_Y(p(\gamma_C)) = L_X(\gamma_C) - C$, while $L_Y(p(f(\gamma_C))) \geq L_X(f(\gamma_C)) - C$ (because $f(\gamma_C) \in \Sigma$). But then

$$\lambda_\phi(Y) \geq \frac{L_Y(p(f(\gamma_C)))}{L_Y(p(\gamma_C))} \geq \frac{L_X(f(\gamma_C)) - C}{L_X(\gamma_C) - C} = \frac{\lambda_\phi(X) L_X(\gamma_C) - C}{L_X(\gamma_C) - C} > \lambda(\phi)$$

where the strict inequality follows from the fact that $\lambda_\phi(X) = \lambda(\phi) > 1$.

We prove now our claim. Consider any $\gamma \in \Sigma$ realising the maximum C. By Proposition 4.11 γ can be modified via turn surgeries to a $\langle\sim_{f^k}\rangle$-legal turn γ' such that $\sum_i \#(f(\gamma'), \tau_i) = 1$. Note that such surgeries occur only at non-free vertices.

It remains to show that the performed surgeries do not affect the p-cancellation of the original loop γ. As τ_j is the unique p-illegal of γ, there exist sub-paths α, β of γ so that $p(\alpha) = p(\beta)$, the first edge of α together with the first edge of β form the turn τ_j, and $L_X(\alpha) = L_X(\beta) = C/2$. That is, α and β are the sub-paths of γ which realise the p-cancellation. By Lemma 5.3, both α and β cross only free turns and so the performed surgeries did not affect neither α, nor β. As the turn τ_j is not affected by the surgeries, as well, it follows that the p-cancellation of γ', is the same as the p-cancellation of γ, that is to say

$$L_X(\gamma') - L_Y(p(\gamma')) = L_X(\gamma) - L_Y(p(\gamma)) = C$$

as we wanted. □

6. Exploring the Minset

Proposition 4.19 tells us that if we want to travel along $\mathrm{Min}(\phi)$, we have to perform only simplex critical turns. Given two simplices Δ, Δ^1, one face on the other, we can easily go from Δ to Δ^1 in few steps by simple folds. However,

even if both simplices intersect $\mathrm{Min}(\phi)$, such folds need not necessarily to be simplex critical. Nonetheless, it may exists a, a priori longer, folding path between them that uses only simplex critical folds.

Definition 6.1. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ and Δ be a simplex in $\mathcal{O}(\mathcal{G})$. We denote by the* **simplex critical neighbourhood** *of Δ of radius 1, all the simplices of $\mathcal{O}(\mathcal{G})$ which can be obtained from Δ via a simplex critical fold, including Δ itself.*

We denote by the simplex critical neighbourhood of Δ of radius $n+1$, the union of all the simplex critical neighbourhoods of radius 1, of all simplices in the simplex critical neighbourhood of Δ of radius n.

Remark 6.2. By Theorem 4.23, if ϕ is irreducible with $\lambda(\phi) > 1$, then any simplex critical neighbourhood of finite radius consists of finitely many simplices.

Definition 6.3. *Let X be a point of $\mathcal{O}(\mathcal{G})$ and let's denote by $\Delta = \Delta_X$ the corresponding simplex. The dimension of $D(\Delta)$ of Δ is the number of edges of X. We denote by $D = D(\mathcal{G})$ the dimension of $\mathcal{O}(\mathcal{G})$, i.e. the maximum number of edges we see in elements of $\mathcal{O}(\mathcal{G})$.*

In $\mathcal{O}_1(\mathcal{G})$ the dimension of the simplex containing X is one less, as is the dimension of the entire space.

Theorem 6.4. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ be an irreducible automorphism with $\lambda(\phi) > 1$. Let Δ, Δ^1 be open simplices of $\mathcal{O}(\mathcal{G})$, both intersecting $\mathrm{Min}(\phi)$, and such that Δ is a face of Δ^1. Then, Δ^1 is contained in the simplex critical neighbourhood of Δ of radius $2D(\mathcal{G})^2$. In particular, $\mathrm{Min}(\phi)$ is locally finite.*

We immediately deduce the same statement for $\mathcal{O}_1(\mathcal{G})$.

Proof. We will show there exists a sequence of simplices, $\Delta_0, \Delta_1, \ldots, \Delta_n$, with $\Delta_0 = \Delta$ and $\Delta_n = \Delta^1$, where $n \leq 2(D(\mathcal{G}))^2$, and such that each Δ_{i+1} is obtained by folding a Δ_i-simplex critical turn.

The underlying graph of Δ is obtained from that of Δ^1 by collapsing a forest $\mathcal{F}$ each of whose tree contains at most one non-free vertex. We apply Lemma 5.2, to get a straight map, $p : \mathrm{unf}(Y) \to Y$ which is locally isometric on edges. Subdivide $\mathrm{unf}(Y)$ so that the p-image of each subdivided edge is a single edge in Y. Proposition 5.4 tells us that it must exist a p-illegal turn that is also simplex critical. Fold this turn; this is an isometric fold directed by p, since p is an isometry on edges, and we get a map $p_1 : X_1 \to Y$ which, by Lemma 5.2, satisfies conditions (1) and (2) of Proposition 5.4, which therefore applies. This process recursively defines a folding path from $\mathrm{unf}(Y)$ to Y, directed by p. The length of such folding path is bounded by $2D(\Delta^1)^2$ by Lemmas 5.1 and 5.2.

To conclude the proof, note that any simplex of $\mathcal{O}(\mathcal{G})$ adjacent to Δ and lying in $\mathrm{Min}(\phi)$ either has Δ as a face, or is a face of Δ. The first case is

dealt with above, and in the second case, we know that there are at most $2^{D(\Delta)} \leq 2^{D(\mathcal{G})}$ faces.

The statement for $\mathcal{O}_1(\mathcal{G})$ follows since the displacement function is invariant under change of volume. ∎

We now show that Theorem 6.4 may be strengthened to show that the minimally displaced set is *uniformly* locally finite. That is, there is a uniform bound (depending only on $\lambda(\phi)$ and $D(\mathcal{G})$) on the number of simplices, adjacent to a given simplex in $\mathrm{Min}(\phi)$ which are also in $\mathrm{Min}(\phi)$. In what follows we are not focused in optimal bounds.

Definition 6.5. *Let Δ be a simplex in $\mathcal{O}(\mathcal{G})$. Then we define the centre $X_\Delta \in \Delta$ to be the graph where all edges have the same length.*

Since we are interested in the function λ_ϕ, which is scale invariant, we may scale the metric on X_Δ as we wish; we will therefore decree that all the edges of X_Δ have length 1.

Lemma 6.6. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$. For any $X, Y \in \mathcal{O}(\mathcal{G})$ we have*

$$\frac{\lambda_\phi(X)}{\lambda_\phi(Y)} \leq \Lambda(X,Y)\Lambda(Y,X).$$

Proof. This immediately follows from the non-symmetric *triangle inequality*:

$$\lambda_\phi(X) = \Lambda(X,\phi X) \leq \Lambda(X,Y)\Lambda(Y,\phi Y)\Lambda(\phi Y,\phi X) = \Lambda(X,Y)\lambda_\phi(Y)\Lambda(Y,X).$$

∎

As above, $D = D(\mathcal{G}) = \dim(\mathcal{O}(\mathcal{G}))$ is the maximum number of (orbits of) edges we see in elements of $\mathcal{O}(\mathcal{G})$.

Lemma 6.7. *Let $[\phi] \in Out(\mathcal{G})$, and Δ, Δ' be simplices in $\mathcal{O}(\mathcal{G})$ such that Δ is a face of Δ'. Then*

$$\frac{1}{2D}\lambda_\phi(X_\Delta) \leq \lambda_\phi(X_{\Delta'}) \leq 2D\lambda_\phi(X_\Delta).$$

Proof. By Lemma 6.6, it is sufficient to prove that

$$\Lambda(X_\Delta, X_{\Delta'})\Lambda(X_{\Delta'}, X_\Delta) \leq 2D.$$

Since the product on the left is scale invariant (even though each factor is not) we are free to choose the volumes for each of the points. Specifically, the two centres have different volumes, as we set every edge to have length 1. In particular,

$$\Lambda(X_{\Delta'}, X_\Delta) \leq 1,$$

as loops become shorter when we collapse a forest (since the length in each case is a count of the number of edges).

To complete the argument, we consider the map $p : \mathrm{unf}(X_{\Delta'}) \to X_{\Delta'}$ given by Lemma 5.2, and we read it as a map from $X_\Delta \to X_{\Delta'}$. It is easy to see that the image of an edge under this map cannot cross the same edge more than twice (usually no more than once, but twice may happen if Δ has some edge-loop). It follows that $L_{X_{\Delta'}}(p(\gamma)) \leq 2DL_X(\gamma)$. Hence $\Lambda(X_\Delta, X_{\Delta'}) \leq 2D$. ∎

The maximum number of vertices we see in elements of $\mathcal{O}(\mathcal{G})$ is bounded by $2D$. Denote by $M = M(\mathcal{G})$ the maximal cardinality of finite vertex groups. Set $K = K(\mathcal{G}) = D + M + 1$.

Lemma 6.8. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$, Δ be a simplex of $\mathcal{O}(\mathcal{G})$, and $\mathcal{C}_\Delta$ be the corresponding set of simplex critical turns. Then*

$$|\mathcal{C}_\Delta| \leq (10K)!\lambda_\phi(X_\Delta)$$

Proof. Since all the edges of X_Δ have length 1, the number of turns crossed by a loop is then equal to its length in X_Δ. Hence, the number of turns crossed by an element of A_Δ is bounded above by

$$\lambda_\phi(X_\Delta)4D|A_\Delta|,$$

where the term $4D$ appears as the maximum length of a loop in A_Δ, as read in X_Δ (see Definition 4.2). Note that the length in X_Δ is simply the number of edges.

To estimate $|A_\Delta|$, we count the number of sequences of edges, where each edge appears at most 4 times - ignoring the incidence relations to simplify matters. The number of sequences of n objects of length k, is $n!/(n-k)! \leq n!$, and so the number of sequences of n objects of length at most n is bounded by $(n+1)!$. For building a loop in A_Δ we have D edges, each of which appear at most 4 times and, taking in account the group elements, the total number of objects we can use to write an element of A_Δ is $8D$. Hence

$$|A_\Delta| \leq (8D+1)!$$

This only bounds the number of turns crossed by elements of A_Δ. In order to bound the simplex critical turns, we need to add all turns based at vertices with finite vertex group, and all the illegal turns for a putative train track map. The former is bounded by the number of possible pairs of germs of edges multiplied by the cardinality of finite vertex group, hence by $M(\mathcal{G})D^2$. The latter, because of Lemma 2.18, is bounded by the number of pairs of edges; hence by D^2. In total we have

$$\lambda_\phi(X_\Delta)4D(8D+1)! + (M+1)D^2 \leq \lambda_\phi(X_\Delta)4K(8K)! + K^3$$

and the result follows. ∎

LEMMA 6.9. *Let $[\phi] \in \mathcal{O}(\mathcal{G})$ be an irreducible element, and suppose that Δ is an open simplex of $\mathcal{O}(\mathcal{G})$ which contains a point of* $\mathrm{Min}(\phi)$. *Then,*

$$\lambda_\phi(X_\Delta) \leq 9 \dim \mathcal{O}(\mathcal{G}) \lambda(\phi)^{3D+2}.$$

PROOF. Let $X_{\min}$ denote a point of Δ which is minimally displaced by ϕ. We will use the fact that $X_{\min}$ is ϵ-thick, as in [2, Proposition 10] (it is proved there in CV_n, but the proof is the same in this context, see also [11, Section 8]). That is, since ϕ is irreducible, there is a lower bound on $L_X(\gamma)/\operatorname{vol}(X)$ that depends only on $\lambda(\phi)$ (and not on $X \in \mathcal{O}(\mathcal{G})$ nor on the non-elliptic element γ). Concretely, this lower bound can be taken to be the reciprocal of $C(\phi) = 3\dim(\mathcal{O}(\mathcal{G}))\lambda(\phi)^{3\dim(\mathcal{O}(\mathcal{G}))+1}$. If we normalise $X_{\min}$ to have volume 1 then $\Lambda(X_\Delta, X_{\min}) \leq 1$. Moreover, since the stretching factor Λ is realised by candidate loops of simplicial length at most 3 (by the Sausage Lemma [11, Theorem 9.10]), we can then deduce that,

$$\Lambda(X_{\min}, X_\Delta)\Lambda(X_\Delta, X_{\min}) \leq 3C(\phi)\Lambda(X_\Delta, X_{\min}) \leq 3C(\phi).$$

The result now follows from Lemma 6.6. □

COROLLARY 6.10. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ be an irreducible element with $\lambda(\phi) > 1$. Then* $\mathrm{Min}(\phi)$ *is uniformly locally finite.*

PROOF. By Theorem 6.4, it is sufficient to show that the simplex critical neighbourhood N of radius $2D^2$, of a minimally displaced simplex Δ_0, contains a uniformly bounded number of simplices (the number of faces of a simplex is always uniformly bounded). Hence it suffices to uniformly bound the cardinality of $\mathcal{C}_\Delta$ for each simplex we encounter in N.

By Lemma 6.7, $\lambda_\phi(X_\Delta) \leq (2D)^{2D^2}\lambda_\phi(X_{\Delta_0})$, which is uniformly bounded by Lemma 6.9. Lemma 6.8 completes the proof. □

7. DEFINITIONS AND BASIC RESULTS USED IN THE PAPER

Our notation and definitions are quite standard. We briefly recall them here, referring the reader to [12] for a detailed discussion.

DEFINITION 7.1. *A free splitting $\mathcal{G}$ of a group G is a decomposition of G as a free product $G_1 * \cdots * G_k * F_n$ where F_n is the free group of rank n. We admit the* trivial *splitting $G = F_n$. We do not require that the groups G_i's are indecomposable.*

DEFINITION 7.2. *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. A G-graph is a graph of groups whose fundamental group, as graph of groups, is isomorphic to G. The action of G on a G-tree, is called marking.*

DEFINITION 7.3. *Let $\mathcal{G}$ be a free splitting of G. In terms of Bass-Serre theory, a $\mathcal{G}$-tree is the tree dual to $\mathcal{G}$, and a $\mathcal{G}$-graph is the corresponding graphs of groups. More explicitly, a $\mathcal{G}$-tree is a simplicial G-tree T such that:*

- *For every G_i there is exactly one orbit of vertices whose stabilizer is conjugate to G_i. Such vertices are called* non-free. *Other vertices are called* free.
- *T has trivial edge stabilisers.*

A $\mathcal{G}$-graph dual to a $\mathcal{G}$-tree. Namely, it is a finite connected G-graph of groups X, along with an isomorphism, $\Psi_X : G \to \pi_1(X)$ - a marking - such that:

- *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, via Ψ_X, to $\mathcal{G}$. That is, Ψ_X restricts to a bijection from the conjugacy classes of the G_i to the vertex groups of X.*

The universal cover of a $\mathcal{G}$-graph is a $\mathcal{G}$-tree and the G-quotient of a $\mathcal{G}$-tree is a $\mathcal{G}$-graph.

Definition 7.4. *Let $\mathcal{G}$ be a splitting of a group G. The Outer Space of $\mathcal{G}$, also known as deformation space of $\mathcal{G}$, and denoted $\mathcal{O}(\mathcal{G})$ is the set of classes of minimal, simplicial, metric $\mathcal{G}$-graphs X with no redundant vertex. (The equivalence relation is given by G-isometries.)*

We denote by $\mathcal{O}_1(\mathcal{G})$ the volume 1 subset of $\mathcal{O}(\mathcal{G})$.

$\mathcal{O}(\mathcal{G})$ can be regarded also as set of $\mathcal{G}$-trees, but in the present paper we adopted the graph view-point. Given a graph of groups X with trivial edge groups, we denote by $\mathcal{O}(X)$ the corresponding deformation space $\mathcal{O}(\pi_1(X))$ (we notice that $X \in \mathcal{O}(X)$ if X is a core graph with no redundant vertex) where $\pi_1(X)$ is endowed with the splitting given by vertex groups. We refer the reader to [11, 12, 17] for more details on deformation spaces.

Definition 7.5. *Let $X \in \mathcal{O}(\mathcal{G})$. The simplex Δ_X is the set of marked metric graphs obtained from X by just changing edge-lengths. Since edge-lengths are strictly positive we think Δ_X as (a cone over) an open simplex. Given a simplex Δ one can consider the the closure of $\Delta \in \mathcal{O}(\mathcal{G})$ or its simplicial bordification. Namely, faces of Δ come in two flavours: that in $\mathcal{O}(\mathcal{G})$, called finitary faces, and that not in $\mathcal{O}(\mathcal{G})$ (typically in other deformation spaces) called faces at infinity.*

We also have a simplex $\Delta_1(X)$ in $\mathcal{O}_1(\mathcal{G})$ - the intersection of $\Delta(X)$ with $\mathcal{O}_1(\mathcal{G})$ - which is a standard open simplex of one dimension less.

There are two topologies on $\mathcal{O}(\mathcal{G})$, both of which restrict to the Euclidean topology on each simplex; these are the weak topology and the axes or Gromov-Hausdorff topology. The topology induced by the Lipschitz metric is the latter one.

Definition 7.6. *Let G be endowed with the splitting $\mathcal{G} : G = G_1 * \cdots * G_i * F_n$. The group of automorphisms of G that preserve the set of conjugacy classes of the G_i's is denoted by $\mathrm{Aut}(\mathcal{G})$. We set $\mathrm{Out}(\mathcal{G}) = \mathrm{Aut}(\mathcal{G})/\mathrm{Inn}(G)$.*

The group $\mathrm{Aut}(\mathcal{G})$ naturally acts on $\mathcal{O}(\mathcal{G})$ by precomposition on marking, and $\mathrm{Inn}(\mathcal{G})$ acts trivially, so $\mathrm{Out}(\mathcal{G})$ acts on $\mathcal{O}(\mathcal{G})$.

Since the volume is invariant under this action, we also get an action of $\mathrm{Aut}(\mathcal{G})$ and $\mathrm{Out}(\mathcal{G})$ on $\mathcal{O}_1(\mathcal{G})$.

Definition 7.7. *Given a splitting $\mathcal{G}$ of G, and $X, Y \in \mathcal{O}(\mathcal{G})$, a map $f : X \to Y$ is called an $\mathcal{O}$-map at the level of the tree, if it is Lipschitz-continuous and G-equivariant (resp., a $\mathcal{G}$-map between graph of groups, it called an $\mathcal{O}$-map if it is the projection of an $\mathcal{O}$-map at level of the universal covers). The Lipschitz constant of f is denoted by* $\mathrm{Lip}(f)$.

Definition 7.8. *Let X, Y be two metric graphs. A Lipschitz-continuous map $f : X \to Y$ 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 edge e is uniformly stretched by a factor $\lambda_e(f)$. A straight map between elements of $\mathcal{O}(\mathcal{G})$ is always supposed to be an $\mathcal{O}$-map.*

Remark 7.9. $\mathcal{O}$-maps always exist and the images of non-free vertices are determined a priori by equivariance (see for instance [11]). For any $\mathcal{O}$-map f there is a unique straight map denoted by $\mathrm{Str}(f)$, which is homotopic, relative to vertices, to f. We have $\mathrm{Lip}(\mathrm{Str}(f)) \leq \mathrm{Lip}(f)$.

Definition 7.10. *Let $f : X \to Y$ be a straight map. We set $\lambda_{\max}(f) = \max_e \lambda_e(f) = \mathrm{Lip}(f)$ and define the* tension graph *of f as the set*

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

Definition 7.11. *Given $X, Y \in \mathcal{O}(\mathcal{G})$ we define $\Lambda(X, Y)$ as the infimum of Lipschitz constants of $\mathcal{O}$-maps from X to Y. That inf is in fact a minimum and coincides with $\max_\gamma \frac{L_Y(\gamma)}{L_X(\gamma))}$ where γ runs on the set of loops in X (seen as a graph). (See for instance* [9, 11, 12].*)*

Definition 7.12. *A gate structure on a graph of groups X is a G-equivariant equivalence relation of germs of edges at vertices of the universal cover, $\widetilde{X}$, of X. A train track structure on a graph of groups X is a gate structure on X with at least two gates at each vertex.*

Remark 7.13. For a straight map $f : X \to Y$, we consider two different gate structures, which we denote by $\sim_f$ and $\langle \sim_{f^k} \rangle$, the latter being defined only if $X = Y$. Two germs of X are $\sim_f$-equivalent, if they have the same non-collapsed f-image and they are $\langle \sim_{f^k} \rangle$-equivalent, if they have the same non-collapsed f^k-image for some positive integer k.

This second gate structure - $\langle \sim_{f^k} \rangle$ - only makes sense if Y is the same topological object as X, so that we may iterate f. However, Y will usually be a different point of $\mathcal{O}(\mathcal{G})$ since the G-action will be different.

We refer to $\sim_f$, as the gate structure which is induced by f.

Definition 7.14. *A turn of $X \in \mathcal{O}(\mathcal{G})$ is (the G_v-orbit of) an unoriented pair of germs of edges based at a vertex v of X. A turn is legal if its germs are not in the same gate. A simplicial path in X is legal if it crosses only legal turns. Legality here depends on the gate structure, which for us will either be the $\sim_f$ or $\langle \sim_{f^k} \rangle$ structure for some $\mathcal{O}$-map f with domain X.*

Definition 7.15. *Let $X, Y \in \mathcal{O}(\mathcal{G})$. A straight map $f : X \to Y$ is said to be optimal if $\mathrm{Lip}(f) = \Lambda(X, Y)$ and every vertex of the tension graph is at least two-gated (i.e. the gate structure $\sim_f$ is a train track structure on the tension graph). An optimal map is minimal if every edge of the tension graph extends to a legal loop in the tension graph (not all optimal maps are minimal, but minimal optimal maps always exist (see* [12]*)).*

Definition 7.16. *Given $[\phi] \in \mathrm{Out}(\mathcal{G})$ and $X \in \mathcal{O}(\mathcal{G})$ we say that an $\mathcal{O}$-map $f : X \to \phi X$ represents ϕ. Note that X and ϕX are the same graph with different markings, so we sometimes abuse notation by saying that f is a map $f : X \to X$ which represents ϕ. In this situation we can speak of the $\langle \sim_{f^k} \rangle$ gate structure.*

Any ϕ is represented by a minimal optimal map (see [12]).

Definition 7.17. *We call $[\phi] \in \mathrm{Out}(\mathcal{G})$* reducible *if there exists an $\mathcal{O}$ map $f : X \to X$ representing ϕ, a lift $\widetilde{f} : \widetilde{X} \to \widetilde{X}$ and a G-subforest, $Y \subsetneq \widetilde{X}$ which is $\widetilde{f}$-invariant and contains the axis of a hyperbolic element. Otherwise $[\phi]$ is called irreducible.*

Remark 7.18. An automorphism ϕ is called *iwip* - irreducible with irreducible powers - if every positive iterate of ϕ is irreducible. We mention this for completeness, but we are concerned with the general irreducible class for this paper.

Definition 7.19. *A straight map $f : X \to X$ representing ϕ is a train track map, if there is a train track structure on X so that:*

- *f maps edges to legal paths*
- *If $f(v)$ is a vertex, then f maps inequivalent germs at v to inequivalent germs at $f(v)$.*

Remark 7.20. (See [11, 12] for more details):

1. If $f : X \to X$ is train track map representing ϕ (with respect to some gate structure), then it is a train track map with respect to the $\langle \sim_{f^k} \rangle$ gate structure.
2. Irreducible elements of $\mathrm{Out}(\mathcal{G})$ admit train track representatives.

Definition 7.21. *Given $[\phi] \in \mathrm{Out}(\mathcal{G})$ the displacement function*

$$\lambda_\phi : \mathcal{O}(\mathcal{G}) \to \mathcal{O}(\mathcal{G})$$

is defined by

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

If $f : X \to X$ is any $\mathcal{O}$-map representing ϕ, then

$$\lambda_\phi(X) = \sup_\gamma \frac{L_X(\phi\gamma)}{L_X(\gamma)}$$

where the sup is taken over all loops γ in X (and it is actually a maximun by [11, 12]*) and $L_X(\gamma)$ denotes the reduced length of γ. For a simplex Δ we define*

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

and

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

Remark 7.22. Note that the displacement function - λ_ϕ - is invariant under change of volume, so one can work interchangeably between $\mathcal{O}(\mathcal{G})$ and $\mathcal{O}_1(\mathcal{G})$.

Theorem 7.23 ([2, 11, 12]). *Given $[\phi] \in \mathrm{Out}(\mathcal{G})$ we define,*

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

(Similarly for $\mathcal{O}_1(\mathcal{G})$.) Then, if ϕ is irreducible, $\mathrm{Min}(\phi)$ *is non-empty and coincides with the set of points which admit a train track map representing ϕ.*

Remark 7.24. There is also a generalisation of the previous theorem for reducible automorphisms, but in that case $\mathrm{Min}(\phi)$ may be empty in $\mathcal{O}(\mathcal{G})$. In any case $\mathrm{Min}(\phi)$ is never emtpy if we add to $\mathcal{O}(\mathcal{G})$ the simplicial bordification at infinity.

While we don't use the following in this paper, it seems worthwhile mentioning that $\mathrm{Min}(\phi)$ coincides with the set of points supporting partial train-tracks (which reduce to classical train-tracks in irreducible case). (See [11, 12] for more details).

Lemma 7.25. *Let $[\phi] \in \mathrm{Out}(\mathcal{G})$ be an irreducible element and let X be a minimally displaced point. Let $f : X \to X$ be an optimal map representing ϕ, then*

1. *The tension graph of f is the whole X.*
2. *If f is train track then it is a minimal optimal map.*

Proof. By [12, Lemma 4.16] (see also[11]), since f is an optimal map representing ϕ, its tension graph contains an invariant sub-graph ϕ. By irreducibility, that sub-graph must be the whole X. In [12] (or also in [11]) it is proved that if f is a train track, and it is not minimal, then any neighborhood of X in the $\mathrm{Min}(\phi)$ supports an optimal map representing ϕ whose tension graph is not the whole X, contradicting point (1). ∎

References

[1] Hyman Bass, *Covering theory for graphs of groups.*, J. Pure Appl. Algebra, 89(1-2):3-47,1993.

[2] 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

[3] Mladen Bestvina and Mark Feighn, *Hyperbolicity of the complex of free factors*, Adv. in Math. **256**, 2014, 104–155.

[4] Mladen Bestvina and Michael Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. (2) **135** (1992), no. 1, 1–51. MR1147956

[5] Mladen Bestvina and Patrick Reynolds *The boundary of the complex of free factors*, Duke Math. J. **164**, (2015), no. 1 2213–2251.

[6] Benjamin Brück and Radhika Gupta *Homotopy type of the complex of free factors of a free group*, preprint arXiv:1810.09380

[7] James Conant, Martin Kassabov, and Karen Vogtmann, *Hairy graphs and the unstable homology of $Mod(g,s)$, $Out(F_n)$ and $Aut(F_n)$*, J. Top. **6** (2013), no. 1, 119–153.

[8] Marc Culler and Karen Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), no. 1, 91–119. MR 830040

[9] Stefano Francaviglia and Armando Martino, *Metric properties of outer space*, Publ. Mat. **55** (2011), no. 2, 433–473. MR2839451

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

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

[12] Stefano Francaviglia and Armando Martino, *Displacements of automorphisms of free groups I: Displacement functions, minpoints and train tracks*, preprint, arXiv:1807.02781.

[13] Stefano Francaviglia and Armando Martino, *Displacements of automorphisms of free groups II: Connectedness of level sets*, preprint, arXiv:1807.02782.

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

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

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

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

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

[19] Michael Handel and Lee Mosher, *The free splitting complex of a free group, II: Loxodromic outer automorphisms*, to appear on Trans. of AMS.

[20] Allen Hatcher, *Homological stability for automorphism groups of free groups*, Comment. Math. Helv. **70** (1995), no. 1, 39–62. MR 1314940

[21] Arnaud Hilion and Camille Horbez, *The hyperbolicity of the sphere complex via surgery paths*, J. Reine Angew. Math. **730** (2017), 135–161.

[22] 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.

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

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

[25] Ilya Kapovich and Kasra Rafi, *On hyperbolicity of free splitting and free factor complexes*, Groups, Geometry, and Dynamics, **8** (2014), no. 2, 391–414.
[26] 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
[27] Martin Lustig, *Conjugacy and centralizers for iwip automorphisms of free groups*, Geometric group theory, Trends Math., Birkhäuser, Basel, 2007, pp. 197–224. MR 2395795
[28] Sebastian Meinert, *The Lipschitz metric on deformation spaces of G-trees*, Algebr. Geom. Topol. **15** (2015), no. 2, 987–1029. MR 3342683
[29] Jean-Pierre Serre. *Trees.*, Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.
[30] John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565. MR 695906

S. Francaviglia
Dipartimento di Matematica of the University of Bologna
E-mail: stefano.francaviglia@unibo.it

A. Martino
Mathematical Sciences, University of Southampton
E-mail: A.Martino@soton.ac.uk

D. Syrigos
Mathematical Sciences, University of Southampton
E-mail: D.Syrigos@soton.ac.uk