UNIVERSITY OF SOUTHAMPTON

Faculty of Social Sciences
School of Mathematics

Groups acting on Graphs: Their Automorphisms and their Length Functions

by

Matthew Peter Collins

ORCiD: 0009-0003-7171-2558

A thesis for the degree of Doctor of Philosophy

January 2024

University of Southampton

Abstract

Faculty of Social Sciences
School of Mathematics

Doctor of Philosophy

Groups acting on Graphs: Their Automorphisms and their Length Functions

by Matthew Peter Collins

Actions on trees are powerful tools for understanding the structure of a group and its automorphisms. In this thesis, we generalise several existing results in this field to larger classes of groups.

This is a three paper thesis; the main body of the work is contained in the following papers:

[1] Matthew Collins. Fixed points of irreducible, displacement one automorphisms of free products. Preprint, May 2023, available at `arXiv:2305.01451`.

[2] Matthew Collins. Growth and displacement of free product automorphisms. Preprint, July 2023, available at `arXiv:2307.13502`.

[3] Matthew Collins and Armando Martino. Length functions on groups and actions on graphs. Preprint, July 2023, available at `arXiv:2307.10760`.

In [1], we prove that an irreducible, growth rate 1 automorphism of a free product fixes a single point in outer space. This can be thought of as a generalisation of Dicks & Ventura's classification of the irreducible, growth rate 1 automorphisms of free groups.

It is well known for an *irreducible* free group automorphism that its growth rate is equal to the minimal Lipschitz displacement of its action on Culler-Vogtmann space. This follows as a consequence of the existence of train track representatives for the automorphism. In [2], we extend this result to the general - possibly reducible - case as well as to the free product situation where growth is replaced by 'relative growth'.

In [3], we study generalisations of Chiswell's Theorem that 0-hyperbolic Lyndon length functions on groups always arise as based length functions of the group acting isometrically on a tree. We produce counter-examples to show that this Theorem fails if one replaces 0-hyperbolicity with δ-hyperbolicity. We then propose a set of axioms for the length function on a finitely generated group that ensures the function is bi-Lipschitz equivalent to a (or any) length function of the group acting on its Cayley graph.

Contents

List of Figures

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

(1) This work was done wholly or mainly while in candidature for a research degree at this University;

(2) Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;

(3) Where I have consulted the published work of others, this is always clearly attributed;

(4) Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;

(5) I have acknowledged all main sources of help;

(6) Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;

(7) Parts of this work have been published as:

[1] Matthew Collins. Fixed points of irreducible, displacement one automorphisms of free products. Preprint, May 2023, available at `arXiv:2305.01451`.

[2] Matthew Collins. Growth and displacement of free product automorphisms. Preprint, July 2023, available at `arXiv:2307.13502`

[3] Matthew Collins and Armando Martino. Length functions on groups and actions on graphs. Preprint, July 2023, available at `arXiv:2307.10760`.

Signed:.. Date:..................

Acknowledgements

First, thank you to my supervisor Armando for his knowledge, his kindness, his optimism, and above all his patience.

Thank you to Sam and Lily for making my office a wonderful environment to work in. You will be pleased to know that the mice with pool are doing very well. Also, thank you to all of my fellow PhD students for your company, support, and the increasingly complex games of noughts and crosses. I only wish the pandemic had not separated us.

Speaking of the pandemic - thank you to my friends outside of university who kept me sane when the world ground to a halt.

Finally, thanks to my parents and my favourite sister, without whom I would not nearly be as clever, kind, witty, or humble.

Definitions and Abbreviations

$\mathbb{Z}$	The integers
$\mathbb{R}$	The real numbers
F_n	The free group on n generators
$G * H$	The free product of G and H
$\mathrm{Aut}(G)$	The automorphism group of G
$\mathrm{Inn}(G)$	The inner automorphisms of G
$\mathrm{Out}(G)$	The outer automorphism group of G, $\mathrm{Aut}(G)/\mathrm{Inn}(G)$
$\mathrm{Out}(G, \mathcal{G})$	The subgroup of $\mathrm{Out}(G)$ preserving the set of subgroups $\mathcal{G}$
G_x or $\mathrm{Stab}(x)$	The subgroup of G stabilising x
$l_T(g)$	The translation length of g in an action on T
CV_n	Culler-Vogtmann space for the free group F_n
$\mathcal{O}(G, \mathcal{G})$	The outer space determined by a group G and the set of subgroups $\mathcal{G}$
$\mathrm{Min}(\alpha)$	The minimally displaced set (in CV_n or $\mathcal{O}(G, \mathcal{G})$) of an automorphism

Introduction

In this introduction we provide background material and context for the three papers that form the main body of the thesis. Note that the ordering of the papers within the thesis does not reflect publication or posting dates but rather the order I worked on them during my PhD.

Papers 1 and 2 are single author papers; Paper 3 is a joint paper with Armando Martino (my supervisor).

I wrote the majority of the first draft of Paper 3; the remaining statements and proofs were formulated and improved together over the course of several meetings. (For example, Armando wrote the final version of Proposition 3.7, whereas I wrote the final version of Theorem 4.8)

All three papers are concerned with the action of groups on metric graphs, and the main result of each paper is a generalisation of an existing result applied to a larger class of groups. Papers 1 and 2 revolve around Bass–Serre theory, making use of the duality between a G-tree and its quotient graph of groups. Paper 3 is more concerned with the metric on these trees - specifically the conditions required for a group to admit a "graph-like" length function.

Sections 1 to 7 of this introduction cover the required background material for this thesis. Afterwards we have three sections devoted to explaining the original results of each paper, and a small section where we discuss some open questions arising from them. After the introduction, we give the three papers themselves.

1 Groups

Let X be a set. A group $F(X)$ is said to be the *free group* on X if there exists a map $\iota : X \longrightarrow F(X)$ (the inclusion map) such that, for any group H and any function $f : X \longrightarrow H$ there exists a unique homomorphism $\phi : F(X) \longrightarrow H$ such that $f(x) = \phi \circ \iota(x)$, $\forall x \in X$. This behaviour can be represented by the following commutative diagram:

$$\begin{tikzcd} X \arrow[r, "\iota"] \arrow[dr, "f"'] & F(X) \arrow[d, dashed, "\exists!\phi"] \\ & H \end{tikzcd}$$

This defining property of free groups is called the *universal property*.

Alternatively, and perhaps more usefully for our purposes, there is a common construction of free groups which can itself serve as a definition. We say that X is an *alphabet*, and by a *word* in X we mean a string $\xi_1 \ldots \xi_w$ for some $\xi_1, \ldots, \xi_k \in X$. A word can be *reduced* by removing a pair $\xi\xi^{-1}$ or $\xi^{-1}\xi$ for some $\xi \in X$, and a *reduced word* is a word which contains no such pairs. We count the trivial word, denoted by 1, as a reduced word. It can be shown that every word can be turned into a unique reduced word by applying a series of reductions - thus we can think of $F(X)$ as the group of reduced words, where we take the operation to be concatenation followed by reduction.

It is a well-known theorem that two free groups $F(X), F(Y)$ are isomorphic if and only if $|X| = |Y|$. Thus, when $|X| = r \in \mathbb{Z}$, we write F_r to denote the free group of *rank* r.

The concept of free groups leads naturally into that of *free products*, which can also be defined using a universal property:

Let $\{G_\alpha\}$ be a family of groups. A group $\ast\, G_\alpha$ is said to be the *free product* of the family $\{G_\alpha\}$ if there exist (inclusion) maps $\iota_\alpha : G_\alpha \longrightarrow F$ such that, for any group H and any family of homomorphisms $\phi_\alpha : G_\alpha \longrightarrow H$ there exists a unique homomorphism $\psi : F(X) \longrightarrow H$ such that the following diagram commutes for all α.

$$\begin{tikzcd} G_\alpha \arrow[r, "\iota_\alpha"] \arrow[dr, "\phi_\alpha"'] & \ast\, G_\alpha \arrow[d, dashed, "\exists!\psi"] \\ & H \end{tikzcd}$$

Once again, there is a way to intuitively construct a free product. By a *word* in $\ast\, G_\alpha$ we mean a string $g_1 \ldots g_k$ where each g_i lies in some G_α. A word can be *reduced* by either

- removing an instance of the identity element of some G_α, or
- if two adjacent letters g_i, g_{i+1} are from the same G_α, we replace the pair $g_i g_{i+1}$ with its product in G_α.

Once again, it can be shown that every word can be turned into a unique reduced word by applying a sequence of reductions. Thus we think of $\ast\, G_\alpha$ as the group of reduced words, where we take the operation to be concatenation followed by reduction.

Papers 1 and 2 of this thesis are focused on taking existing results regarding free groups - specifically in relation to their actions on trees - and generalizing them to free products.

We will be focusing on free products of the form $G_1 * \ldots * G_k * F_r$, where $k + r \geqslant 1$ and where, if $k \geqslant 1$, the G_i's are non-trivial. If F_r is trivial, then we omit it from the notation.

Remark 1.1. We must acknowledge this free product's similarity to the Grushko decomposition:

Theorem 1.2. *Any finitely generated group G can be decomposed as a free product $G = G_1 * \ldots * G_k * F_r$, where the G_i are non-trivial, freely indecomposable and not infinite cyclic, and F_r is a free group of rank r. Further, the G_i are unique up to conjugacy, and the rank of F_r is unique.*

This decomposition theorem is a well-known consequence of Grushko's theorem on the rank of free products [22, 24] and the Kurosh subgroup theorem; see for example [27]. We, however, are ***not*** restricting ourselves to Grushko decompositions of groups. The only condition we impose upon the G_i's is non-triviality. This means that the valid decompositions of a group are no longer unique.

In addition, if one of our G_i's is itself a free group, then we can relabel it as the final group F_r instead, and we consider this relabelled decomposition to be distinct from the original because it gives rise to a different free factor system (See Example 5.13).

Example 1.3. Let $G = \langle a, b, c \rangle$, a free group of rank 3. There are several possible decompositions of this group, including but not limited to:

- The trivial free product decomposition consisting of a single group $G = \langle a, b, c \rangle$. Since we adopt the convention that trivial free groups F_0 are omitted from the notation, this decomposition is actually two different decompositions in disguise:
  - We can take $k = 1$ and $r = 0$ - so $G_1 = \langle a, b, c \rangle$ is the entire group, and $F_0 = 1$ is omitted.
  - Alternatively, we can take $k = 0$ and $r = 3$ - so $F_3 = \langle a, b, c \rangle$ is the entire group, and there are no G_i's.
- The free product decomposition $G = \langle a \rangle * \langle b \rangle * \langle c \rangle$. Since every group in this decomposition is free, there are several different ways of labelling this decomposition, such as:
  - $G_1 = \langle a \rangle$, $G_2 = \langle b \rangle$, $G_3 = \langle c \rangle$ - so $k = 3$, $r = 0$.
  - $G_1 = \langle a \rangle$, $G_2 = \langle b \rangle$, $F_1 = \langle c \rangle$ - so $k = 2$, $r = 1$.
  - $G_1 = \langle a \rangle$, $F_1 = \langle b \rangle$, $G_2 = \langle c \rangle$ - so $k = 2$, $r = 1$ once again, but this time we have take a different free factor to be F_r.

2 Graphs and trees

Graphs

The focus of papers 1 and 2 is the study of the actions of free products on metric trees. There exist several definitions of trees and graphs in the literature. We shall use the one attributed to Serre.

Definition 2.1. [8, p.113] A *(Serre) graph* Y consists of the following:

- Two disjoint sets $V(Y)$ and $E(Y)$, called the *vertex* and *edge* sets of Y respectively.
- A function $^{-} : E(Y) \to E(Y)$, called *involution*, such that for all $e \in E(Y)$, $\bar{e} \neq e$ and $\bar{\bar{e}} = e$.
- A function $\iota : E(Y) \to V(Y)$, and another function $\tau : E(Y) \to V(Y)$ defined by $\tau e := \iota\bar{e}$. We call ιe the *initial vertex* of e, and τe the terminal vertex of e.

We say Y is *finite* if $V(Y)$ and $E(Y)$ are both finite.

A Serre graph is a combinatorial object, but one usually thinks of it geometrically as a CW-complex. The vertices are 0-cells, the edge pairs are 1-cells, and ι and τ are the attaching maps. When visualising a graph in this way, each element of the edge pair $\{e, \bar{e}\}$ corresponds to travelling in different directions along the corresponding 1-simplex. Note that these CW-complexes are not necessarily simplicial complexes, since they can contain looping edges or multiple edges between a single pair of vertices.

One can also think of a graph as a metric space by identifying each 1-simplex e with a closed, non-trivial, real interval $[0, L_e]$ and then taking the path metric. We say that L_e is the *length* of e, and we impose the condition that $L_e = L_{\bar{e}}$ for all e.

Remark 2.2. The $\{e, \bar{e}\}$ notation is only useful to us when we are specifying the orientation of an edge, which is mostly restricted to our work on Bass-Serre Theory. If we are only concerned with the *length* of our edges (for example, in Paper 3), we will generally avoid writing $\bar{e}$.

Definition 2.3.

- We say that an edge e is *incident* to a vertex v if $\iota e = v$ or $\tau e = v$.
- The *valence* or *degree* of a vertex v is the number of edges e such that $\iota e = v$.

Trees

Much of this thesis works with the metric spaces known as *metric simplicial trees*. These are a type of graph, but one can define them without using Serre's combinatorial structure:

Definition 2.4. An $\mathbb{R}$-*tree* is a non-empty metric space in which any two points are joined by a unique arc, and in which every arc is isometric to a closed interval in the real line.

Definition 2.5. Let p be a point in a non-trivial $\mathbb{R}$-tree T.

- p is called a *branch point* if $T - p$ has three or more components.
- p is called *regular* if $T - p$ has exactly two components.
- p is called *external* otherwise

Points which are not regular are called *non-regular*.

Definition 2.6. A *metric simplicial tree* is an $\mathbb{R}$-tree whose set of non-regular points is discrete.

Remark 2.7. A subset of a metric space is *discrete* if the subspace topology is the discrete topology. One might instead define a metric simplicial tree as an $\mathbb{R}$-tree with a global lower bound on the distance between non-regular points. This is a stronger condition than discreteness, but it is equivalent in the presence of an isometric, cocompact group action.

This approach to building metric simplicial trees is a purely metric approach, with no combinatorial structure. However, the structure of a Serre graph is a useful tool, so we give an additional, second construction:

Definition 2.8. Given a Serre graph Y, we say an *edge path* is a sequence of edges $e_1 \dots e_n$ in $V(Y)$ such that $\tau e_i = \iota e_{i+1}$ for all $1 \leqslant i \leqslant n-1$.

A *cycle* is an edge path where in addition $\tau e_n = \iota e_1$.

A edge path or cycle is called *reduced* if no pair $e_i e_{i+1}$ is of the form $e\bar{e}$.

Definition 2.9. We say a Serre graph Y is *connected* if for every pair of vertices $v \neq w$ in $V(Y)$, there is an edge path with $\iota e_1 = v$ and $\tau e_n = w$.

Definition 2.10. A (Serre) *forest* is a graph with no cycles; a (Serre) *tree* is a connected forest.

To what degree are these two constructions equivalent? It is not hard to see that metric Serre trees built in this way are $\mathbb{R}$-trees as described in Definition 2.4; the external points are the vertices of degree 1 (or 0, in the case of the trivial tree), the branch points are the vertices of degree 3 or more, and all other points are regular. Furthermore, if the vertex set of a metric Serre tree is discrete, then it will be a metric simplicial tree.

Definition 2.11. Let Y, Y' be Serre graphs. A *graph morphism* from Y to Y' is a continuous map $Y \to Y'$ which sends vertices to vertices and edges to edges. A *graph isomorphism* is a bijective graph morphism.

(The word "continuous" here implies that we are visualising our graphs as CW-complexes.)

Definition 2.12. Let X, X' be metric spaces. An *isometry* from X to X' is a map $f : X \to X'$ such that for all $x_1, x_2 \in X$, $d_X(x_1, x_2) = d_{X'}(f(x_1), f(x_2))$. An *isomorphism of metric spaces* is a bijective isometry. If such a map exists, we say the two spaces are *isometric*.

A metric simplical tree can be thought of as the underlying space of a 1-dimensional simplicial complex, hence the name. This is in turn a 1-dimensional CW-complex, and hence it is a Serre tree.

There are some natural ways to construct a simplicial structure on a metric simplicial tree which will be covered in the next chapter. Crucially, however, this structure/Serre tree is not unique. Given any Serre tree T, one may create a new Serre tree by subdividing any edge pair and introducing a new vertex. Conversely, one may remove a degree 2 vertex by combining two edges. Assuming one adjusts the edge lengths correctly, these operations produce trees which are isometric (as metric spaces) but not, in general, isomorphic as graphs.

To put it another way, a "metric simplical tree" is simply a name we use to describe an entire isometry class of metric Serre trees. If we want to talk about the edges and vertices of a simplicial tree, we must choose a representative Serre tree from this isomorphism class. As we will show in the next chapter, the representative we choose will be motivated by the action of a group on our simplicial tree.

3 Groups Acting on Graphs

Definition 3.1. Let G be a group, and let X be a set. We say that G *acts* on X *on the right* if there exists a map $X \times G \to X, (x, g) \mapsto x \cdot g$ such that for all $x \in X$,

(i) $x \cdot 1 = x$, and

(ii) $\forall g, h \in G$, $(x \cdot g) \cdot h = x \cdot gh$

The definition of a *left action* on X is similar; the only difference being the order of composition of the group elements in (ii). Our notation will represent this by writing, for example, $g \cdot x$ instead of $x \cdot g$.

For the sake of brevity, we may omit the symbol $\cdot$ from the notation. Unless stated otherwise, all our actions are assumed to be right actions.

Definition 3.2. Suppose G acts on a metric space X. We say this action is *isometric* if $\forall x, y \in X, \forall g \in G, d_X(x, y) = d_X(x \cdot g, y \cdot g)$.

Suppose a group G acts isometrically on a metric simplicial tree T. As stated in the previous chapter, every metric simplicial tree is isomorphic to a whole family of metric Serre trees. Our goal here is to find which of these Serre trees has the structure which is most "useful" for studying the action. We start by giving some of the properties we find desirable:

Definition 3.3. Let Y be a Serre graph. A *subgraph* of Y is a graph Y' such that $V(Y') \subseteq V(Y)$, $E(Y') \subseteq E(Y)$, and such that for all $e \in E(Y')$, we have $\bar{e} \in E(Y')$ and $\iota e, \tau e \in V(Y')$.

A subgraph is called a *subforest* if it is a disjoint union of Serre trees.

Every subgraph of a Serre tree is a *subforest*. A connected subforest of a Serre tree is called a *subtree*.

Definition 3.4. Suppose G acts on a metric simplicial tree T, and suppose we have chosen an isomorphic metric Serre tree whose structure we can apply to T.

- The action is said to be *simplicial* if it maps vertices to vertices and edges to edges. (Equivalently, we can say that G acts *via graph automorphisms*)
- If no edge of T is sent to its inverse by any element of G, we say that G *acts without inversions*.
- We say that both T and the action are *minimal* if T contains no proper, G-invariant subtree.
-
  - Let $x \in T$. The *stabiliser* $\mathrm{stab}(x)$ of x is defined to be the subgroup $\{g \mid x \cdot g = x\}$ of G. More generally, given any subset A of T, we have $\mathrm{stab}(A) = \{g \in G \mid A \cdot g = A\}$. Note that this definition is setwise, not pointwise; for an edge e of T the stabiliser of e is the subgroup which fixes e but does not necessarily preserve the orientation of e.
  - If every edge in T has trivial stabiliser, we say that T is *edge-free*.
  - Let p be a vertex in T. If $\mathrm{stab}(p) = 1$, we say that p is *free*. Otherwise, it is *non-free*.

We would like our actions to be simplicial, without inversions, minimal and edge-free.

Let T be a metric simplicial tree acted upon by G. Then the "simplest" structure on T - that is, the structure containing the fewest vertices - is obtained by defining the vertex set to be the set of non-regular points of T. One then takes the edge set to be the set of simple arcs between elements of the vertex set which do not contain any other vertices.

Using the simplest structure, the action of G on T is simplicial. However, some edges may be sent to their inverses by elements of G. Future calculations will be easier if we have an action without inversions, therefore we shall instead use the following structure:

- We define the vertex set to be the set of non-regular points of T, together with the midpoints of all the edges of the simplest structure which were inverted by an element of G. We denote this vertex set by $V(T)$.
- We then define the edge set to be the set of simple arcs between elements of the vertex set which do not contain any other vertices. We denote this edge set by $E(T)$.

Essentially, we divide each inverted edge into two edges by placing a new vertex at its midpoint. With this new structure the action is still simplicial and, in addition, it is without inversions.

As for edge freeness - this is not a condition we can guarantee from our choice of simplicial structure. It is a condition we will have to impose ourselves.

Definition 3.5. A *G-tree* is a triple $(T, d_T, \cdot)$, where T is a metric simplicial tree, d_T is the metric on T, and $\cdot$ is an isometric group action $T \times G \to T,\ (x, g) \mapsto x \cdot g$.

If the metric and action are obvious from context, we may choose to omit one or both of them from the notation.

We will always give G-trees the specific simplicial structure described above.

We list some of the properties of G-trees which arise from our choice of structure. These results are stated without proof in Paper 1, but here we take the opportunity to provide more detail.

Remark 3.6. Note that we have chosen to define G-trees with a *right* action. This ensures that

- $\forall x \in T, \forall g \in G, \mathrm{stab}(x \cdot g) = \mathrm{stab}(x)^g$
- $\forall H \leqslant G, \forall g \in G, \mathrm{Fix}(H^g) = \mathrm{Fix}(H) \cdot g$

Had we chosen to act on the left, acting by g would have caused the stabilisers to be conjugated by g^{-1}, and similar for the fixed point sets.

Proposition 3.7. *Minimal G-trees do not contain any degree 1 vertices (and hence the set of non-regular points is exactly the set of branch points).*

Proof. Let T be a G-tree. If T is trivial, the result immediately follows, so we can assume there exists an edge e of T with incident vertices v and w. Suppose that v is degree 1. Then, since our action is simplicial and isometric, every vertex in the orbit of v will also be degree 1, and eG will be the set of edges incident to this orbit. Thus $T - (vG \cup eG)$ is connected.

Suppose that $w \in vG$. Then w is also degree 1, and hence the element of G which mapped v to w would also invert e. Since we chose our simplicial structure such that this does not occur, this is a contradiction. Thus $w \in T - (vG \cup eG)$, and hence $T - (vG \cup eG)$ is non-empty.

It is clear that vG and eG are G-invariant. Therefore $T - (vG \cup eG)$ is a proper G-invariant subtree of T, contradicting the minimality of T. Hence T does not contain any degree 1 vertices. □

Proposition 3.8. *G-trees do not contain any free vertices of degree 2.*

Proof. By definition of a non-regular point, none of the vertices in the simplest structure are degree 2. Therefore the only degree 2 vertices in $V(T)$ must have been added as the midpoints of inverted edges. The elements of G which inverted these edges will fix these midpoints, thus making them non-free vertices. □

4 Bass-Serre Theory

Bass-Serre theory is the study of the relationship between a G-tree and its quotient graph of groups.

Definition 4.1 (Graph of Groups)**.** [8, p.198] A *graph of groups X* consists of:

(i) A connected (Serre) graph Y.

(ii) A group G_v for each vertex v of Y, and a group G_e for each edge e of Y such that $G_{\bar{e}} = G_e$.

(iii) For each (oriented) edge e of Y, a monomorphism $\rho_e : G_e \to G_{\tau e}$.

If Y is a metric graph, then we say X is a *metric graph of groups*.

Remark 4.2. We are using Serre's notation for edge pairs here, so (iii) implies the existence of monomorphisms $\rho_{\bar{e}} : G_e \to G_{\iota e}$.

As mentioned in the previous chapter, we will be working exclusively with edge-free G-trees, which will correspond to graphs of groups with trivial edge groups (and hence trivial monomorphisms). With this in mind, we will restrict our explanations of the fundamentals of Bass-Serre Theory to this subset of G-trees and graphs of groups. Bass-Serre Theory in the more general case has been covered extensively by the existing literature, for example in [8] and [15].

Let X be a graph of groups with trivial edge groups on a graph Y. One can define the fundamental group of X in a similar manner to that of a standard graph, by thinking of elements of the group as reduced loops in the graph. However, some additional structure is added by the edge and vertex groups.

By a *path* in X, we mean an alternating string $g_0e_1g_1 \dots g_{n-1}e_ng_n$ where for all i, $g_{i-1} \in G_{\iota e_i}$ and $g_i \in G_{\tau e_i}$. (If some $g_i = 1$, then we usually omit it from the notation.) The set of all paths X is a groupoid $\Pi(X)$ where the operation is concatenation; two paths $g_0e_1g_1 \dots g_{n-1}e_ng_n$, $g'_0e'_1g'_1 \dots g'_{n-1}e'_mg'_m \in \Pi(X)$ can be concatenated if and only if $\tau e_n = \iota e'_1$.

We say a path $g_0e_1g_1 \dots g_{n-1}e_ng_n \in \Pi(X)$ is a *loop* if $\iota e_1 = \tau e_n$.

We say a path in $\Pi(X)$ is *reduced* if it does not contain a subpath of the form $e\bar{e}$ (or $e(gg^{-1}\bar{e})$).

Definition 4.3. Let X be a graph of groups on a graph Y. Then we define the *fundamental group* $\pi_1(X, v)$ of X to be the group of reduced loops in X which start and end at a particular vertex v of Y, called the base point. It can be shown that the isomorphism class of this group does not depend on our choice of base point; therefore, unless v is required, we shall simply denote the group by $\pi_1(X)$.

We will be working with graphs of groups whose fundamental group is isomorphic to a particular group G. Thus we consider pairs (X, ϕ), where X is a graph of groups and $\phi : G \to \pi_1(X)$ is an isomorphism. Such a pair is called a *marked graph of groups*, and ϕ is called the *marking*.

The Quotient Graph of Groups

Given a G-tree T, we can construct from it a marked, metric graph of groups called the *quotient graph of groups*. A method for constructing a quotient graph of groups from an arbitrary connected graph acted upon by G can be found on pages 204-205 of [8]. We shall restrict this construction to edge-free G-trees.

Let T be a G-tree. Take the quotient graph T/G, let $p : T \to T/G$ be the projection map, and Y_0 a maximal tree of T/G. Let $j : Y_0 \to T$ be a map such that $p \circ j$ is the identity on Y_0 (i.e j is a lift of Y_0 to T). We call $j(Y_0)$ a *representative tree* for the action.

We then define a graph of groups X on T/G as follows: For any vertex x of T/G, we define the vertex group G_x to be $\mathrm{stab}(j(x))$. We take all edge groups, and hence all edge monomorphisms, to be trivial, and edges inherit their lengths from T. This completely defines X.

Theorem 4.4. *[8, p.210] Let T be a G-tree, and let X be a quotient graph of groups for T. Then the fundamental group of X is isomorphic to G.*

This isomorphism gives a marking on X, and hence we can think of the quotient graph of groups as a marked graph of groups.

The Universal Cover

Conversely, let (X, ϕ) be a marked metric graph of groups with trivial edge groups and with $G \overset{\phi}{\cong} \pi_1(X, v)$. Then we can construct from X a G-tree called the *Bass-Serre tree*, or *universal cover* of X, denoted by $\tilde{X}$.

We shall denote the universal cover by $\widetilde{X}$, and we shall begin by defining it as a graph. Recall the definition of the path groupoid $\Pi(X)$. Within this groupoid we can take 'cosets' $G_w\gamma$ of the vertex groups, where $\gamma = g_0e_1g_1 \dots g_{n-1}e_ng_n$ is a path from w to v (i.e. $w = \iota e_1$ and $v = \tau e_n$). Without loss of generality, we shall always choose the representative γ of a coset such that $g_0 = 1$. We take the vertex set $V(\widetilde{X})$ to be the set of all these cosets.

To define the edge set $E(\widetilde{X})$, we say two vertices $G_{w_1}\gamma_1$ and $G_{w_2}\gamma_2$ are joined by an edge pair if:

- the vertices w_1 and w_2 are joined by an edge pair $(e, \bar{e})$ in X, and
- $\gamma_1 = eg_{w_2}\gamma_2$ or $\gamma_2 = \bar{e}g_{w_1}\gamma_1$.

We define the length of the edge pair in $\widetilde{X}$ to be the length of e.

It can be shown that $\widetilde{X}$ is a tree [8, p.206-207]. Thus we can think of this graph as a simplicial structure.

We define a right action of $\pi_1(X, v)$ on $\widetilde{X}$ via the multiplication in $\Pi(X)$. Since each γ ends at v, and elements of $\pi_1(X, v)$ are loops at v, this action is well-defined and respects adjacency. We can then induce an action of G on $\widetilde{X}$ via ϕ, making $\widetilde{X}$ a G-tree.

Equivalence and the Fundamental Theorem

Definition 4.5. Let $(T, d_T, \cdot)$, $(S, d_S, *)$ be G-trees. We say a map of trees $f : T \to S$ is a *G-equivariant* map from $(T, d_T, \cdot)$ to $(S, d_S, *)$ if $f(x \cdot g) = f(x) * g$ for all $x \in T$, for all $g \in G$.

Definition 4.6. Two G-trees $(T, d_T, \cdot)$, $(S, d_S, *)$ are said to be *equivalent* if there exists a G-equivariant isometry between them. We write $(T, d_T, \cdot) \sim (S, d_S, *)$ to denote equivalence.

Definition 4.7. We say that two marked metric graphs of groups are *equivalent* if their universal covers are equivalent G-trees.

Theorem 4.8 (Fundamental Theorem of Bass-Serre Theory)**.** *The process of lifting to the universal cover and the process of descending to a quotient graph of groups are mutually inverse, up to equivalence of the structures involved.*

This theorem allows us to think of a G-tree/graph of groups as a single object with two different forms. We can move between these two forms as much as we want, depending on which structure is most useful at the time.

Definition 4.9. We say a map of G-trees is *simplicial* if it maps vertices to vertices. Note that it does not have to map edges to edges, and hence this definition differs from that of a simplicial group action.

5 Free Factor Systems and Deformation Spaces

Let G be a group, and let T be a G-tree.

Definition 5.1. An element $g \in G$ is said to be *elliptic* (with respect to T) if it fixes a point in T. If g is not elliptic, we say it is *hyperbolic* (with respect to T).

We shall say a subgroup H of G is *elliptic* (with respect to T) if there exists a point $x \in T$ such that $x \cdot H = x$.

Definition 5.2 (Free Factor System)**.** Let T be a minimal, cocompact, edge-free G-tree, and let $\mathcal{G}_T$ denote the set of elliptic subgroups for T. We say $\mathcal{G}_T$ is a *free factor system* for G.

Note that this is not the usual definition of a free factor system. The usual definition can be found in [4, p.530-531], and we shall refer to it as a *traditional free factor system*:

Definition 5.3 (Traditional Free Factor System)**.** If $G_1 * \ldots * G_k * F_r$ is a free product decomposition for a group G, and each G_i is nontrivial, then we say that the collection $\{[G_1], \ldots, [G_k]\}$ of conjugacy classes is a *traditional free factor system*. The empty set $\varnothing$ is the trivial traditional free factor system.

Corollory 5.11 will show that the Fundamental Theorem of Bass-Serre Theory provides a natural way to construct a free factor system from a traditional free factor system, and vice versa, and that these constructions are mutually inverse. In this sense, the two definitions are equivalent.

We have chosen our definition for two reasons. Firstly, the trivial free factor system must be defined separately when using the traditional definition. Secondly, our definition allows us to order our free factor systems by inclusion, which is equivalent to the somewhat more complicated ordering used for the traditional free factor systems, as defined in [4, p.532].

For now, we return to our definition of a free factor system and we observe the following properties:

Lemma 5.4. *Free factor systems are closed under conjugation and taking subgroups.*

Lemma 5.5. *Let $(T,\cdot),(S,*)$ be equivalent minimal, cocompact, edge-free G-trees. Then $\mathcal{G}_T = \mathcal{G}_S$.*

Lemma 5.6. *Let $(T,\cdot)$ be an edge-free G-tree. A nontrivial element of G cannot fix more than one point in T, and the fixed point will always be a vertex.*

The final lemma tells us that a subgroup of G is elliptic with respect to a minimal, cocompact, edge-free G-tree, and hence is an element of the corresponding free factor system, if and only if it is a vertex stabiliser or a subgroup of a vertex stabiliser.

Definition 5.7. Let $\mathcal{G}$ be a free factor system for G. The *deformation space* $\mathcal{O} = \mathcal{O}(G,\mathcal{G})$ is the space of equivalence classes of minimal, cocompact, edge-free G-trees T such $\mathcal{G}_T = \mathcal{G}$.

By Lemma 5.6, a subgroup of G is in $\mathcal{G}$ if and only if it is a vertex stabiliser or subgroup of a vertex stabiliser for some (and hence every) G-tree in $\mathcal{O}$. Additionally, by general properties of group actions, two vertices lie in the same orbit if and only if they have conjugate stabilisers. Thus we can define a *minimal generating set* for $\mathcal{G}$:

Definition 5.8. We say a subset of a free factor system $\mathcal{G}$ is a *minimal generating set* for $\mathcal{G}$ if and only if it contains exactly one vertex stabiliser from every orbit of non-free vertices in some (and hence every) $T \in \mathcal{O}$.

Remark 5.9. If $\mathcal{G}$ is non-trivial, then this definition is the conventional definition of a minimal generating set under the operations of conjugation and taking subgroups.

If $\mathcal{G} = \{1\}$ - the trivial free factor system - then trees in $\mathcal{O}(G,\mathcal{G})$ do not contain any non-free vertices. Thus the minimal generating set for $\mathcal{G} = \{1\}$ is the empty set $\varnothing$.

G-trees in $\mathcal{O}$ are cocompact; in particular, they each have a finite number of vertex orbits. Hence a minimal generating set for $\mathcal{G}$ will always be finite.

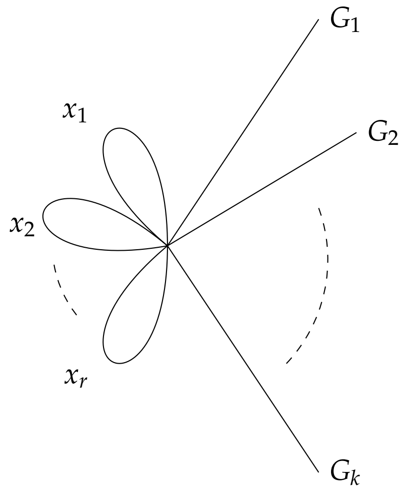

FIGURE 1: A graph of groups with fundamental group $G_1 * \ldots * G_k * \langle x_1, \ldots, x_r \rangle$

Theorem 5.10. *The following are equivalent:*

(i) There exists a minimal, cocompact, edge-free G-tree T containing a representative tree T_0 in T such that the non-trivial vertex stabilisers in T_0 are exactly $G_1, \ldots, G_k$.

(ii) There exists a minimal, cocompact, edge-free G-tree T and a quotient graph of groups on T/G whose non-trivial vertex groups are exactly $G_1, \ldots, G_k$.

(iii) G can be written as a free product $G = G_1 * \ldots * G_k * F_r$, where F_r is a free group of rank $r \geqslant 0$.

Proof. $(i) \Leftrightarrow (ii)$ Follows from the Fundamental Theorem of Bass-Serre Theory.

$(ii) \Rightarrow (iii)$ Follows immediately from the definition of the fundamental group of a graph of groups and Theorem 4.4.

$(iii) \Rightarrow (ii)$ When $k = 1$ and $r = 0$ we take X to be the graph of groups consisting of a single vertex with vertex group G. Otherwise we can take X to be the graph of groups depicted in Figure 1.

□

We can now demonstrate the correspondence between free factor systems and traditional free factor systems.

Let $\mathcal{G}$ be a free factor system for a group G. This means that $\mathcal{G}$ is the set of elliptic subgroups of some $T \in \mathcal{O}(G, \mathcal{G})$. Take a representative tree in T, and let $\{G_1, \ldots, G_k\}$ be the nontrivial vertex groups of this representative tree. Then by Theorem 5.10, G can be written as a free product $G = G_1 * \ldots * G_k * F_r$. Thus the set $\{[G_1], \ldots, [G_k]\}$ is a traditional free factor system. Since each $[G_i]$ is a conjugacy class, this set does not depend on the choice of representative tree.

Conversely, let $\{[G_1], \ldots, [G_k]\}$ be a traditional free factor system. Then we have $G = G_1 * \ldots * G_k * F_r$, so by Theorem 5.10, there exists a minimal, cocompact, edge-free G-tree T containing a representative tree T_0 in T such that the non-trivial vertex stabilisers in T_0 are exactly $G_1, \ldots, G_k$. These vertex groups give a minimal generating set for a free factor system $\mathcal{G} = \mathcal{G}_T$.

Corollary 5.11. *The process of constructing a free factor system from a traditional free factor system, and the process of constructing a traditional free factor system from a free factor system, are mutually inverse. In this sense the two definitions are equivalent.*

Proof. Follows from Theorem 5.10. □

It follows from Corollary 5.11 that, given a free product $G = G_1 * \ldots * G_k * F_r$, we can construct a free factor system $\mathcal{G}$ and hence a corresponding deformation space $\mathcal{O}(G, \mathcal{G})$. This space can then be used to study the automorphisms of the free product.

Remark 5.12 (Culler-Vogtmann Space)**.** It is a well known result that a group admits a free action on a tree if and only if it is a free group. Of particular interest then is the case when G is a free group of rank n and $\mathcal{G}$ is the trivial free factor system. The deformation space which corresponds to this case is *Culler-Vogtmann space CV_n*, the space of equivalence classes of minimal, cocompact, F_n-trees acted upon freely by F_n. Much of the original work in this thesis revolves around taking existing results in Culler-Vogtmann space and generalising them to arbitrary $\mathcal{O}(G, \mathcal{G})$.

Example 5.13. In Example 1.3, we gave some examples of possible free product decompositions of the group $G = \langle a, b, c \rangle$, a free group of rank 3. We also noted that several of these decompositions may appear the same but would be treated differently depending on which free factor was labelled as F_r. This is because they would give rise to different free factor systems, which we can now demonstrate. Additionally, since each free factor system $\mathcal{G}$ corresponds to an outer space $\mathcal{O}(G, \mathcal{G})$, we can also give some examples of points in these spaces. We shall use graphs of groups to represent these points, since they are easier to represent visually than G-trees.

- We looked at the free product decomposition $G = \langle a, b, c \rangle$ consisting of only a single group, and how it can be expressed in two ways: either by taking $G_1 = \langle a, b, c \rangle$ and $r = 0$, or by taking $F_r = \langle a, b, c \rangle$ and $k = 0$. We can now see that these give distinct free factor systems:
  - First consider the case where $G_1 = \langle a, b, c \rangle$ and $r = 0$. The free factor system $\mathcal{G}$ corresponding to this decomposition is the one generated by all the G_i's - therefore $\mathcal{G}$ is the set of all subgroups of G. The only possible graph of groups is a single vertex with vertex group $\langle a, b, c \rangle$.

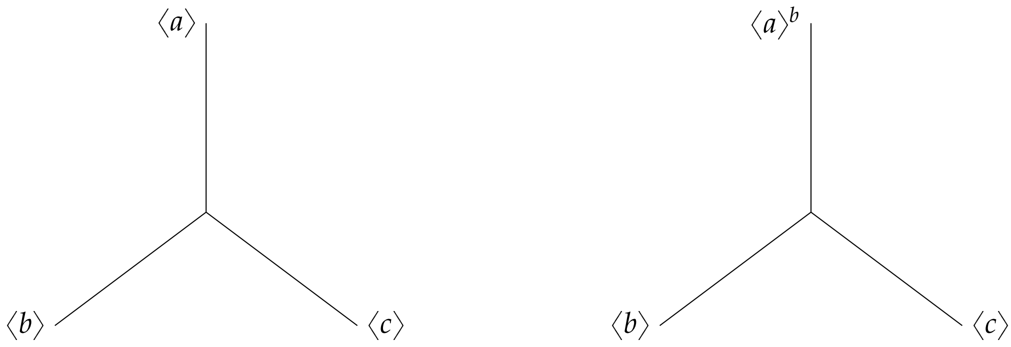

FIGURE 2: Two possible graphs of groups in $\mathcal{O}(G, \mathcal{G}_{abc})$

- Now consider the case where $F_r = \langle a, b, c\rangle$ and $k = 0$. There are no G_i's, so the corresponding free factor system is the trivial free factor system $\varnothing$. The graphs of groups in $\mathcal{O}(G, \varnothing)$ are those with trivial vertex groups and whose underlying graph has fundamental group F_3. Thus $\mathcal{O}(G, \varnothing)$ is in fact the Culler-Vogtmann space CV_3.

- Take the decomposition $G = \langle a\rangle * \langle b\rangle * \langle c\rangle$, and suppose we take
  $G_1 = \langle a\rangle$, $G_2 = \langle b\rangle$, $G_3 = \langle c\rangle$ - so $k = 3$, $r = 0$. This gives rise to the free factor system $\mathcal{G}_{abc}$ generated by $\langle a\rangle$, $\langle b\rangle$ and $\langle c\rangle$. The system $\mathcal{G}_{abc}$ is distinct from the system $\mathcal{G}$ from the first example - observe that the element ab does not appear in any subgroup in $\mathcal{G}_{abc}$.

  We can use this example to demonstrate why we focus on free factor systems and not just free product decompositions. The group G can also be decomposed as $G = \langle a\rangle^b * \langle b\rangle * \langle c\rangle$, where $G_1' = \langle a\rangle^b$, $G_2' = \langle b\rangle$, $G_3' = \langle c\rangle$. This corresponds to the free factor system generated by $\langle a\rangle^b$, $\langle b\rangle$ and $\langle c\rangle$, but it is not hard to see that this is also $\mathcal{G}_{abc}$. Thus the two decompositions produce the same space $\mathcal{O}(G, \mathcal{G}_{abc})$.

  Figure 2 depicts two graphs of groups in $\mathcal{O}(G, \mathcal{G}_{abc})$. The fundamental group of each of these can be written naturally as one of these free product decompositions, but as points in $\mathcal{O}(G, \mathcal{G}_{abc})$ they would each come equipped with a marking - an iso- morphism from the fundamental group to G - so the natural decomposition may not always be used. In fact, when given the correct markings, the two graphs of groups may actually represent the same point in $\mathcal{O}(G, \mathcal{G}_{abc})$.

6 Automorphisms

For the duration of this section, let $\mathcal{G}$ denote a free factor system for a group G, and let $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$.

Acting on the Deformation Space

Notation. The outer automorphism group of G is defined as $\mathrm{Out}(G) := {}^{\mathrm{Aut}(G)}\!\big/_{\mathrm{Inn}(G)}$; elements of $\mathrm{Out}(G)$ are equivalence classes of automorphisms, where two automorphisms are equivalent if they differ by an inner automorphism.

When we write $\alpha \in \mathrm{Out}(G)$, we mean that α is an automorphism in $\mathrm{Aut}(G)$ representing an equivalence class in $\mathrm{Out}(G)$.

In this paper, the automorphisms of G will act on G on the right.

Definition 6.1. Let $\alpha \in \mathrm{Aut}(G)$, and let $\mathcal{G}\alpha = \{(H)\alpha \mid H \in \mathcal{G}\}$. We say that $\mathcal{G}$ is α-*invariant* if $\mathcal{G}\alpha = \mathcal{G}$.

Free factor systems are closed under conjugation by elements of G, hence α-invariance depends only on the outer automorphism class of α. Thus we can make a similar definition for $\mathrm{Out}(G)$:

Definition 6.2. Let $\alpha \in \mathrm{Out}(G)$, and let $\mathcal{G}\alpha = \{(H)\alpha \mid H \in \mathcal{G}\}$. We say that $\mathcal{G}$ is α-*invariant* if $\mathcal{G}\alpha = \mathcal{G}$.

The set of outer automorphisms of G leaving $\mathcal{G}$ invariant forms a group, which we shall denote $\mathrm{Out}(G, \mathcal{G})$.

The group $\mathrm{Out}(G, \mathcal{G})$ admits a left action on the deformation space $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$: Let $(T, d_T, \cdot) \in \mathcal{O}$, let $\alpha \in \mathrm{Out}(G, \mathcal{G})$. Then $\alpha(T, d_T, \cdot) := (T, d_T, \cdot_\alpha)$, the G-tree with the same underlying simplicial tree and metric, but with 'twisted' action given by $x \cdot_\alpha g = x \cdot (g)\alpha$ for all $x \in T$.

Observe the following:

Lemma 6.3. *Let $T \in \mathcal{O}$ be a G-tree, and let $\alpha \in Out(G, \mathcal{G})$. Then the G-orbits of T are the same as those of αT. That is, for all $x \in T$, $x \cdot G = x \cdot_\alpha G$.*

Proof. $x \cdot G = \{x \cdot g \mid g \in G\} = \{x \cdot (g)\alpha \mid g \in G\} = x \cdot_\alpha G$ □

Thus we see that the 'twists' α applies to the action only occur within each orbit. This means that if we are working with both $(T, \cdot)$ and $(T, \cdot_\alpha)$, we are able to simply refer to 'a G-orbit of T' without having to state which action is being used.

Definition 6.4. We can partially order the set of all free factor systems of G by inclusion. Let $\mathcal{G}$ be a proper, α-invariant free factor system. We say $\alpha \in \mathrm{Out}(G, \mathcal{G})$ is $\mathcal{G}$-*irreducible*, or *irreducible with respect to* $\mathcal{G}$, if $\mathcal{G}$ is a maximal, proper α-invariant free-factor system.

Otherwise, we say α is *reducible* with respect to $\mathcal{G}$.

Topological Representatives

Definition 6.5. [19, p.16] Let $T, S \in \mathcal{O}(G, \mathcal{G})$. An *$\mathcal{O}$-map* $f : T \to S$ is a G-equivariant, Lipschitz continuous function. The Lipschitz constant of f is denoted $\mathrm{Lip}(f)$.

Note that an $\mathcal{O}$-map does not have to send vertices to vertices, and hence does not need to be a graph morphism.

Definition 6.6. [19, p.16] We say an $\mathcal{O}$-map $f : T \to S$ is *straight* if it has constant speed on edges - that is, for each edge e in T, there exists a non-negative number $l_e(f)$ such that for any $a, b \in e$ we have $d_T(f(a), f(b)) = l_e(f) d_S(a, b)$.

Definition 6.7. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, and let $T \in \mathcal{O}$ be a G-tree. Then a map $f : T \to \alpha T$ is said to *topologically represent* α if it is a straight $\mathcal{O}$-map.

The authors of [19, p. 16] make the following remark:

Remark 6.8. Any two trees $T, S \in \mathcal{O}$ have an $\mathcal{O}$-map between them. Furthermore, any $\mathcal{O}$-map $f : T \to S$ can be uniquely 'straightened' - that is to say, there exists a unique straight $\mathcal{O}$-map $\mathrm{Str}(f) : T \to S$, such that $\mathrm{Str}\big(f\big)(v) = f(v)$ for every vertex $v \in T$. We have $\mathrm{Lip}(\mathrm{Str}(f)) \leqslant \mathrm{Lip}(f)$.

From this remark it follows that $\forall \alpha \in \mathrm{Out}(G, \mathcal{G})$, $\forall T \in \mathcal{O}$, there exists a topological representative $f : T \to \alpha T$.

Definition 6.9. Let F be a subforest of some $T \in \mathcal{O}$, and let A be a component of F. We define the *stabiliser* of A to be the set $\mathrm{stab}(A) = \{g \in G \mid A \cdot g = A\}$ - that is, we are taking the setwise stabiliser, not the pointwise stabiliser.

We say that F is $\mathcal{G}$-elliptic if, for every component A of F, $\mathrm{stab}(A) \in \mathcal{G}$. Otherwise we say that F is $\mathcal{G}$-hyperbolic.

Using topological representatives, we can construct a test for the reducibility of an automorphism:

Theorem 6.10. *Let $\mathcal{G}$ be a proper free factor system for a group G, let $\alpha \in Out(G, \mathcal{G})$, and let $T \in \mathcal{O}$.*

Suppose that α can be topologically represented by a G-equivariant simplicial map $f : T \to \alpha T$, and there exists a proper f-invariant, G-invariant, $\mathcal{G}$-hyperbolic subforest of T. Then α is reducible with respect to $\mathcal{G}$.

Proof. We shall prove that α is reducible by constructing a new cocompact, minimal, edge-free G-tree S, from which we shall retrieve another proper α-invariant free factor system $\mathcal{H}$ such that $\mathcal{G} \subset \mathcal{H}$.

Let F denote the subforest of T described above, together with all the remaining vertices of T. This extended subforest is still proper, f-invariant, and G-invariant. (In addition, recall that we defined subforests such that they respect the simplicial structures of our G-trees; therefore the complement of F is a set of edges.)

We obtain S from T by collapsing each component A of F to a point p_A. Edges which were not collapsed inherit their lengths from T, giving us a metric on S. Since F is G-invariant, this collapse induces a minimal, isometric action of G on S. Thus S is a G-tree.

Furthermore, if we declare the vertex set to be the set of p_A, we induce a new simplicial structure on S. (This is a well-defined vertex set, and the simplicial structure given by this vertex set is exactly the same as the usual simplicial structure we give to all G-trees).

S inherits cocompactness and edge-freeness from T. Hence, by definition, the set $\mathcal{H}$ of elliptic subgroups for S is a free factor system for G.

Let $v \in T$ be a vertex. Since F contains every vertex of T, v must lie in some component A of F. Therefore, since F is G-invariant, all of A must be fixed (setwise, not necessarily pointwise) by $\mathrm{stab}(v)$. Hence $\mathrm{stab}(v) \leqslant \mathrm{stab}(p_A)$. This holds for all v, which is enough to tell us that $\mathcal{G} \sqsubseteq \mathcal{H}$.

Some component of F has $\mathcal{G}$-hyperbolic stabiliser. This means that $\mathrm{stab}(p_A) \notin \mathcal{G}$ but $\mathrm{stab}(p_A) \in \mathcal{H}$. Thus $\mathcal{G} \sqsubset \mathcal{H}$.

Suppose that $G \in \mathcal{H}$ - that is to say, a vertex of S is stabilised by G. Then the component of F corresponding to this vertex is a G-invariant subtree of T, contradicting the minimality of T. Hence $G \notin \mathcal{H}$, so $\mathcal{H}$ is a proper free factor system.

Finally, we must show that $\mathcal{H}$ is α-invariant:

Let p_A be a vertex of S. We first want to show that $(\mathrm{stab}(p_A))\alpha$ lies in $\mathcal{H}$ - that is to say, it fixes a point in $(S, *)$. F is f-invariant, therefore $f(A)$ is a component of F and $p_{f(A)}$ is a vertex of S. Furthermore, $\forall (g)\alpha \in (\mathrm{stab}(v_A))\alpha$,
$p_{f(A)} * (g)\alpha = p_{f(A)\cdot(g)\alpha} = p_{f(A\cdot g)} = p_{f(A)}$, and hence $(\mathrm{stab}(p_A))\alpha \in \mathcal{H}$. This holds for all $p_A \in S$. This tells us that $\mathcal{H}\alpha \sqsubseteq \mathcal{H}$, which in turn is enough to show that $\mathcal{H} = \mathcal{H}\alpha$.

To summarize, $\mathcal{H}$ is a proper, α-invariant free-factor system for G, and $\mathcal{G} \sqsubset \mathcal{H}$. Hence, by Definition 6.4, α is reducible with respect to $\mathcal{G}$. □

Isometric topological representatives

Our main results will make use of *isometric* topological representatives, which allow us to make some additional observations:

Remark 6.11. Recall that $\mathcal{O}$ is a space of equivalence classes of G-trees, where two trees are equivalent if there exists an equivariant isometry between them. Topological representatives are equivariant; therefore, if an isometric topological representative $f : T \to \alpha T$ exists, the two G-trees T and αT are representing the same point in $\mathcal{O}$.

Proposition 6.12. *Let $f : T \to \alpha T$ be a topological representative for some $T \in \mathcal{O}(G, \mathcal{G})$, for some $\alpha \in Out(G, \mathcal{G})$. If f is an isometry, then it is also a graph automorphism.*

Proof. It is sufficient to show that $f(v)$ is a vertex if and only if v is a vertex. f is an isometry - in particular it is bijective - therefore $f(v)$ is a branch point if and only if v is a branch point. The only vertices which remain are the degree 2 vertices. Recall that these were introduced as the midpoints of inverted edges, hence they all have stabiliser of order 2. f is equivariant, therefore $\mathrm{stab}(v)$ is order 2 if and only if $\mathrm{stab}(f(v))$ is order 2. Thus the set of degree 2 vertices is also preserved, and f is a graph automorphism. □

Definition 6.13. Let Y be a metric graph. The *volume* of Y, denoted $\mathrm{Vol}(Y)$ is defined to be the sum of the lengths of the edges of Y.

Let $T \in \mathcal{O}$. The *covolume* of T, denoted $\mathrm{Covol}(T)$, is defined to be the volume of the graph T/G.

Proposition 6.14. *Let $f : T \to \alpha T$ be a topological representative for some $T \in \mathcal{O}(G, \mathcal{G})$, for some $\alpha \in Out(G, \mathcal{G})$. Then $Lip(f) = 1$ if and only if f is an isometry.*

Proof. If f is an isometry, then $\mathrm{Lip}(f) = 1$ follows immediately. It remains to prove the converse.

Let D be a subforest of T consisting of exactly one edge from each orbit. Then $\mathrm{Covol}(T) = \mathrm{Vol}(D)$. Without loss of generality, we may assume that $\mathrm{Covol}(T) = 1$. Since T and αT have the same metric, this means that $\mathrm{Covol}(\alpha T) = 1$

Since f is equivariant, $f(D)$ contains a fundamental domain for $f(T)$. Since αT does not contain any proper invariant subtrees, we must have $f(T) = \alpha T$, hence $f(D)$ contains a fundamental domain for αT. It follows that $\mathrm{Vol}(f(D)) \geqslant 1$. In addition,

$$\begin{aligned}
\mathrm{Vol}(f(D)) &\leqslant \sum_{\text{edges } e \in D} L(f(e)) && \text{(*)} \\
&\leqslant \sum_{\text{edges } e \in D} L(e) && \text{(as } \mathrm{Lip}(f) = 1\text{)} \\
&= \mathrm{Vol}(D) \\
&= \mathrm{Covol}(T) \\
&= 1.
\end{aligned}$$

We split into two cases:

Case 1: f is not locally injective This is equivalent to saying that f 'folds' a pair of edges - that is, there is a vertex v, neighbourhoods U_1, U_2 of v, and edges e_1, e_2 incident to v such that $f(e_1 \cap U_1) = f(e_2 \cap U_2)$. (The neighbourhoods are required because f may not fold the entirety of the edges, only the initial segments. Since f may stretch these segments, the neighbourhoods are not the same size in general).

The two edges can only be folded if they lie in different orbits; observe that if $e_1 \cdot g = e_2$, then $f(e_1 \cap U_1)$ must be fixed by $\alpha(g)$, contradicting edge freeness. Therefore we are free to choose D such that it contains a pair of folded edges.

If e_1, e_2 are a pair of folded edges in D, then the volume of their image under f will be strictly less then the sum of their original lengths. This means that (*) is a strict inequality, so $\mathrm{Vol}(f(D)) < 1$. This contradicts $\mathrm{Vol}(f(D)) \geqslant 1$, hence Case 1 cannot occur.

Case 2: f is locally injective. Then (*) is an equality, so $\mathrm{Vol}(f(D)) = 1$. This is enough to tell us that $L(f(e)) = L(e)$ for every edge e of D, and hence all the edges of T; thus f is an isometry on every edge of T. This, combined with local injectivity, means that f is an isometry on all of T.

□

7 Distance on $\mathcal{O}$

Stretching Factors

Definition 7.1. Let $g \in G$, and let $T \in \mathcal{O}(G, \mathcal{G})$. The *translation length* of g in T, denoted $l_T(g)$, is defined as

$$l_T(g) = \inf_{x \in T}\{d_T(x, x \cdot g)\}.$$

Remark 7.2. This infimum is in fact a minimum, and is obtained for some x. If g is elliptic then this is observed to be true from the definition of an elliptic element, and we have $l_T(g) = 0$.

If g is hyperbolic then the translation length will be non-zero, and the set of elements realising this length will form a line through T called the *hyperbolic axis of g*. Points on the axis will be translated along the axis by $l_T(g)$.

Remark 7.3. An equivalence class of G-trees in $\mathcal{O}$ is uniquely determined by its translation length function [13] - thus one can think of $\mathcal{O}$ as being embedded in the space $\mathbb{R}^G$.

Definition 7.4. Let $\mathrm{Hyp}(\mathcal{G})$ denote the set of elements of G which do not lie in any subgroup of $\mathcal{G}$. (In other words, $\mathrm{Hyp}(\mathcal{G})$ is the set of elements which are hyperbolic with respect to some, and hence all, G-trees in $\mathcal{O}(G,\mathcal{G})$).

Definition 7.5. [18, p.8] Let $T, S \in \mathcal{O}(G,\mathcal{G})$. Then we define the *left and right stretching factor* from T to S as

$$\Lambda_L(T,S) := \sup_{g\in \mathrm{Hyp}(\mathcal{G})} \frac{l_T(g)}{l_S(g)}, \qquad \Lambda_R(T,S) := \sup_{g\in \mathrm{Hyp}(\mathcal{G})} \frac{l_S(g)}{l_T(g)} = \Lambda_L(S,T),$$

respectively. We also define the *symmetric stretching factor* from T to S to be

$$\Lambda(T,S) := \Lambda_L(T,S)\Lambda_R(T,S).$$

The next Theorem follows from [18, Corollary 6.8, p.18 and Theorem 6.11, p.19].

Theorem 7.6. *Let $T, S \in \mathcal{O}$. Then there exists a Lipschitz continuous map $f : T \to S$ such that $Lip(f) = \Lambda_R(T,S)$.*

The Displacement of an Automorphism

Definition 7.7. Let $\alpha \in \mathrm{Out}(G,\mathcal{G})$. Then we define the *displacement* of α to be

$$l_\alpha := \inf_{T\in\mathcal{O}} \Lambda_R(T,\alpha T).$$

Theorem 7.8. *[18, p. 25] For any $\mathcal{G}$-irreducible $\alpha \in Out(G,\mathcal{G})$, the displacement of α is a minimum and obtained for some $T \in \mathcal{O}$.*

Definition 7.9. For any $\alpha \in \mathrm{Out}(G,\mathcal{G})$, we define

$$\mathrm{Min}(\alpha) = \{T \in \mathcal{O} \mid \Lambda_R(T,\alpha T) = \lambda_\alpha\}.$$

That is to say, $\mathrm{Min}(\alpha)$ is the set of all T which realise the above infimum.

Theorem 7.10. *Let $\alpha \in Out(G,\mathcal{G})$ be a $\mathcal{G}$-irreducible, displacement 1 automorphism. Then for all $T \in Min(\alpha)$, there exists an isometric topological representative for α on T.*

Proof. Let $T \in \mathrm{Min}(\alpha)$. Then by definition of the minimally displaced set, $\Lambda_R(T,\alpha T) = \lambda_\alpha = 1$, and hence by Theorem 7.6 there exists a Lipschitz continuous map $f : T \to \alpha T$ with $\mathrm{Lip}(f) = 1$. Therefore, by Proposition 6.14, f is an isometric topological representative for α. □

Corollary 7.11. *Let $\alpha \in Out(G,\mathcal{G})$ be a $\mathcal{G}$-irreducible, displacement 1 automorphism. Then $Min(\alpha) = Fix(\alpha)$.*

Proof. Let $T \in \mathrm{Fix}(\alpha)$. Then T and αT are equivalent G-trees, so $\Lambda(T, \alpha T) = 1 = \lambda_\alpha$. Thus $\mathrm{Fix}(\alpha) \subseteq \mathrm{Min}(\alpha)$.

Conversely, let $T \in \mathrm{Min}(\alpha)$. By Theorem 7.10, there exists an equivariant isometry from T to αT. Points in $\mathcal{O}$ are equivalence classes of G-trees under equivariant isometry, hence T and αT represent the same point in $\mathcal{O}$. Thus $\mathrm{Min}(\alpha) \subseteq \mathrm{Fix}(\alpha)$. □

8 Summary of paper 1

Background

The first paper in this thesis can be thought of as a generalisation of [15].

Definition 8.1. Let α be an automorphism of a finitely generated group G. Let E be a finite generating set of G, and let l_E denote the conjugacy length in the alphabet E. Then for $g \in G$ we define the *growth rate of α with respect to (the conjugacy class of) g* as

$$\mathrm{GR}(\alpha, l_E, g) = \limsup_{k \to \infty} \sqrt[k]{l_E(g\alpha^k)}$$

The *growth rate of α* is then

$$\mathrm{GR}(\alpha, l_E) = \sup\{\mathrm{GR}(\alpha, l_E, g) \mid g \in G\}$$

Remark 8.2. It can be shown that any two finite generating sets of a finitely generated group G give Lipschitz equivalent conjugacy length functions, and from this it can be further shown that Definition 8.1 does not depend on the choice of E.

In [16], Dicks & Ventura classify the irreducible growth rate 1 automorphisms of free groups. For prime numbers $p < q$, the free group of rank $(p-1)(q-1)$ can be presented as

$$\langle x_{i,j} (i \in \mathbb{Z}_p, j \in \mathbb{Z}_q) \mid x_{i,0} = x_{j,0} = 1 (i \in \mathbb{Z}_p, j \in \mathbb{Z}_q) \rangle$$

Using this presentation, the authors show that the following automorphisms of free groups are all irreducible and growth rate 1:

- $\alpha_{pq} \in \mathrm{Aut}(F_{(p-1)(q-1)})$ defined by $x_{i,j} \mapsto x_{1,1} x_{i+1,1} x_{i+1,j+1} x_{1,j+1}$,
- $\alpha_q := \alpha_{2q}^2 \in \mathrm{Aut}(F_{q-1})$ (so taking $p = 2$) for q an arbitrary odd prime,
- α_0 and α_1, denoting the identity automorphisms of the free groups of ranks 0 and 1 respectively,

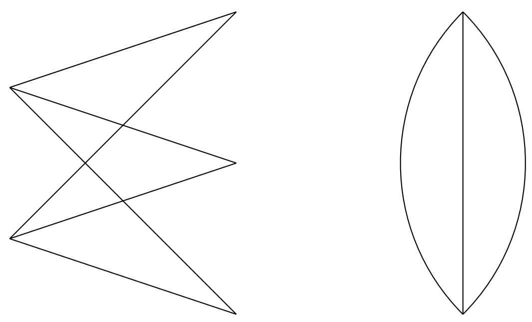

FIGURE 3: Y_{23} and Y_2

- α_2, denoting the inverting automorphism of the free group of rank 1.

They further prove that, up to outer automorphism class, these are the *only* irreducible, growth rate 1 automorphisms of free groups. They prove this by topologically representing these automorphisms as the action of the free group F on a graph whose fundamental group is equal to F. They show that the properties of irreducibility and growth rate limit the shape the graph can take and how the free group acts upon it [16, Prop 3.4] - thus they can find a graph and topological representative for every irreducible growth rate 1 automorphism.

- For $p < q \in \mathbb{Z}$, let Y_{pq} denote the complete bipartite graph on vertex sets $\mathbb{Z}_p$ and $\mathbb{Z}_q$. For $(i, j) \in \mathbb{Z}_p \times \mathbb{Z}_q$, let $e_{i,j}$ denote the corresponding edge. Let β_{pq} be the graph automorphism which sends each $e_{i,j}$ to $e_{i+1,j+1}$.

  Y_{pq} has fundamental group $F_{(p-1)(q-1)}$, and when $p < q$ are primes, β_{pq} represents the automorphism α_{pq}.

- For $q \in \mathbb{Z}$, let Y_q denote the graph with vertex set $\{0, 1\}$ and with p edges from 0 to 1. Let β_q be the graph automorphism which fixes both vertices and cyclically permutes the edges.

  Y_q has fundamental group F_{q-1}, and β_q represents the automorphism α_q for $q = 0, 1, 2$ and all odd primes.

Example 8.3. As an example, the free group F_2 has two irreducible growth rate 1 automorphisms, α_{23} and $\alpha_2 := \alpha_{23}^2$, which can be represented by cyclically permuting the edges of the graphs depicted in Figure 3

However, there is more to explore when taking this topological approach. In particular, we are interested in this corollary to Dicks' & Ventura's classification:

Corollary 8.4. *An irreducible, growth rate 1 automorphism of a free group fixes exactly one point in Culler-Vogtmann space.*

This corollary is not stated in Dicks' & Ventura's paper, but it can be proved fairly swiftly thanks to the existing library of work on Culler-Vogtmann space.

Proof. The maps described by Dicks & Ventura in [16] are graph automorphisms - hence they are also train track maps (more detail on train track maps is given in the chapter on Paper 2).

The minimally displaced set in CV_n of an irreducible automorphism coincides exactly with the set of points which support train track maps [18, p.32, Thm 8.19]. Additionally, it can be shown that the growth rate of an irreducible automorphism is equal to its displacement, and if the automorphism has displacement 1, then the minimally displaced set is actually a fixed point set. It follows that the graphs supporting the train track maps found by Dicks & Ventura are fixed points in Culler-Vogtmann space.

The final step is to show the uniqueness of these points, which is done in the proof of [21, p.10, Theorem 3.8]. □

The goal of paper 1 is to generalise this corollary to free products.

Dicks & Ventura classifies the fixed points of irreducible, growth rate 1 automorphisms in Culler-Vogtmann space, so the natural generalisation would be a classification of the fixed points of irreducible, growth rate 1 automorphisms of *free products* in *outer space*. However, there is one minor issue which we should address. For our proofs in Culler-Vogtmann space, we actually use the displacement of automorphisms, not their growth rate. For irreducible automorphisms of free groups, these two values are equal, and this can be proved fairly swiftly thanks to the existing body of work on Culler-Vogtmann space. In the more general case though, the proof is not quite so simple, and we do not give it in Paper 1. Thus what we *actually* do in Paper 1 is classify the fixed points of irreducible, *displacement* 1 automorphisms of free products.

Ultimately, however, this distinction does not matter, because patching this hole in the literature became my motive for Paper 2. In that second paper we do indeed prove that the growth rate of irreducible automorphisms of free products is equal to their displacement in Outer Space. Thus we reach a true generalisation of Corollary 8.4:

Theorem 8.5. *An irreducible, displacement 1 (or equivalently growth rate 1) automorphism of a free product fixes exactly one point in the corresponding outer space.*

Remark 8.6. The converse is not necessarily true. See the final section of this introduction.

Method

As for how we generalise this result in Paper 1: Let G be a free product with corresponding free factor system $\mathcal{G}$. One can think of $\mathcal{O}_1$ as a union of open simplices, where a G-tree's position in its simplex is determined by the lengths of its edges. We show that the action of α on each $T \in \mathrm{Min}_1(\alpha)$ can be topologically represented by an isometry of T, and that this isometry must cyclically permute the G-orbits of edges in T. It follows that the edges of T must all have the same length, and hence T must lie at the centre of its open simplex - thus the set of points in $\mathcal{O}_1$ fixed by α must consist solely of simplex centres. However, it is shown in [20, p.19, Cor 5.4] that $\mathrm{Min}_1(\alpha)$ is connected by so-called simplicial paths. Since a nontrivial simplicial path between the centres of two simplices must pass through a point which is not at the centre of a simplex, it follows that α cannot have more than one fixed point. It can also be shown that, for an irreducible automorphism, $\mathrm{Min}_1(\alpha)$ is non-empty [18, Theorem 8.4]. Thus $\mathrm{Min}_1(\alpha)$ must consist of exactly one point.

9 Summary of Paper 2

As mentioned in the previous chapter, Paper 1 can be thought of as a generalisation of a paper by Dicks & Ventura [16], but where they used growth rate, we used the displacement in outer space. This difference prompted the questions which eventually formed Paper 2:

(1) In [16], the growth rate was defined for free groups. What is the "correct" generalisation of this definition to free products?

(2) When we have found this generalisation, does the growth rate of an *irreducible* free product automorphism equal its displacement in outer space?

(3) More generally, does the growth rate of *any* free product automorphism equal its displacement in outer space?

(1)

First we address the definition of the growth rate. In the case of free groups, Dicks & Ventura [16] attribute the following definition of the growth rate in arbitrary groups to Thurston:

Definition 9.1. Let α be an automorphism of a finitely generated group G. Let E be a finite generating set of G, and let l_E denote the conjugacy length in the alphabet E. Then for $g \in G$ we define the *growth rate of α with respect to (the conjugacy class of) g* as

$$\mathrm{GR}(\alpha, l_E, g) = \limsup_{k\to\infty} \sqrt[k]{l_E(g\alpha^k)}$$

The *growth rate of α* is then

$$\mathrm{GR}(\alpha, l_E) = \sup\{\mathrm{GR}(\alpha, l_E, g) \mid g \in G\}$$

Remark 9.2. It can be shown that Lipschitz equivalent length functions will produce the same growth rate. In particular, any two finite generating sets of a finitely generated group G give Lipschitz equivalent conjugacy length functions - thus we may omit l_E from the notation, writing only $GR(\alpha)$.

Using this definition, it can be fairly swiftly proved that the growth rate of an irreducible automorphism of a free group is equal to its displacement in Culler-Vogtmann space. We will not give this proof just yet, since our later proof in the case of free products is essentially identical, only more general. When we say that we are searching for the "correct" definition of growth rate in free products, it is this proof we have in mind.

The displacement of an automorphism is determined entirely by the hyperbolic elements of G. When G is free this poses no issue, as all non-trivial elements will be hyperbolic. In the case of free products, however, the subgroups in the induced free factor system $\mathcal{G}$ are all elliptic. Our definition of the growth rate should somehow ignore any growth contributed by the elliptic elements. Thus we elect to use *relative generating sets*. (This method was inspired by Osin's paper on relatively hyperbolic groups [25]).

Definition 9.3. We say that $E \subseteq G$ is a *relative generating set of G with respect to $\mathcal{G}$* if G is generated by the set

$$\left(\bigcup_{i=1}^{k} G_i\right) \cup E,$$

where $G = G_1 * \ldots * G_k * F_r$ is a free product decomposition corresponding to $\mathcal{G}$.

From this we can essentially follow the definition of the growth rate in free groups by defining *relative conjugacy length*, *relative Lipschitz equivalence*, and finally the *relative growth rate*. We give a brief overview of these three:

The *relative conjugacy length* is the conjugacy length in the alphabet $E \cup G_1 \cup \ldots \cup G_k$. This ensures that elliptic elements all have length 1.

In the free group case, Lipschitz equivalent length functions produce the same growth rate. We need to generalise this to free products - thus we define relative Lipschitz equivalence: Two length functions l_1, l_2 on the conjugacy classes of G are *Lipschitz equivalent relative to $\mathcal{G}$* if they are Lipschitz equivalent when restricted to $\mathrm{Hyp}(\mathcal{G})$. We write $l_1 \sim_{\mathcal{G}} l_2$.

Definition 9.4. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, and let l_E be a relative conjugacy length function. Then for $g \in G$ we define the *relative growth rate of α with respect to (the conjugacy class of) g* as

$$\mathrm{GR}_{\mathcal{G}}(\alpha, l_E, g) = \limsup_{k\to\infty} \sqrt[k]{l_E(g\alpha^k)}$$

The *relative growth rate of* α is then

$$\mathrm{GR}_{\mathcal{G}}(\alpha, l_E) = \sup\{\mathrm{GR}(\alpha, l_E, g) \mid g \in G\}.$$

Remark 9.5. If we consider the case where G is a free group and $\mathcal{G}$ is the trivial free factor system (that is to say, we take the free product decomposition of G consisting of a single free factor - G itself), then all of these "relative" definitions restrict to their original counterparts, as they should.

(2)

It turns out that (2) is true. We prove this using *train track maps*. These are maps $f : T \to \alpha T$ which do not "fold" any edge back onto itself under any iterate f^k. More specifically...

Definition 9.6. Let $T \in \mathcal{O}$, and let v be a vertex in T. A *turn* at v is a pair of directed edges (e_1, e_2) such that $\tau(e_1) = v = \iota(e_2)$. We say the turn is *degenerate* if $e_2 = \overline{e_1}$.

A simplicial topological representative $f : T \to \alpha T$ induces a map Df on the turns of T. $Df(e_1, e_2)$ is the turn consisting of the first edges in the edge paths $f(e_1), f(e_2)$. A turn is *illegal* with respect to f if its image under some iterate of Df is degenerate. Otherwise, it is *legal*.

We say an edge path γ in T is *legal* if it does not contain any illegal turns.

Definition 9.7. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, $T \in \mathcal{O}(G, \mathcal{G})$. We say that a topological representative $f : T \to \alpha T$ is a *train track map* if for every edge $e \in T$ the path $f(e)$ is legal with respect to f.

Since the definition of a legal turn considers images under *iterations* of Df, the defining property of train track maps extends naturally to the iterates f^k:

Lemma 9.8. *Let $f : T \to \alpha T$ be a train track map. Then for all edges e in T and for all $k > 0$, $f^k(e)$ is a legal path.*

Proof. We prove this by induction on k. Recall that a path is legal if it does not contain any illegal turns.

$k = 0 : f^0(e) = e$ is a single edge. It does not contain any turns - in particular illegal turns - hence it is a legal path.

Now suppose that $f^{n-1}(e)$ is legal for $n \geqslant 1$. We want to show that every turn in $f^n(e)$ is legal. $f^{n-1}(e)$ is an edge path, so we have edges $\varepsilon_1, \ldots, \varepsilon_r$ such that $f^{n-1}(e) = \varepsilon_1 \ldots \varepsilon_r$. Therefore we can write $f^n(e) = f(\varepsilon_1) \ldots f(\varepsilon_r)$. The turns contained inside each $f(\varepsilon_i)$ are legal by the definition of a train track map. The remaining turns are those where the subpaths $f(\varepsilon_i)$ and $f(\varepsilon_{i+1})$ meet for each i - that is to say, the turns $Df(\overline{e_i}, e_{i+1})$. Each turn $(\overline{e_i}, e_{i+1})$ appears in $f^{n-1}(e)$, so each turn $Df(\overline{e_i}, e_{i+1})$ is legal by induction. □

Some further properties of train tracks can be seen by examining the corresponding transition matrix:

Definition 9.9. Topological representatives (such as train tracks) are simplicial, so they map edges to edge paths. Thus every topological representative $f : T \to \alpha T$ has an associated *transition matrix* $M = (m_{ij})$, where m_{ij} is the number of times the f-image of the j-th edge-orbit crosses the i-th edge-orbit in any direction.

A transition matrix is a non-negative, integer-valued square matrix, and in the specific case of irreducible automorphisms the transition matrix will also be irreducible; therefore, by the Perron-Frobenius theorem, one of its eigenvalues, called the *PF-eigenvalue*, is a positive real number μ which is greater than or equal to the absolute value of all other eigenvalues. There is a positive real eigenvector $\mathbf{v}$ corresponding to μ, which is called the *PF-eigenvector*.

The metric on this particular tree T has not been specified, so we are free to choose the edge lengths in the way that is most convenient for us. We do this using the PF-eigenvector, declaring the length of any representative of the ith edge to be the ith entry in this eigenvector. This ensures that f scales the length of every edge by exactly μ - that is to say, $\forall e \in T, l(f(e)) = \mu l(e)$.

Furthermore, iterations f^k of the map f correspond to iterations of the transition matrix M^k: the transition matrix of f^k is M^k, with PF-eingenvalue μ^k and PF-eigenvector $\mathbf{v}$. Thus $\forall e \in T, \forall k > 0$, we have $l(f^k(e)) = \mu^k l(e)$.

Lemma 9.10. *Let $f : T \to \alpha T$ be a train track map and take the metric on T to be the one determined by the PF-eigenvector. Then $\exists g \in Hyp(\mathcal{G})$ such that $l_T(g\alpha^k) = \mu^k l_T(g)$ for all $k > 0$.*

Proof. Take an edge e in T and consider the path $f^k(e)$. We observe that T, being co-compact, contains finitely many edge orbits, so $f^k(e)$ must cross the orbit of some edge ε at least three times. Furthermore, at least two of these edges will point in the same direction.

Let us label these edges ε and $\varepsilon \cdot g$. Then the path $f^k(e)$ will pass through the vertices $\iota\varepsilon$, $\tau\varepsilon$, $\iota\varepsilon \cdot g$ and $\tau\varepsilon \cdot g$ in this order. This can only occur when ε, and hence $\varepsilon \cdot g$, lie on the hyperbolic axis of g (see Figures 4 and 5). It follows that the path $f^k(e)$ contains a subpath which is a fundamental domain of the axis of g. As a part of the main path, this subpath will be rescaled by μ with each iteration of f, which concludes the proof.

□

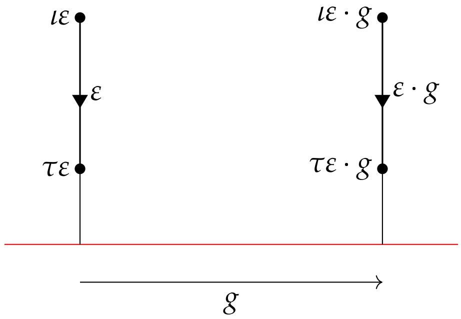

FIGURE 4: If an edge ε does not lie on the axis of g (in red), the edges ε and $\varepsilon \cdot g$ will have different orientations on a path passing through them

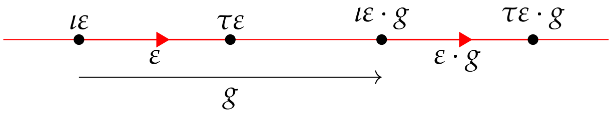

FIGURE 5: If an edge ε lies on the axis of g (in red), the edges ε and $\varepsilon \cdot g$ will have the same orientation on a path passing through them

We can now prove our answer to Question **(2)**:

Theorem 9.11. *Let $\mathcal{G}$ be a free factor system for a group G, let E be any relative generating set for G, and let $\alpha \in Out(G, \mathcal{G})$ be irreducible. Then $GR_{\mathcal{G}}(\alpha, l_E) = \lambda_\alpha$.*

Proof. $\mathrm{Min}(\alpha)$ is equal to the train-track bundle $\mathrm{TT}(\alpha)$ - the set of $T \in \mathcal{O}$ admitting train track representatives $f : T \to \alpha T$ with $\mathrm{Lip}(f) = \Lambda_R(T, \alpha T)$ [18, Thmm 8.19, Thm 6.11]. In addition, since α is irreducible, $\mathrm{Min}(\alpha)$ is non-empty [18, Theorem 8.4].

Thus we can guarantee the existence of a train track map $f : T \to \alpha T$ on some T such that:

$$\mu \overset{f \text{ is train track}}{=} \mathrm{Lip}(f) \overset{f \in \mathrm{TT}(\alpha)}{=} \Lambda_R(T, \alpha T) \overset{\text{by defn of } \mathrm{Min}(\alpha)}{=} \lambda_\alpha$$

In addition, since $l_E \sim_{\mathcal{G}} l_T$, these length functions will produce the same growth rate.

Therefore in order to prove that $\mathrm{GR}_{\mathcal{G}}(\alpha, l_E) = \lambda_\alpha$ it suffices to prove that $\mathrm{GR}_{\mathcal{G}}(\alpha, l_T) = \mu$. The proof of this follows from the definition of the right stretching factor Λ_R. Recall,

$$\Lambda_R(T, \alpha^k T) := \sup_{g \in \mathrm{Hyp}(\mathcal{G})} \frac{l_T(g\alpha^k)}{l_T(g)}$$

$\Rightarrow$ For all $g \in G$, $l_T(g\alpha^k) \leqslant \mu^k l_T(g)$.

$\Rightarrow$ For all $g \in G$, $GR_{\mathcal{G}}(\alpha, g, l_T) = \limsup_{k\to\infty} \sqrt[k]{l_T(g\alpha^k)} \leqslant \limsup_{k\to\infty} \sqrt[k]{\mu^k l_T(g)} = \mu$.

Additionally, by the train track property, there exists $h \in G$ such that $\mu^k l_T(h) = l_T(h\alpha^k)$

$\Rightarrow GR_{\mathcal{G}}(\alpha, h, l_T) = \limsup_{k\to\infty} \sqrt[k]{l_T(h\alpha^k)} = \limsup_{k\to\infty} \sqrt[k]{\mu^k l_T(h)} = \mu$

Thus $GR_{\mathcal{G}}(\alpha, l_T) := \sup_g GR_{\mathcal{G}}(\alpha, g, l_T) = \mu$, and we are done. □

(3)

If we drop the irreducibility condition, a problem arises which prevent us from copying the previous proof outright: we cannot guarantee that $\text{Min}(\alpha)$ will be non-empty, and hence we cannot guarantee the existence of an optimal train track map. Happily, however, we can guarantee the existence of a weaker set of maps known as *relative train tracks*.

Remark 9.12. Relative train track maps were introduced by Bestvina & Handel in the case of free groups [5], and were generalised to free products by Collins & Turner [9]. Collins and Turner defined their maps on graphs of complexes - graphs with 2-complexes assigned to the vertices. However, this definition can be transferred to G-trees by replacing each 2-complex with its fundamental group to give a graph of groups and then lifting to the Bass-Serre tree.

Let $f : T \mapsto \alpha T$ be any topological representative, and consider the associated transition matrix $M = (m_{ij})$. By relabelling edges appropriately, it is always possible to write the transition matrix in block upper triangular form:

$$M = \begin{pmatrix} M_1 & ? & ? & ? \\ 0 & M_2 & ? & ? \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & M_n \end{pmatrix},$$

where the matrices $M_1, \ldots, M_n$ are either zero matrices or irreducible matrices.

Writing M in this form determines a partition of the edges of T: The *rth stratum* H_r of T is the subgraph of T given by closure of the union of the edge orbits corresponding to the rows/columns in M_r.

This also determines a filtration $\emptyset = T_0 \subset \ldots \subset T_n = T$ of T, where $T_r = \bigcup_{i\leqslant r} H_r$. Observe that each T_i is f-invariant, but the H_i are not, in general.

Definition 9.13. We say an edge path γ in T_r is *r-legal* if no component of $\gamma \cap H_r$ contains an illegal turn.

Definition 9.14 (Relative train track)**.** Let $T \in \mathcal{O}(G, \mathcal{G})$, let $\alpha \in \text{Out}(G, \mathcal{G})$, and let $f : T \to \alpha T$ be a simplicial topological representative for α. Use this map to divide T into strata as described above. We say that f is a *relative train track map* if the following hold:

(1) *f preserves r-germs*: For every edge $e \in H_r$, the path $f(e)$ begins and ends with edges in H_r.

(2) *f is injective on r-connecting paths*: For each nontrivial path $\gamma \in T_{r-1}$ joining points in $H_r \cap T_{r-1}$, the homotopy class $[f(\gamma)]$ is nontrivial.

(3) *f is r-legal*: If a path γ is r-legal, then $f(\gamma)$ is r-legal.

Theorem 9.15. *For any automorphism $\alpha \in Out(G, \mathcal{G})$, there exists a relative train track map $f : T \to \alpha T$ on some $T \in \mathcal{O}$. [9, Thm 2.12]*

We observe that M_r is the transition matrix of H_r, and each of these submatrices will have its own PF-eigenvalue μ_r. It can be shown that, even though a relative train track map will not, in general, satisfy Lemma 9.10, it *will* satisfy a similar property on each stratum of T using these μ_r. This gives us the tools we require to prove that question (3) is true:

Theorem 9.16. *Let $\alpha \in Out(G, \mathcal{G})$. Then the following are equal:*

- *The relative growth rate of α, $GR_{\mathcal{G}}(\alpha)$.*
- *The largest PF-eigenvalue μ_R of any relative train track map $f : T \to \alpha T$, on any $T \in \mathcal{O}(G, \mathcal{G})$.*
- *The displacement λ_α of α in $\mathcal{O}$.*

We prove this by proving three inequalites:

A: $\mu_R \leqslant \mathrm{GR}_{\mathcal{G}}(\alpha)$

B: $\mathrm{GR}_{\mathcal{G}}(\alpha) \leqslant \lambda_\alpha$

C: $\lambda_\alpha \leqslant \mu_R$

A follows from existing properties of relative train tracks. **B** can be proved explicitly using the definition of the right stretching factor. **C** requires more thought.

Recall the definition of the displacement: $\lambda_\alpha := \inf_{S \in \mathcal{O}} \Lambda_R(S, \alpha S)$. Ideally we would prove inequality **C** by finding a G-tree in $\mathcal{O}$ whose right stretching factor is exactly μ_R. However, unless α is irreducible, this is not always possible. Thus we instead find a *sequence* of G-trees whose right stretching factors *tend* towards μ_R.

The lengths of edges in T are determined by the PF-eigenvectors, but these are only determined up to scalar multiplication, so we are free to rescale the edges in each stratum by a constant of our choosing. We choose to rescale each H_r by N^r. As N tends to infinity, we observe that the growth in the stratum with the largest PF-eigenvalue becomes greater than that of all other strata. From this the result follows.

10 Summary of Paper 3

One of the key insights of geometric group theory is that one can obtain information on a group by viewing it as a metric space, via the word metric on its Cayley graph. More generally if a group, G, acts isometrically on a metric space (X, d) one can elucidate properties of the group from this action. For instance, the class of hyperbolic groups is precisely the class of those groups admitting a proper, co-compact isometric action on some locally compact, geodesic δ-hyperbolic space X.

Given a (right) isometric action of G on (X, d), and a point p in X, one can define a G-invariant pseudo-metric - which we denote by d_p - on G via $d_p(g, h) := d(pg, ph)$, which is a metric precisely when the stabiliser of p is trivial. In fact, this metric on G can be encoded via the *based length function*.

Definition 10.1. Let G act isometrically on the metric space (X, d). Then the *based length function* of G based at some point, $p \in X$ is the function, $l_p : G \to \mathbb{R}$, given by:

$$l_p(g) := d(p, pg)$$

It is straightforward to see that one can recover the invariant (pseudo) metric from the based length function via $d_p(g, h) = l_p(gh^{-1})$.

Of course, in order to obtain properties of the group it is helpful to impose conditions on the space and the action, just as for hyperbolicity above. A key area where one can recover a great deal of information about G is when X is a tree.

The source of inspiration for this paper is a striking result of Chiswell, that one can axiomatise the based length functions arising from actions on trees - sometimes called Lyndon length functions, following results from [23] - and, from the axioms, always recover an isometric action. Specifically,

Theorem 10.2 ([6])**.** *Suppose $l : G \to \mathbb{R}_{\geqslant 0}$ satisfies the following axioms:*

A1′: $l(g) = 0$ *if* $g = 1$

A2 : $l(g^{-1}) = l(g)$

A3 : $c(g, h) \geqq 0$

H_0 *: For all* $g_1, g_2, g_3 \in G$,
$c(g_1, g_2) \geqslant m$, $c(g_2, g_3) \geqslant m$ *implies that* $c(g_1, g_3) \geqslant m$,

where

$$c(g, h) := \frac{1}{2}(l(g) + l(h) - l(gh^{-1})).$$

Then there exists an $\mathbb{R}$-tree, (X, d), admitting an isometric G-action and a point, $p \in X$, such that $l_p(g) = l(g)$. Moreover, if the images of l and c lie in $\mathbb{Z}$, then the tree will be simplicial.

Remark 10.3. As noted above a function $d : G \times G \to \mathbb{R}$ can be defined from l and, from this point of view, A1′ says that d vanishes on the diagonal, A2 says that it is symmetric and A3 says that it satisfies the triangle inequality.

The function $c(g, h)$ is then really the Gromov product and axiom H_0 should be thought of as a 0-hyperbolicity condition (see, for example, [1] for a discussion on hyperbolic groups, spaces and the Gromov product). Chiswell's Theorem can then be summarised as saying that a 0-hyperbolic Lyndon length function (that is, one satisfying $A1, A2, A3, H_0$) is always a based length function on a 0-hyperbolic space.

With this in mind, we are motivated to ask the following:

Questions.

- *Is there a generalisation of Chiswell's Theorem for isometric group actions on metric graphs?*
- *In particular, is there a generalisation of Chiswell's Theorem for isometric actions on δ-hyperbolic graphs?*

Remark 10.4. In the spirit of Chiswell's result, we will consider graphs whose edge lengths may not be integers. For instance, one could take the Cayley graph of a group with respect to some generating set, and then equivariantly assign positive real lengths to edges.

It turns out that these questions are somehow too broad in their scope. Given a (strictly positive) length function on G there is always a metric graph and a point p whose based length function is equal to this function: take the complete graph on G where the edge between g and h has length $l(hg^{-1})$. The based length function on this graph, with respect to the basepoint 1, is equal to l. However, this action is not particularly useful.

In order to rule out this kind of example we will add some restrictions.

Questions. *Let us suppose that G is finitely generated and let us restrict ourselves to isometric, co-compact actions on locally compact graphs, X.*

- *Given a (strictly positive) length function, l, does G admit an isometric, co-compact action on a locally compact metric graph, X, such that $l = l_p$ for some $p \in X$?*
- *What if we add the hypothesis that l is δ-hyperbolic*

It turns out that the answer to both of these questions is **no**; there exists a δ-hyperbolic length function which cannot arise as the based length function associated to any isometric, co-compact action on a locally compact graph at any base point p.

However, that example is bi-Lipschitz equivalent to a length function on a Cayley graph. (Note that, for finitely generated groups, all based length functions on Cayley graphs with respect to finite generating sets are bi-Lipschitz equivalent). But we also produce examples of δ-hyperbolic length functions which are not bi-Lipschitz equivalent to any based length function on a Cayley graph. In fact, *every* finitely generated group admits a hyperbolic length function.

Theorem 10.5. *There exists a finitely generated group, G, with a hyperbolic length function, $l : G \to \mathbb{R}_{\geqslant 0}$ such that $l \neq l_p$ for any co-compact, metric G-graph and any point $p \in T$.*

Moreover, any *finitely generated group admits a (free) hyperbolic length function. In particular, we can find an example of a group G with a hyperbolic length function, l, which is not quasi-isometric to any based length function arising from an isometric action of G on a geodesic and proper δ-hyperbolic metric space.*

This leads us to the following.

Questions. *Suppose that G is finitely generated.*

- *Can one axiomatise those length functions which are bi-Lipschitz equivalent to some (and hence all) based length functions on a Cayley graph for G (with respect to a finite generating set)?*
- *Can we make these axioms apply to - for instance - any free F_n action on a simplicial tree as well as Cayley graphs?*
- *Does this axiomatisation define a connected/contractible/finite dimensional subspace of $\mathbb{R}^G$ on which $\mathrm{Aut}(G)$ acts?*

Remark 10.6. We do come up with an axiom scheme, below, and we observe that these axioms hold for all sufficiently well behaved actions - and in particular to all points of Culler-Vogtmann space.

The third question here arises from the fact that one key use of Chiswell's Theorem is in the study of group actions on trees, and the definition of the space of such actions which are then encoded via functions (usually the translation length function, which is related to the Lyndon length function). See [14] for the seminal paper on the 'Outer Space' of free actions on trees, encoded by length functions (amongst other things).

It is clear that the space of all length functions which are bi-Lipschitz to one arising from a Cayley graph is a contractible space (because a linear combination of such functions is another such function). Therefore, this provides a contractible space on which $\mathrm{Aut}(G)$ acts. However, it is far too large and so one might hope that an axiomatisation could provide a more reasonable subspace.

With these questions in mind, we propose the following axioms for our length functions:

Definition 10.7. Let G be a group. We say that $l : G \to \mathbb{R}_{\geqslant 0}$ is a *graph-like length function* if it satisfies the following axioms:

A1: $l(g) = 0$ if and only if $g = 1$

A2: $l(g^{-1}) = l(g)$

A3: $c(g, h) \geqq 0$

A4: For all $R \geqslant 0$, the closed ball $B_R := \{g \in G \mid l(g) \leqslant R\}$ is finite

A5: There exists $0 \leqslant \epsilon < 1$ and $K > 0$ such that, for any $g \in G$, if $l(g) > K$ then there exists an $x \in G$ with:

  (i) $0 < l(x) \leqslant K$, and
  (ii) $c(gx^{-1}, x^{-1}) \leqslant \frac{\epsilon l(x)}{2}$.

Remark 10.8. We have added two additional axioms, A4 and A5, to the definition of a length function. The purpose of these new axioms is to describe the action of a group on a “sensible” graph purely in terms of its based length function.

A4 is really a statement about the action being properly discontinuous. Furthermore, in the presence of A5, A4 is equivalent to the statement that B_K is finite.

One should view A5 as a co-compactness condition. Specifically, if one has a reasonable action on a graph, then one can approximate geodesics in the graph with uniform quasi-geodesics built from the translates of finitely many paths. The following example illustrates this behaviour. In Paper 3 itself, we prove that A5 holds for the based length function of every isometric, co-bounded action on a based geodesic metric space.

We also note that if G acts on its Cayley graph then one easily gets that the based length function satisfies these axioms with $K = 1$ and $\epsilon = 0$. However, if once considers actions on graphs with more than one orbit of vertices, then one quickly discovers that the correct condition is A5(ii) with $\epsilon \neq 0$. Moreover, scaling the graph by a constant clearly changes the value of K. For these reasons, to allow these kinds of deformations, we consider these axioms for more general K and ϵ.

Example 10.9. We will use the free group $F_2 =< a, b >$ as an example. Consider first the Cayley graph on this generating set. This is a tree acted on by F_2 - specifically, it is the universal cover $\tilde{X}$ of a graph X with one vertex v and two edges a, b:

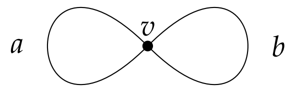

Let every edge have length 1, and let the base point p be the vertex corresponding to the identity. Then for all $g \in F_2$ the based length function $l_p(g)$ is equal to the word length $||g||$. We can easily prove that $||g||$ satisfies A4 and A5 with $K = 1$ and $\epsilon = 0$: A4 holds because F_2 is finitely generated, and A5 holds when we take x to be the last letter in the reduced word representing g.

Our focus, however, is on the behaviour of the graph $\tilde{X}$. Consider the word $ab^2 \in F_2$. This word corresponds to a geodesic path $[p, p \cdot ab^2]$ in the Cayley graph $\tilde{X}$, depicted below.

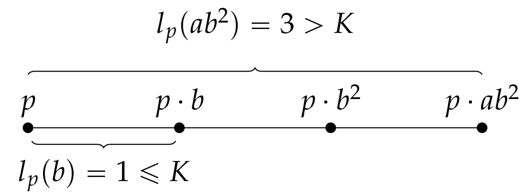

We are taking $\varepsilon = 0$, so axiom A5 is essentially stating that, given a based length longer than $K = 1$, there exists a group element whose based length is no greater than $K = 1$ and which corresponds to a point *on the geodesic*. The final part holds in this graph (and indeed in all Cayley graphs) because there is only one orbit of vertices; every vertex on the geodesic corresponds to a group element.

If we consider more general graphs with more than one orbit of vertices, it becomes necessary to take $\varepsilon > 0$. Consider a second graph Y with two vertices u, v and three edges: an edge-loop E_u at u, an edge loop E_v at v, and an edge E_{uv} from u to v. Just like X, Y has fundamental group $F_2 =< a, b >$: we will assign the generators a and b to the loops E_u and $E_{uv}E_vE_{uv}^{-1}$ respectively. To ensure that each generator has length 1, we set the length of E_u to be 1, and the lengths of E_v and E_{uv} to be $\frac{1}{3}$.

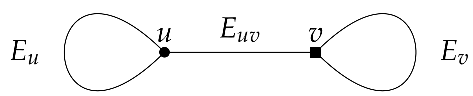

The universal cover $\tilde{Y}$ is a tree acted upon by F_2 with two orbits of vertices. Take the base point q of Y to be a lift of u, and consider the word $b^3 \in F_2$. In $\tilde{Y}$ the geodesic $[q, q \cdot b^3]$ in $\tilde{Y}$ (highlighted in red) only meets the orbit of the base point at its endpoints:

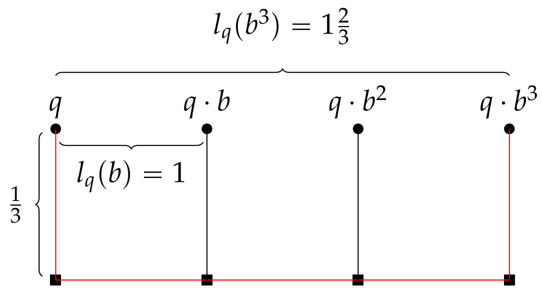

Thus, unlike the previous example, when $K = 1$ there are no $1 \neq g \in F_2$ with $l_q(g) \leqslant K$ which correspond to points on the geodesic. If we take a larger value of K, then we can have the endpoint of the geodesic satsifying $l_q(b^3) \leqslant K$, but upon taking an arbitrarily large product b^n the problem will repeat.

However, there is a sequence of points, starting with $q \cdot b$, which lies *uniformly close* to the geodesic. Every word in F_2 must start with either an a or a b, so from here it is not hard to see that the function l_q will satisfy axiom A5 for $K = 1$ and $\varepsilon = \frac{1}{3}$.

It turns out that axioms A1 to A5 are indeed sufficient to prove the following:

Theorem 10.10. *Let $l : G \to \mathbb{R}_{\geqslant 0}$ be a graph-like length function on a group G. Then l is bi-Lipschitz equivalent to some (and hence to all) based length function l_p arising from a locally compact, co-compact, metric G-graph and with $Stab(p) = 1$.*

Note that in view of Theorem 10.5, since any finitely generated group admits a hyperbolic length function, the extra axioms are clearly necessary.

Remark 10.11. We should note that another length function one can extract from an action is the *translation length function*, which has the advantage of not relying on a basepoint. This is the point of view of [14]. An important result here, building on the work of [14], is that of [26] which states that a translation length function (which is 0-hyperbolic) always arises from an action on a tree. However, this builds crucially on Chiswell's Theorem 10.2 so it seems reasonable to start with Lyndon length functions.

11 Open problems

To conclude the introduction, we remark upon some open problems arising from these papers which could be avenues for future research.

Paper 1

As mentioned above, Paper 1 is a partial generalization of a paper by Dicks & Ventura [15]. In Paper 1 we prove that if α is an irreducible, displacement 1 automorphism of a free product, then any topological representative $f : T \to \alpha T$ on a tree T in $\mathcal{O}$ must cyclically permute the edge orbits of T. We then use this to prove that there is only one fixed point in $\mathcal{O}$, but we do not give any additional descriptions of T or α. Dicks & Ventura, on the other hand, do.

Dicks & Ventura work with irreducible, growth rate 1 automorphisms of free groups. They are not explicitly working in outer space or Culler-Vogtmann space, and they prefer to work at the level of graphs instead of trees, but like us they prove that their group automorphisms must be represented by some graph automorphism $\beta : X \to X$, and that this graph automorphism must cyclically permute the edges of X. However, unlike us, they also prove some additional properties which are required of β and X [15, Prop 3.6]. This allows them to give an exhaustive list of all β and X which satisfy their conditions and from there prove the converse: prove that every $\beta : X \to X$ on this list is indeed representing an irreducible growth rate 1 automorphism. Thus they arrive at a true classification of irreducible growth rate 1 automorphisms of free groups.

We have not generalized all of Dicks & Ventura's conditions to outer space: just the one concerning cyclically permuted edge orbits, which was enough to tell us that our fixed points lie at the centre of simplices in $\mathcal{O}$. That converse statement - that every automorphism which fixes a simplex centre is irreducible, growth rate 1 - is not necessarily true. To rectify this and truly classify our irreducible, displacement 1 automorphisms, we would need to prove additional properties of the fixed points in $\mathcal{O}$, but there are a few obstacles.

Firstly, we are working with free products of the form $G = G_1 * \ldots G_k * F_r$. The only conditions we have imposed on these free factors is their number, k, and how they are permuted (used in the definition of irreducibility). Therefore any reasonable classification we performed would ultimately have to be done in terms of k and r, ignoring the effect the automorphisms had inside each factor.

Secondly, when generalized to free products, the conditions Dicks & Ventura used may not be restrictive enough. For example, they are able to discount all graphs with a degree 1 vertex, but we cannot discount all graphs of groups with this property, since we must consider the possibility that the vertex has a non-trivial vertex group.

Paper 3

In Paper 3 we propose the axioms for a *graph-like length function*. These are elements of $\mathbb{R}^G$, the space of functions from G to $\mathbb{R}$, and we prove that they hold for a reasonable collection of based length functions of metric spaces: A4 is a cocompactness condition, and A5 holds for the based length function of every isometric, co-bounded action on a based geodesic metric space. Taking this further, we can think of the space of graph-like length functions as a subspace of $\mathbb{R}^G$. This raises questions which we asked in the introduction to Paper 3, but which we ultimately did not answer:

What are the properties of this subspace? For example, is it a connected/contractible/finite dimensional subspace of $\mathbb{R}^G$ on which $\mathrm{Aut}(G)$ acts?

References

[1] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short. Notes on word hyperbolic groups. In *Group theory from a geometrical viewpoint. Proceedings of a workshop, held at the International Centre for Theoretical Physics in Trieste, Italy, 26 March to 6 April 1990*, pages 3–63. Singapore: World Scientific.

[2] Naomi Andrew. Serre's property (FA) for automorphism groups of free products. *Journal of Group Theory*, 24(2):385–414, 2021.

[3] Hyman Bass. Covering theory for graphs of groups. 89(1-2):3–47.

[4] Mladen Bestvina, Mark Feighn, and Michael Handel. The Tits alternative for $\mathrm{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms. *Ann. of Math. (2)*, 151(2):517–623, 2000.

[5] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *The Annals of Mathematics*, 135(1):1–51, jan 1992.

[6] I. M. Chiswell. Abstract length functions in groups. 80:451–463.

[7] Ian Chiswell. *Introduction to Λ-trees*. Singapore: World Scientific.

[8] Daniel E. Cohen. *Combinatorial group theory: a topological approach*, volume 14 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1989.

[9] D. J. Collins and E. C. Turner. Efficient representatives for automorphisms of free products. 41(3):443–464.

[10] Matthew Collins. Fixed points of irreducible, displacement one automorphisms of free products. May 2023. Preprint available at https://arxiv.org/abs/2305.01451.

[11] Matthew Collins. Growth and displacement of free product automorphisms. July 2023. Preprint available at https://arxiv.org/abs/2307.13502.

[12] Matthew Collins and Armando Martino. Length functions on groups and actions on graphs. July 2023. Preprint available at https://arxiv.org/abs/2307.10760.

[13] Marc Culler and John W. Morgan. Group actions on **R**-trees. *Proceedings of the London Mathematical Society*, s3-55(3):571–604, nov 1987.

[14] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones Mathematicae*, 84(1):91–119, feb 1986.

[15] Warren Dicks and M. J. Dunwoody. *Groups acting on graphs*, volume 17 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1989.

[16] Warren Dicks and Enric Ventura. Irreducible automorphisms of growth rate one. *Journal of Pure and Applied Algebra*, 88(1-3):51–62, 1993.

[17] Max Forester. Deformation and rigidity of simplicial group actions on trees. *Geometry and Topology*, 6:219–267, 2002.

[18] Stefano Francaviglia and Armando Martino. Stretching factors, metrics and train tracks for free products. *Illinois Journal of Mathematics*, 59(4):859–899, 2015.

[19] Stefano Francaviglia and Armando Martino. Displacements of automorphisms of free groups I: Displacement functions, minpoints and train tracks. *Transactions of the American Mathematical Society*, 374(5):3215–3264, 2021.

[20] Stefano Francaviglia and Armando Martino. Displacements of automorphisms of free groups II: Connectivity of level sets and decision problems. *Transactions of the American Mathematical Society*, 375(4):2511–2551, 2022.

[21] Stefano Francaviglia, Armando Martino, and Dionysios Syrigos. The minimally displaced set of an irreducible automorphism of F_N is co-compact. *Archiv der Mathematik*, 116(4):369–383, 2021.

[22] I. Gruschko. Über die Basen eines freien Produktes von Gruppen. *Rec. Math. [Mat. Sbornik] N.S.*, 8 (50):169–182, 1940.

[23] R. C. Lyndon. Length functions in groups. 12:209–234.

[24] B. H. Neumann. On the number of generators of a free product. *J. London Math. Soc.*, 18:12–20, 1943.

[25] Denis V. Osin. Relatively hyperbolic groups: Intrinsic geometry, algebraic properties, and algorithmic problems. 179(843):vi+100.

[26] Walter Parry. Axioms for translation length functions. In *Arboreal group theory. Proceedings of a workshop, held September 13-16, 1988, in Berkeley, CA (USA)*, pages 295–330. New York etc.: Springer-Verlag.

[27] John Stallings. Coherence of 3-manifold fundamental groups. In *Séminaire Bourbaki, Vol. 1975/76, 28 ème année, Exp. No. 481*, pages 167–173. Lecture Notes in Math., Vol. 567. 1977.

Paper 1: **Fixed points of irreducible, displacement one automorphisms of free products**

Matthew Collins

ABSTRACT. We consider the action of outer automorphisms on the deformation space $\mathcal{O}$ of G-trees given by a free product decomposition of a group G. We show that an irreducible, displacement 1 automorphism fixes exactly one point in $\mathcal{O}_1$ (the covolume 1 slice of $\mathcal{O}$).

1 Introduction

This paper can be thought of as a generalisation of a paper by Dicks & Ventura [6], in which the authors classify the irreducible, growth rate 1 automorphisms of free groups F_n. In the process of doing so, they show that each of these automorphisms can be represented by a graph automorphism of a graph with fundamental group F_n. When combined with the results of [2] and [8], this means that an irreducible, growth rate 1 automorphism of a free group F_n fixes a single point in Culler-Voghtmann space CV_n. The main result of this paper is a generalisation of this result: we prove that an irreducible, growth rate 1 automorphism of a free product $G = G_1 * \ldots * G_k * F_r$ fixes a single point in the deformation space $\mathcal{O}_1$.

The free group version of this result is stated explicitly in [11, p.10, Thm 3.8], and it follows from Dicks & Ventura's classification like so: Every irreducible outer automorphsim of F_n is topogically represented by an irreducible train track map f on a graph in Culler-Vogtmann space CV_n [2, p.9, Thm 1.7]. If the automorphism is growth rate 1, then f is a finite order homeomorphism - in this case, a graph automorphism. Thus the graph automorphisms found by Dicks & Ventura in [6] are in fact train track maps.

The minimally displaced set in CV_n of an irreducible automorphism coincides exactly with the set of points which support train track maps [8, p.32, Thm 8.19]. Additionally, it can be shown that the growth rate of an irreducible automorphism is equal to its displacement, and if the automorphism has displacement 1, then the minimally displaced set is actually a fixed point set. It follows that the graphs supporting the train track maps found by Dicks & Ventura are fixed points in Culler-Voghtmann space.

The final step is to show the uniqueness of these points, which is done in the proof of [11, p.10, Theorem 3.8].

Our generalisation to free products follows a similar outline - however, we instead use the deformation space $\mathcal{O}$, otherwise known as *outer space*, which is a generalisation of

Culler-Vogtmann space to free products $G = G_1 * \ldots * G_k * F_r$. The notion of a deformation space was first introduced by Forester [7], and they have since been studied in [5] and [4]. Given a group G, one considers minimal, cocompact, isometric actions of G on metric simplicial trees. These trees, together with their actions, are called G-trees. Two G-trees are said to be equivalent if there exists an equivariant isometry between them, and one defines $\mathcal{O}$ to be the space of equivalence classes of G-trees which share the same set of elliptic subgroups - that is, subgroups which fix a point in the tree.

The group of outer automorphisms which preserve the set of conjugacy classes $\{[G_1], \ldots, [G_k]\}$ acts on $\mathcal{O}$ by "twisting" the actions of the G-trees. This group is denoted $\mathrm{Out}(G, \mathcal{G})$, and we study its action on the covolume one slice of $\mathcal{O}$ (denoted $\mathcal{O}_1$) by using the asymmetric Lipschitz metric: For any two G-trees $T, S \in \mathcal{O}_1$, we write $\Lambda_R(T, S)$ to denote the asymmetric Lipschitz distance, or stretching factor, between them.

For an automorphism $\alpha \in \mathrm{Out}(G, \mathcal{G})$, one can define the displacement of α as $\lambda_\alpha = \inf\{\Lambda(T, \alpha T) \mid T \in \mathcal{O}_1\}$. The *minimally displaced set* of α, $\mathrm{Min}_1(\alpha)$, is the set of G-trees T in $\mathcal{O}_1$ which realise this infimum. It can shown that if α is irreducible, then the displacement λ_α is not just an infimum, but it is a minimum, and hence the set $\mathrm{Min}_1(\alpha)$ is non-empty. In addition, it can be shown that $\Lambda_R(T, S) = 1$ if and only if T and S represent the same equivalence class in $\mathcal{O}_1$ - hence $\lambda_\alpha = 1$ implies that $\mathrm{Min}_1(\alpha)$ is the fixed-point set of α.

One can think of $\mathcal{O}_1$ as a union of open simplices, where a G-tree's position in its simplex is determined by the lengths of its edges. In Theorem 7.6, we show that the action of α on each $T \in \mathrm{Min}_1(\alpha)$ can be toplogically represented by an isometry of T, and that this isometry must cyclically permute the G-orbits of edges in T. It follows that the edges of T must all have the same length, and hence T must lie at the centre of its open simplex - thus $\mathrm{Min}_1(\alpha)$ must consist solely of simplex centres. However, it is shown in [10, p.19, Cor 5.4] that $\mathrm{Min}_1(\alpha)$ is connected by so-called simplicial paths. Since a non-trivial simplicial path between the centres of two simplices must pass through a point which is not at the centre of a simplex, our main result follows:

Theorem 8.7. *Let $\alpha \in Out(G, \mathcal{G})$ be irreducible and displacement 1. Then $Min_1(\alpha) = Fix_1(\alpha)$ is a single point.*

2 Groups acting on trees

For the duration of this chapter, let G be a group.

2.1 Metric simplicial trees

Definition 2.1. An $\mathbb{R}$-*tree* is a non-empty metric space in which any two points are joined by a unique arc, and in which every arc is isometric to a closed interval in the real line.

Definition 2.2. Let p be a point in a non-trivial $\mathbb{R}$-tree T.

- p is called a *branch point* if $T - p$ has three or more components.
- p is called *regular* if $T - p$ has exactly two components.
- p is called *external* otherwise

Points which are not regular are called *non-regular*.

Definition 2.3. A *metric simplicial tree* is an $\mathbb{R}$-tree whose set of non-regular points is discrete.

It will be useful to give these metric simplicial trees a combinatorial structure. Let C be a 1-dimensional simplicial complex. The 1-simplices will be called *edges*, and the 0-simplices will be called *vertices*. One can construct metric simplicial trees from simplicial complexes as follows:

Definition 2.4. Let C be a 1-dimensional simplicial complex. The *geometric realisation* of C is the metric space obtained from C by assigning the length of every edge in C to be 1. We give the geometric realisation the path metric topology.

Theorem 2.5. *An $\mathbb{R}$-tree T is metric simplicial if and only if it is homeomorphic to the geometric realisation of a connected* 1*-dimensional simplicial complex C with trivial fundamental group.*

Remark 2.6. There exists an alternative way to view this construction. Let T and C be as in Theorem 2.5. Then one can think of T as being obtained from C by assigning a length $L(e)$ (not necessarily 1) to every edge e in C, and to ensure discreteness of the non-regular points we impose the condition that for every vertex $v \in C$,
$\inf\{L(e) \mid e \text{ is incident to } v\} > 0$.

It is important to note that if T is a metric simplicial tree, then the geometric realisation to which T is homeomorphic to is not unique: Dividing any edge of C into two by adding a new vertex will result in a new simplicial complex whose geometric realisation is also homeomorphic to T. Thus we have a degree of choice over the structure of our trees. Once a choice of C has been made, we shall simply say "e is an edge of T" to mean that e is an edge of C, and similar for vertices.

When T is acted upon by a group G, the action can be used to determine our structure, as described in the next section. The conventional way of doing this, which we shall also be using, is described in the next section.

2.2 G-trees

Definition 2.7. Let T be a metric simplicial tree with an underlying simplicial complex. We define a *subforest* of T to be a subspace given by a set E' of edges and a set V' of vertices, such that the incident vertices of every edge in E' lie in V'.

A *subtree* of T is a connected subforest.

What this definition means is that we are defining subforests and subtrees so that they respect the underlying simplicial complex. In general this is not required, but it suits our purposes for this paper.

Definition 2.8. Suppose G acts on a metric simplicial tree T with an underlying simplicial complex.

- The action is said to be *simplicial* if it maps vertices to vertices and edges to edges.
- If no edge of T is sent to its inverse by any element of G, we say that G *acts without inversions*.
- We say that both T and the action are *minimal* if T contains no proper, G-invariant subtree.
- Let $x \in T$. The *stabiliser stab*(x) *of* x is defined to be the subgroup $\{g \mid x \cdot g = x\}$ of G.

  Let e be an edge of T. Then we define the *stabiliser stab*(e) *of* e to be the subgroup of G which fixes e but does not necessarily preserve the orientation of e.

  If every edge in T has trivial stabiliser, we say that T is *edge-free*.

  Let p be a vertex in T. If $\text{stab}(p) = 1$, we say that p is *free*. Otherwise, it is *non-free*.

Definition 2.9. A *G-tree* is a triple $(T, d_T, \cdot)$, where T is a metric simplicial tree, d_T is the metric on T, and $\cdot$ is an isometric group action $T \times G \to T,\ (x, g) \mapsto x \cdot g$.

If the metric and action are obvious from context, we may choose to omit one or both of them from the notation.

Remark 2.10. We have chosen to define G-trees with a right action so that $\forall x \in T$, $\forall g \in G$, and $\forall H \leqslant G$, $\text{stab}(x \cdot g) = \text{stab}(x)^g$ and $\text{Fix}(H^g) = \text{Fix}(H) \cdot g$.

Had we chosen to act on the left, acting by g would have caused the stabilisers to be conjugated by g^{-1}, and similar for the fixed point sets.

We are now ready to choose a simplicial structure for our G-trees. Let T be a G-tree. Then the *simplest* structure on T - that is, the structure containing the fewest vertices - is obtained by defining the vertex set to be the set of non-regular points of T. One then takes the edge set to be the set of simple arcs between elements of the vertex set which do not contain any other vertices.

Using the simplest structure, the action of G on T is simplicial. However, some edges may be sent to their inverses by elements of G. Future calculations will be easier if we have an action without inversions, therefore we shall instead use the following structure:

- We define the vertex set to be the set of non-regular points of T, together with the midpoints of all the edges of the simplest structure which were inverted by an element of G. We denote this vertex set by $V(T)$.
- We then define the edge set to be the set of simple arcs between elements of the vertex set which do not contain any other vertices. We denote this edge set by $E(T)$.

Essentially, we divide each inverted edge into two edges by placing a new vertex at its midpoint. With this new structure the action is still simplicial and, in addition, it is without inversions. This is the simplicial structure we shall be giving to all G-trees throughout this paper.

Remark 2.11. $V(T)$ as defined above is a discrete set, and hence is a well defined vertex set.

The following two propositions follow immediately from this choice of simplicial structure.

Proposition 2.12. *Minimal G-trees do not contain any degree 1 vertices (and hence the set of non-regular points is exactly the set of branch points).*

Proposition 2.13. *G-trees do not contain any free vertices of degree 2.*

2.3 Equivalence

Definition 2.14. Let $(T, d_T, \cdot)$, $(S, d_S, *)$ be G-trees. We say a map of trees $f : T \to S$ is a *G-equivariant* map from $(T, d_T, \cdot)$ to $(S, d_S, *)$ if $f(x \cdot g) = f(x) * g$ for all $x \in T$, for all $g \in G$.

Definition 2.15. Two G-trees $(T, d_T, \cdot)$, $(S, d_S, *)$ are said to be *equivalent* if there exists a G-equivariant isometry between them. We write $(T, d_T, \cdot) \sim (S, d_S, *)$ to denote equivalence.

Definition 2.16. We say a map of G-trees is *simplicial* if it maps vertices to vertices. Note that it does not have to map edges to edges, and hence this definition differs from that of a simplicial group action.

3 Bass-Serre Theory

3.1 Graphs of Groups

Definition 3.1. [3, p.113] A *graph* Y consists of the following:

- Two disjoint sets $V(Y)$ and $E(Y)$, called the *vertex* and *edge* sets of Y respectively.
- A function $^{-} : E(Y) \to E(Y)$ such that, for all $e \in E(Y)$, $\bar{e} \neq e$ and $\bar{\bar{e}} = e$.
- A function $\iota : E(Y) \to V(Y)$, and another function $\tau : E(Y) \to V(Y)$ defined by $\tau e := \iota\bar{e}$. We call ιe the *initial vertex* of e, and τe the terminal vertex of e.

We say Y is *finite* if $V(Y)$ and $E(Y)$ are both finite.

Graphs defined in this way - by considering each unoriented edge as a pair of oriented edges $(e, \bar{e})$ - are often referred to as *Serre graphs*.

Definition 3.2. A *metric graph* is a graph Y together with a length function $L : E(Y) \to \mathbb{R}$ such that, for all edges e of Y, $L(e) = L(\bar{e})$. This length function induces a metric d_Y on Y.

Diverting briefly back to the previous chapter, we remark that a 1-dimensional simplicial complex can be thought of as a Serre graph by considering each 1-simplex to be an edge pair. Thus G-trees (and their quotients) can be thought of as metric graphs, and depending on context we may treat them as such. This allows us to make the following observations:

Proposition 3.3. *G acts on G-trees via graph automorphisms (without inversions).*

Proposition 3.4. *Let T be a metric simplicial tree. Then T/G is finite if and only if it is compact (under the path metric topology).*

We now return to defining graphs of groups.

Definition 3.5. [3, p.198] A *graph of groups* X consists of:

(i) An connected graph Y

(ii) A group G_v for each vertex v of Y, and a group G_e for each edge e of Y such that $G_{\bar{e}} = G_e$.

(iii) For each edge e of Y, a monomorphism $\rho_e : G_e \to G_{\tau e}$, where ιe and τe are the endpoints of e.

If Y is a metric graph, then we say X is a *metric* graph of groups.

Let X be a graph of groups on a graph Y. One can define the fundamental group of X in a similar manner to that of a standard graph, by thinking of elements of the group as reduced loops in the graph. However, some additional structure is added by the edge and vertex groups. We shall use a definition adapted from [3, p.198], restricted to the case where X has trivial edge groups (and hence trivial monomorphisms ρ_e).

Let Y_0 be a spanning tree in Y. Then one can define the fundamental group of X to be

$$\pi_1(X) = \frac{\left(*_{v \in V(Y)} G_v \right) * F(E(Y))}{N}$$

where N is the normal closure of the set $\{e\bar{e} \mid e \in E(Y)\} \cup \{e \in Y_0\}$. This can be simplified into the form $\pi_1(X) \cong G_1 * \ldots * G_k * F_r$, where $G_1, \ldots, G_k$ are the nontrivial vertex groups, and $F_r \cong \pi_1(Y)$ is a free group of rank $r \geqslant 0$. This definition does not depend upon the choice of Y_0.

We will be working with graphs of groups whose fundamental group is isomorphic to a particular group G. Thus we consider pairs (X, ϕ), where X is a graph of groups and $\phi : G \to \pi_1(X)$ is an isomorphism. Such a pair is called a *marked graph of groups*, and ϕ is called the *marking*.

3.2 The Quotient Graph of Groups

Given a G-tree T, we can construct from it a metric graph of groups called the *quotient graph of groups*. A comprehensive method for constructing a quotient graph of groups from an arbitrary connected graph acted upon by G can be found on pages 204-205 of [3]. For this paper, we shall restrict his construction to edge-free G-trees.

Let T be a G-tree. Take the quotient graph T/G, and let $p : T \to T/G$ be the projection map, and Y_0 a maximal tree of T/G. Let $j : Y_0 \to T$ be a map such that $p \circ j$ is the identity on Y_0 (i.e j is a lift of Y_0 to T). We call $j(Y_0)$ a *representative tree* for the action.

We then define a graph of groups X on T/G as follows: For any vertex x of T/G, we define the vertex group G_x to be $\mathrm{stab}(j(x))$. We take all edge groups, and hence all edge monomorphisms, to be trivial, and edges inherit their lengths from T. This completely defines X.

The metric on X is given by assigning the length of each edge e of X to be the length of its corresponding edge orbit in T.

Theorem 3.6. *[3, p.210, Thm 26 (iii)] Let T be a G-tree, and let X be a quotient graph of groups for T. Then the fundamental group of X is isomorphic to G.*

This isomorphism gives a marking on X, and hence we can think of the quotient graph of groups as a marked graph of groups.

3.3 The Universal Cover

Conversely, let (X, ϕ) be a marked metric graph of groups with $G \overset{\phi}{\cong} \pi_1(X, v)$. Then we can construct from X a G-tree called the *Bass-Serre tree*, or *universal cover* of X, denoted by $\tilde{X}$. The process of constructing the universal cover is well-documented in the literature (e.g. [3, p.205]), so we shall not cover it here.

Definition 3.7. We say that two marked metric graphs of groups are *equivalent* if their universal covers are equivalent G-trees.

Theorem 3.8 (Fundamental Theorem of Bass-Serre Theory)**.** *The process of lifting to the universal cover and the process of descending to a quotient graph of groups are mutually inverse, up to equivalence of the structures involved.*

Let $T \in \mathcal{O}$. We observe that a quotient graph of groups (X, ϕ) of T is not unique: while the underlying graph is always T/G, the vertex groups depend on our choice of j, and hence X is not unique. Additionally, given a choice of X, the marking $\phi : G \to \pi_1(X)$ is not unique.

However, it follows from the fundamental theorem of Bass-Serre Theory that all possible choices of j and ϕ give equivalent marked graphs of groups. Thus, for brevity of notation, we shall simply denote a marked graph of groups (X, ϕ) by X.

4 Free Factor Systems and the Deformation Space

Let G be a group, and let T be a G-tree.

Definition 4.1. An element $g \in G$ is said to be *elliptic* (with respect to T) if it fixes a point in T. If g is not elliptic, we say it is *hyperbolic* (with respect to T).

We shall say a subgroup H of G is *elliptic* (with respect to T) if there exists a point $x \in T$ such that $x \cdot H = x$.

Definition 4.2 (Free Factor System)**.** Let T be a minimal, cocompact, edge-free G-tree, and let $\mathcal{G}_T$ denote the set of elliptic subgroups for T. We say $\mathcal{G}_T$ is a *free factor system* for G.

Note that this is not the usual definition of a free factor system. The usual definition can be found in [1, p.530-531], and we shall refer to it as a *traditional free factor system*:

Definition 4.3 (Traditional Free Factor System)**.** If $G_1 * \ldots * G_k * F_r$ is a free product decomposition for a group G, and each G_i is nontrivial, then we say that the collection $\{[G_1], \ldots, [G_k]\}$ of conjugacy classes is a *traditional free factor system*. The empty set $\varnothing$ is the trivial traditional free factor system of a free group F_r.

Corollary 4.11 will show that the Fundamental Theorem of Bass-Serre Theory provides a natural way to construct a free factor system from a traditional free factor system, and vice versa, and that these constructions are mutually inverse. In this sense, the two definitions are equivalent.

We have chosen our definition for two reasons. Firstly, the trivial free factor system must be defined separately when using the traditional definition. Secondly, our definition allows us to order our free factor systems by inclusion, and this corresponds to the somewhat more complicated ordering used for the traditional free factor systems, as defined in [1, p.532].

For now, we return to our definition of a free factor system and we observe the following properties:

Lemma 4.4. *Free factor systems are closed under conjugation and taking subgroups.*

Lemma 4.5. *Let $(T, \cdot), (S, *)$ be equivalent minimal, cocompact, edge-free G-trees. Then $\mathcal{G}_T = \mathcal{G}_S$.*

Lemma 4.6. *Let $(T, \cdot)$ be an edge-free G-tree. A nontrivial element of G cannot fix more than one point in T, and the fixed point will always be a vertex.*

The final lemma tells us that a subgroup of G is elliptic with respect to a minimal, cocompact, edge-free G-tree, and hence is an element of the corresponding free factor system, if and only if it is a vertex stabiliser or a subgroup of a vertex stabiliser.

Definition 4.7. Let $\mathcal{G}$ be a free factor system for G. The *deformation space* $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$ is the space of equivalence classes of minimal, cocompact, edge-free G-trees T such $\mathcal{G}_T = \mathcal{G}$.

By Lemma 4.6, a subgroup of G is in $\mathcal{G}$ if and only if it is a vertex stabiliser or subgroup of a vertex stabiliser for some (and hence every) G-tree in $\mathcal{O}$. Additionally, by general properties of group actions, two vertices lie in the same orbit if and only if they have conjugate stabilisers. Thus we can define a *minimal generating set* for $\mathcal{G}$:

Definition 4.8. We say a subset of a free factor system $\mathcal{G}$ is a *minimal generating set* for $\mathcal{G}$ if and only if it contains exactly one vertex stabiliser from every orbit of non-free vertices in some (and hence every) $T \in \mathcal{O}$.

Remark 4.9. If $\mathcal{G}$ is non-trivial, then this definition is the conventional definition of a minimal generating set under the operations of conjugation and taking subgroups.

If $\mathcal{G} = \{1\}$ - the trivial free factor system - then trees in $\mathcal{O}(G, \mathcal{G})$ do not contain any non-free vertices. Thus the minimal generating set for $\mathcal{G} = \{1\}$ is the empty set $\varnothing$.

G-trees in $\mathcal{O}$ are cocompact, so by Proposition 3.4, T will have a finite number of vertex orbits. Hence a minimal generating set for $\mathcal{G}$ will always be finite.

Theorem 4.10. *The following are equivalent:*

(i) *There exists a minimal, cocompact, edge-free G-tree T containing a representative tree T_0 in T such that the non-trivial vertex stabilisers in T_0 are exactly $G_1, \ldots, G_k$.*

(ii) *There exists a minimal, cocompact, edge-free G-tree T and a quotient graph of groups on T/G whose non-trivial vertex groups are exactly $G_1, \ldots, G_k$.*

(iii) *G can be written as a free product $G = G_1 * \ldots * G_k * F_r$, where F_r is a free group of rank $r \geqslant 0$.*

Proof. $(i) \Leftrightarrow (ii)$ Follows from the Fundamental Theorem of Bass-Serre Theory.

$(ii) \Rightarrow (iii)$ Follows immediately from the definition of the fundamental group of a graph of groups and Theorem 3.6.

$(iii) \Rightarrow (ii)$ We must first consider the case where $k = 1$ and $r = 0$ - that is to say, the trivial free product decomposition. In this case, we take X to be the graph of groups consisting of a single vertex with vertex group G.

Otherwise, we take X to be the graph of groups which consists of:

- a rose with central vertex v_∞ (with trivial vertex group) and r petals.
- for each $i \in \{1, \ldots, k\}$:
  - a vertex v_i with associated vertex group G_i
  - an edge e_i between v_i and v_∞
- trivial edge groups for every edge

Then $\pi_1(X) = G_1 * \ldots * G_k * F_r = G$. Upon lifting to the universal cover, it can be seen that this is a quotient graph of groups for a minimal, cocompact, edge-free G-tree. $\square$

We can now demonstrate the correspondence between free factor systems and traditional free factor systems.

Let $\mathcal{G}$ be a free factor system for a group G. This means that $\mathcal{G}$ is the set of elliptic subgroups of some $T \in \mathcal{O}(G, \mathcal{G})$. Take a representative tree in T, and let $\{G_1, \ldots, G_k\}$ be the nontrivial vertex groups of this representative tree. Then by Theorem 4.10, G can be written as a free product $G = G_1 * \ldots * G_k * F_r$. Thus the set $\{[G_1], \ldots, [G_k]\}$ is a traditional free factor system. Since each $[G_i]$ is a conjugacy class, this set does not depend on the choice of representative tree.

Conversely, let $\{[G_1], \ldots, [G_k]\}$ be a traditional free factor system. Then we have $G = G_1 * \ldots * G_k * F_r$, so by Theorem 4.10, there exists a minimal, cocompact, edge-free G-tree T containing a representative tree T_0 in T such that the non-trivial vertex stabilisers in T_0 are exactly $G_1, \ldots, G_k$. These vertex groups give a minimal generating set for a free factor system $\mathcal{G} = \mathcal{G}_T$.

Corollary 4.11. *The process of constructing a free factor system from a traditional free factor system, and the process of constructing a traditional free factor system from a free factor system, are mutually inverse. In this sense, the two definitions are equivalent.*

Proof. Follows from Theorem 4.10. □

It follows from Corollary 4.11 that, given a free product $G = G_1 * \ldots * G_k * F_r$, we can construct a free factor system $\mathcal{G}$, and hence a corresponding deformation space $\mathcal{O}(G, \mathcal{G})$. This space can then be used to study the automorphisms of the free product.

5 Automorphisms

For the duration of this section, let $\mathcal{G}$ denote a free factor system for a group G, and let $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$.

5.1 Acting on the Deformation Space

Notation. The outer automorphism group of G is defined as $\mathrm{Out}(G) := {\mathrm{Aut}(G)}/{\mathrm{Inn}(G)}$; elements of $\mathrm{Out}(G)$ are equivalence classes of automorphisms, where two automorphisms are equivalent if they differ by an inner automorphism.

When we write $\alpha \in \mathrm{Out}(G)$, we mean that α is an automorphism in $\mathrm{Aut}(G)$ representing an equivalence class in $\mathrm{Out}(G)$.

In this paper, the automorphisms of G will act on G on the right.

Definition 5.1. Let $\alpha \in \mathrm{Aut}(G)$, and let $\mathcal{G}\alpha = \{(H)\alpha \mid H \in \mathcal{G}\}$. We say that $\mathcal{G}$ is α-*invariant* if $\mathcal{G}\alpha = \mathcal{G}$.

Free factor systems are closed under conjugation by elements of G, hence α-invariance depends only on the outer automorphism class of α. Thus we can make a similar definition for $\mathrm{Out}(G)$:

Definition 5.2. Let $\alpha \in \mathrm{Out}(G)$, and let $\mathcal{G}\alpha = \{(H)\alpha \mid H \in \mathcal{G}\}$. We say that $\mathcal{G}$ is α-*invariant* if $\mathcal{G}\alpha = \mathcal{G}$.

The set of $\alpha \in \mathbf{Out}(G)$ leaving $\mathcal{G}$ invariant forms a group, which we shall denote $\mathrm{Out}(G, \mathcal{G})$.

The group $\mathrm{Out}(G, \mathcal{G})$ admits a left action on the deformation space $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$: Let $(T, d_T, \cdot) \in \mathcal{O}$, let $\alpha \in \mathrm{Out}(G, \mathcal{G})$. Then $\alpha(T, d_T, \cdot) := (T, d_T, \cdot_\alpha)$, the G-tree with the same underlying simplicial tree and metric, but with 'twisted' action given by $x \cdot_\alpha g = x \cdot (g)\alpha$ for all $x \in T$.

Observe the following:

Lemma 5.3. *Let $T \in \mathcal{O}$ be a G-tree, and let $\alpha \in Out(G, \mathcal{G})$. Then the G-orbits of T are the same as those of αT. That is, for all $x \in T$, $x \cdot G = x \cdot_\alpha G$.*

Proof. $x \cdot G = \{x \cdot g \mid g \in G\} = \{x \cdot (g)\alpha \mid g \in G\} = x \cdot_\alpha G$. ☐

Thus we see that the 'twists' α applies to the action only occur within each orbit. This means that if we are working with both $(T, \cdot)$ and $(T, \cdot_\alpha)$, we are able to simply refer to 'a G-orbit of T' without having to state which action is being used.

Definition 5.4. We can partially order the set of all free factor systems of G by inclusion. Let $\mathcal{G}$ be a proper, α-invariant free factor system. We say $\alpha \in \mathrm{Out}(G, \mathcal{G})$ is $\mathcal{G}$-*irreducible*, or *irreducible with respect to* $\mathcal{G}$, if $\mathcal{G}$ is a maximal, proper α-invariant free-factor system.

Otherwise, we say α is *reducible* with respect to $\mathcal{G}$.

5.2 Topological Representatives

Definition 5.5. [9, p.16] Let $T, S \in \mathcal{O}(G, \mathcal{G})$. An $\mathcal{O}$-*map* $f : T \to S$ is a G-equivariant, Lipschitz continuous function. The Lipschitz constant of f is denoted $\mathrm{Lip}(f)$.

Note that an $\mathcal{O}$-map does not have to send vertices to vertices, and hence does not need to be a graph morphism.

Definition 5.6. [9, p.16] We say an $\mathcal{O}$-map $f : T \to S$ is *straight* if it has constant speed on edges - that is, for each edge e in T, there exists a non-negative number $l_e(f)$ such that for any $a, b \in e$ we have $d_T(f(a), f(b)) = l_e(f) d_S(a, b)$.

Definition 5.7. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, and let $T \in \mathcal{O}$ be a G-tree. Then a map $f : T \to \alpha T$ is said to *topologically represent* α if it is a straight $\mathcal{O}$-map.

The authors of [9, p. 16] make the following remark:

Remark 5.8. Any two trees $T, S \in \mathcal{O}$ have an $\mathcal{O}$-map between them. Furthermore, any $\mathcal{O}$-map $f : T \to S$ can be uniquely 'straightened' - that is to say, there exists a unique straight $\mathcal{O}$-map $\mathrm{Str}(f) : T \to S$, such that $\mathrm{Str}\big(f\big)(v) = f(v)$ for every vertex $v \in T$. We have $\mathrm{Lip}(\mathrm{Str}(f)) \leqslant \mathrm{Lip}(f)$.

From this remark it follows that $\forall \alpha \in \mathrm{Out}(G, \mathcal{G})$, $\forall T \in \mathcal{O}$, there exists a topological representative $f : T \to \alpha T$.

Definition 5.9. Let F be a subforest of some $T \in \mathcal{O}$, and let A be a component of F. We define the *stabiliser* of A to be the set $\mathrm{stab}(A) = \{g \in G \mid A \cdot g = A\}$ - that is, we are taking the setwise stabiliser, not the pointwise stabiliser.

We say that F is $\mathcal{G}$-elliptic if, for every component A of F, $\mathrm{stab}(A) \in \mathcal{G}$. Otherwise we say that F is $\mathcal{G}$-hyperbolic.

Using topological representatives, we can construct a test for the reducibility of an automorphism:

Theorem 5.10. *Let $\mathcal{G}$ be a proper free factor system for a group G, let $\alpha \in Out(G, \mathcal{G})$, and let $T \in \mathcal{O}$.*

Suppose that α can be topologically represented by a G-equivariant simplicial map $f : T \to \alpha T$, and there exists a proper f-invariant, G-invariant, $\mathcal{G}$-hyperbolic subforest of T. Then α is reducible with respect to $\mathcal{G}$.

Proof. We shall prove that α is reducible by constructing a new cocompact, minimal, edge-free G-tree S, from which we shall retrieve another proper α-invariant free factor system $\mathcal{H}$ such that $\mathcal{G} \subset \mathcal{H}$.

Let F denote the subforest of T described above, together with all the remaining vertices of T. This extended subforest is still proper, f-invariant, and G-invariant. (In addition, recall that we defined subforests such that they respect the simplicial structures of our G-trees; therefore the complement of F is a set of edges.)

We obtain S from T by collapsing each component A of F to a point p_A. Edges which were not collapsed inherit their lengths from T, giving us a metric on S. Since F is G-invariant, this collapse induces a minimal, isometric action of G on S. Thus S is a G-tree.

Furthermore, if we declare the vertex set to be the set of p_A, we induce a new simplicial structure on S. (This is a well-defined vertex set, and the simplicial structure given by this vertex set is exactly the same as the usual simplicial structure we give to all G-trees, as defined in Section 2.2).

S inherits cocompactness and edge-freeness from T. Hence, by definition, the set $\mathcal{H}$ of elliptic subgroups for S is a free factor system for G.

Let $v \in T$ be a vertex. Since F contains every vertex of T, v must lie in some component A of F. Therefore, since F is G-invariant, all of A must be fixed (setwise, not necessarily pointwise) by $\text{stab}(v)$. Hence $\text{stab}(v) \leqslant \text{stab}(p_A)$. This holds for all v, which is enough to tell us that $\mathcal{G} \sqsubseteq \mathcal{H}$.

Some component of F has $\mathcal{G}$-hyperbolic stabiliser. This means that $\text{stab}(p_A) \notin \mathcal{G}$ but $\text{stab}(p_A) \in \mathcal{H}$. Thus $\mathcal{G} \sqsubset \mathcal{H}$.

Suppose that $G \in \mathcal{H}$ - that is to say, a vertex of S is stabilised by G. Then the component of F corresponding to this vertex is a G-invariant subtree of T, contradicting the minimality of T. Hence $G \notin \mathcal{H}$, so $\mathcal{H}$ is a proper free factor system.

Finally, we must show that $\mathcal{H}$ is α-invariant:

Let p_A be a vertex of S. We first want to show that $(\text{stab}(p_A))\alpha$ lies in $\mathcal{H}$ - that is to say, it fixes a point in $(S, *)$. F is f-invariant, therefore $f(A)$ is a component of F and $p_{f(A)}$ is a vertex of S. Furthermore, $\forall (g)\alpha \in (\text{stab}(v_A))\alpha$,
$p_{f(A)} * (g)\alpha = p_{f(A)\cdot(g)\alpha} = p_{f(A\cdot g)} = p_{f(A)}$, and hence $(\text{stab}(p_A))\alpha \in \mathcal{H}$. This holds for all $p_A \in S$. This tells us that $\mathcal{H}\alpha \sqsubseteq \mathcal{H}$, which in turn is enough to show that $\mathcal{H} = \mathcal{H}\alpha$.

To summarize, $\mathcal{H}$ is a proper, α-invariant free-factor system for G, and $\mathcal{G} \sqsubset \mathcal{H}$. Hence, by Definition 5.4, α is reducible with respect to $\mathcal{G}$. □

5.3 Isometric topological representatives

Our main results will make use of *isometric* topological representatives, which allow us to make some additional observations:

Remark 5.11. Recall that $\mathcal{O}$ is a space of equivalence classes of G-trees, where two trees are equivalent if there exists an equivariant isometry between them. Topological representatives are equivariant; therefore, if an isometric topological representative $f : T \to \alpha T$ exists, the two G-trees T and αT are representing the same point in $\mathcal{O}$.

Proposition 5.12. *Let $f : T \to \alpha T$ be a topological representative for some $T \in \mathcal{O}(G, \mathcal{G})$, for some $\alpha \in Out(G, \mathcal{G})$. If f is an isometry, then it is also a graph automorphism.*

Proof. It is sufficient to show that $f(v)$ is a vertex if and only if v is a vertex. f is an isometry - in particular it is bijective - therefore $f(v)$ is a branch point if and only if v is a branch point. The only vertices which remain are the degree 2 vertices. Recall that these were introduced as the midpoints of inverted edges, hence they all have stabiliser of order 2. f is equivariant, therefore $\mathrm{stab}(v)$ is order 2 if and only if $\mathrm{stab}(f(v))$ is order 2. Thus the set of degree 2 vertices is also preserved, and f is a graph automorphism. □

Definition 5.13. Let Y be a metric graph. The *volume* of Y, denoted $\mathrm{Vol}(Y)$ is defined to be the sum of the lengths of the edges of Y.

Let $T \in \mathcal{O}$. The *covolume* of T, denoted $\mathrm{Covol}(T)$, is defined to be the volume of the graph T/G.

Proposition 5.14. *Let $f : T \to \alpha T$ be a topological representative for some $T \in \mathcal{O}(G, \mathcal{G})$, for some $\alpha \in Out(G, \mathcal{G})$. Then $Lip(f) = 1$ if and only if f is an isometry.*

Proof. If f is an isometry, then $\mathrm{Lip}(f) = 1$ follows immediately. It remains to prove the converse.

Let D be a subforest of T consisting of exactly one edge from each orbit. Then $\mathrm{Covol}(T) = \mathrm{Vol}(D)$. Without loss of generality, we may assume that $\mathrm{Covol}(T) = 1$. Since T and αT have the same metric, this means that $\mathrm{Covol}(\alpha T) = 1$.

Since f is equivariant, $f(D)$ contains a fundamental domain for $f(T)$. Since αT does not contain any proper invariant subtrees, we must have $f(T) = \alpha T$, hence $f(D)$ contains a fundamental domain for αT. It follows that $\mathrm{Vol}(f(D)) \geqslant 1$. In addition,

$$\begin{aligned}
\mathrm{Vol}(f(D)) &\leqslant \sum_{\text{edges } e \in D} L(f(e)) && (*) \\
&\leqslant \sum_{\text{edges } e \in D} L(e) && (\text{as } \mathrm{Lip}(f) = 1) \\
&= \mathrm{Vol}(D) \\
&= \mathrm{Covol}(T) \\
&= 1
\end{aligned}$$

.

We split into two cases:

Case 1: f is not locally injective This is equivalent to saying that f 'folds' a pair of edges - that is, there is a vertex v, neighbourhoods U_1, U_2 of v, and edges e_1, e_2 incident to v such that $f(e_1 \cap U_1) = f(e_2 \cap U_2)$. (The neighbourhoods are required because f may not fold the entirety of the edges, only the initial segments. Since f may stretch these segments, the neighbourhoods are not the same size in general).

The two edges can only be folded if they lie in different orbits; observe that if $e_1 \cdot g = e_2$, then $f(e_1 \cap U_1)$ must be fixed by $\alpha(g)$, contradicting edge freeness. Therefore we are free to choose D such that it contains a pair of folded edges.

If e_1, e_2 are a pair of folded edges in D, then the volume of their image under f will be strictly less then the sum of their original lengths. This means that (*) is a strict inequality, so $\mathrm{Vol}(f(D)) < 1$. This contradicts $\mathrm{Vol}(f(D)) \geqslant 1$, hence Case 1 cannot occur.

Case 2: f is locally injective. Then (*) is an equality, so $\mathrm{Vol}(f(D)) = 1$. This is enough to tell us that $L(f(e)) = L(e)$ for every edge e of D, and hence all the edges of T; thus f is an isometry on every edge of T. This, combined with local injectivity, means that f is an isometry on all of T.

□

6 Distance on $\mathcal{O}$

6.1 Stretching Factors

Definition 6.1. Let $g \in G$, and let $T \in \mathcal{O}(G, \mathcal{G})$. The *translation length* of g in T, denoted $l_T(g)$, is defined as

$$l_T(g) = \inf_{x \in T}\{d_T(x, x \cdot g)\}.$$

Remark 6.2. This infimum is in fact a minimum, and is obtained for some x. If g is elliptic then this is observed to be true from the definition of an elliptic element, and we have $l_T(g) = 0$.

If g is hyperbolic then the translation length will be non-zero, and the set of elements realising this length will form a line through T called the *hyperbolic axis of g*. Points on the axis will be translated along the axis by $l_T(g)$.

Remark 6.3. An equivalence class of G-trees in $\mathcal{O}$ is uniquely determined by its translation length function [4] - thus one can think of $\mathcal{O}$ as being embedded in the space $\mathbb{R}^G$.

Definition 6.4. Let $\mathrm{Hyp}(\mathcal{G})$ denote the set of elements of G which do not lie in any subgroup of $\mathcal{G}$. (In other words, $\mathrm{Hyp}(\mathcal{G})$ is the set of elements which are hyperbolic with respect to some, and hence all, G-trees in $\mathcal{O}(G, \mathcal{G})$).

Definition 6.5. [8, p.8] Let $T, S \in \mathcal{O}(G, \mathcal{G})$. Then we define the *left and right stretching factor* from T to S as

$$\Lambda_L(T, S) := \sup_{g \in \mathrm{Hyp}(\mathcal{G})} \frac{l_T(g)}{l_S(g)}, \qquad \Lambda_R(T, S) := \sup_{g \in \mathrm{Hyp}(\mathcal{G})} \frac{l_S(g)}{l_T(g)} = \Lambda_L(S, T),$$

respectively. We also define the *symmetric stretching factor* from T to S to be

$$\Lambda(T, S) := \Lambda_L(T, S)\Lambda_R(T, S).$$

The next Theorem follows from [8, Corollary 6.8, p.18 and Theorem 6.11, p.19].

Theorem 6.6. *Let $T, S \in \mathcal{O}$. Then there exists a Lipschitz continuous map $f : T \to S$ such that $Lip(f) = \Lambda_R(T, S)$.*

6.2 The Displacement of an Automorphism

Definition 6.7. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$. Then we define the *displacement* of α to be

$$l_\alpha := \inf_{T \in \mathcal{O}} \Lambda_R(T, \alpha T),$$

Theorem 6.8. *[8, p. 25] For any $\mathcal{G}$-irreducible $\alpha \in Out(G, \mathcal{G})$, the displacement of α is a minimum and obtained for some $T \in \mathcal{O}$.*

Definition 6.9. For any $\alpha \in \mathrm{Out}(G, \mathcal{G})$, we define

$$\mathrm{Min}(\alpha) = \{T \in \mathcal{O} \mid \Lambda_R(T, \alpha T) = \lambda_\alpha\},$$

That is to say, $\mathrm{Min}(\alpha)$ is the set of all T which realise the above infimum.

Theorem 6.10. *Let $\alpha \in Out(G, \mathcal{G})$ be a $\mathcal{G}$-irreducible, displacement 1 automorphism. Then for all $T \in Min(\alpha)$, there exists an isometric topological representative for α on T.*

Proof. Let $T \in \mathrm{Min}(\alpha)$. Then by definition of the minimally displaced set, $\Lambda_R(T, \alpha T) = \lambda_\alpha = 1$, and hence by Theorem 6.6 there exists a Lipschitz continuous map $f : T \to \alpha T$ with $\mathrm{Lip}(f) = 1$. Therefore, by Proposition 5.14, f is an isometric topological representative for α. □

Corollary 6.11. *Let $\alpha \in Out(G, \mathcal{G})$ be a $\mathcal{G}$-irreducible, displacement 1 automorphism. Then $Min(\alpha) = Fix(\alpha)$.*

Proof. Let $T \in \mathrm{Fix}(\alpha)$. Then T and αT are equivalent G-trees, so $\Lambda(T, \alpha T) = 1 = \lambda_\alpha$. Thus $\mathrm{Fix}(\alpha) \subseteq \mathrm{Min}(\alpha)$.

Conversely, let $T \in \mathrm{Min}(\alpha)$. By Theorem 6.10, there exists an equivariant isometry from T to αT. Points in $\mathcal{O}$ are equivalence classes of G-trees under equivariant isometry, hence T and αT represent the same point in $\mathcal{O}$. Thus $\mathrm{Min}(\alpha) \subseteq \mathrm{Fix}(\alpha)$. □

7 Secondary Theorem

We extend some of our terminology for G-trees to graphs of groups:

Definition 7.1. Let $T \in \mathcal{O}(G, \mathcal{G})$, and let X be a quotient graph of groups on T/G. We shall say that a vertex of T/G is *free* if it has trivial vertex group in X. Otherwise, it is *non-free*. Note that this definition does not depend on our choice of X.

Definition 7.2. Let $T \in \mathcal{O}(G, \mathcal{G})$, and let X be a quotient graph of groups on T/G. We say that a subgraph-of-groups of X is $\mathcal{G}$-elliptic if and only if the fundamental group of all its components lies in $\mathcal{G}$. Otherwise, we say it is $\mathcal{G}$-hyperbolic.

Similarly, we say that a subgraph of T/G is $\mathcal{G}$-elliptic/hyperbolic if the corresponding subgraph-of-groups of X is $\mathcal{G}$-elliptic/hyperbolic. Observe that this definition does not depend on the choice of marking on X.

It follows from the Fundamental Theorem of Bass-Serre Theory that a G-invariant subforest of T is $\mathcal{G}$-elliptic if and only if it collapses to a $\mathcal{G}$-elliptic subgraph of T/G.

We also observe that, by the definition of the fundamental group, a subgraph of T/G will be G-elliptic if and only if each component is a tree containing at most one non-free vertex.

Let $f : T \to \alpha T$ be a topological representative. Topological representatives are equivariant, hence f induces a well-defined map $\varphi : T/G \mapsto T/G$. (Observe that since orbits in T and αT are the same, $T/G = \alpha T/G$).

Suppose that f is an isometry. Then by Proposition 5.12, f is a graph automorphism. It follows that φ is also an isometric graph automorphism - in particular, it is invertible. Thus we can think of the cyclic group $\langle \varphi \rangle$ as acting on T/G.

φ can be used in an equivalent form of the reducibility test (Theorem 5.10), this time using the quotient graph:

Theorem 7.3. *Let $\mathcal{G}$ be a proper free factor system for a group G, let $\alpha \in Out(G, \mathcal{G})$, and let $T \in \mathcal{O}$.*

Suppose that α can be topologically represented by a G-equivariant simplicial map $f : T \to \alpha T$, and there exists a proper φ-invariant, $\mathcal{G}$-hyperbolic subgraph of T/G. Then α is reducible with respect to $\mathcal{G}$.

This form of the reducibility test eliminates the need to check for G-invariance. Note that a subgraph of T/G is φ-invariant if and only if it is invariant under the action of $\langle \varphi \rangle$.

Definition 7.4. We say a graph Y is a *star* if it is a tree and there exists a vertex w which is incident to every edge of Y.

Lemma 7.5. *Let $T \in \mathcal{O}$, and let X be a quotient graph of groups for T. Then all the vertices of degree 1 or 2 in X will have non-trivial vertex groups.*

The proof of this lemma follows directly from Propositions 2.12 and 2.13.

Theorem 7.6. *Let $\mathcal{G}$ be a free factor system for a group G, and let $\alpha \in Out(G, \mathcal{G})$ be an irreducible automorphism with $l_\alpha = 1$. Let $T \in Min(\alpha)$, and let $f : T \to \alpha T$ be an equivariant isometry (f exists by Theorem 6.10).*

Then f cyclically permutes the G-orbits of edges in T.

Proof. f induces a map φ on T/G, and the cyclic group $\langle\varphi\rangle$ acts on T/G. The theorem statement is equivalent to saying that φ cyclically permutes the edges of T/G.

Let e be an edge in T/G. We define two subgraphs of T/G:

- Let A be the subgraph of T with $E(A) = e \cdot \langle\varphi\rangle$ and with $V(A)$ equal to the set of vertices incident to $E(A)$.
- Let B be the subgraph of T with edge set $E(B) = E(T) - e \cdot \langle\varphi\rangle$ and with $V(B)$ equal to the set of vertices incident to $E(B)$.

Observe that these are both $\langle\varphi\rangle$-invariant.

Suppose that φ does not cyclically permute the edges of T/G. This means that A and B are both proper subgraphs of T/G. We shall show that at least one of the two subgraphs is $\mathcal{G}$-hyperbolic. This will mean that α is reducible by Theorem 7.3, giving us a contradiction.

In more detail, we assume that A is $\mathcal{G}$-elliptic. Then A is a forest such that every component contains at most one non-free vertex. If B is not a forest, then B is immediately $\mathcal{G}$-hyperbolic, so assume that it is a forest. We shall show that some component of B contains at least 2 non-free vertices, and hence B is $\mathcal{G}$-hyperbolic.

To begin we note that by Lemma 7.5, X does not contain any free vertices of degree 1 or 2. Thus we make the following claim:

Claim (i): Let v be a free vertex of T/G. Then $\deg_{T/G}(v) = \deg_A(v) + \deg_B(v) \geqslant 3$. It follows that, if $\deg_A(v) = 1$ or 2 *or* $\deg_B(v) = 1$ or 2, then $v \in A \cap B$. Additionally, if $\deg_A(v) = 1$, then $\deg_B(v) \geqslant 2$, and if $\deg_B(v) = 1$, then $\deg_A(v) \geqslant 2$.

Now, $\langle\varphi\rangle$ acts via isometries, and since A is the $\langle\varphi\rangle$-orbit of a single edge, $\langle\varphi\rangle$ acts transitively on the components of A. Hence the components of A are all isometric to each other, and we can divide A into two cases:

Case 1: Each component of A is a single edge

B is a finite forest, therefore each component of B has at least two vertices of B-degree 1. By Claim (i), if any of these vertices are free, then they must lie in $A \cap B$ and have

A-degree of at least 2. However, all vertices in A have A-degree 1. Hence the B-degree 1 vertices are all non-free, and B is $\mathcal{G}$-hyperbolic.

Case 2: Each component of A contains more than one edge

By definition, A contains at most two $\langle\varphi\rangle$-orbits of vertices. Since A is a finite forest, some vertices in $V(A)$ will have A-degree 1, and since each component of A contains more than one edge, some vertices in $V(A)$ will have A-degree strictly greater than 1. $\langle\varphi\rangle$ acts via graph automorphisms (by Proposition 5.12), and A is $\langle\varphi\rangle$-invariant, therefore vertices in the same $\langle\varphi\rangle$-orbit will have the same A-degree. Thus A contains exactly two $\langle\varphi\rangle$-orbits of vertices.

Furthermore, $E(A)$ is the $\langle\varphi\rangle$-orbit of a single edge, so the incident vertices of this edge are representatives for our two vertex orbits. Thus every edge in A must have exactly one incident vertex with A-degree 1, and hence A is in fact a disjoint union of stars. We shall refer to the A-degree 1 vertices as the *spoke* vertices. The remaining vertices, at the centre of each star, shall be called the *hub* vertices.

By the equivariance of f, $\langle\varphi\rangle$ sends free vertices to free vertices, and non-free vertices to non-free vertices. The spoke vertices all lie in the same $\langle\varphi\rangle$-orbit, and each component of A contains at least 2 spoke vertices. Therefore, since A is $\mathcal{G}$-elliptic, the spoke vertices must all be free. Thus, by Claim (i), they must lie in B, and have B-degree at least 2.

$\langle\varphi\rangle$ acts transitively on the spoke vertices. Therefore $\langle\varphi\rangle$ acts transitively on the components of B which contain the spoke vertices, and hence these components are all isometric to each other. We shall write B' to denote the subforest of B consisting of these components. (Observe that, for any vertex $v \in B'$, $\deg_B(v) = \deg_{B'}(v)$).

We divide into two cases once again:

Subcase 1: Each component of B' contains exactly one spoke vertex

Let s be the number of spoke vertices, and let l be the number of components of A. Then $s \geqslant 2l$.

B' has exactly s components. These components are finite trees, so they will each contain at least 2 vertices of B'-degree 1. Thus B' has at least $2s$ vertices with B'-degree 1. By Claim (i), any of these vertices which are free must lie in A and have A-degree at least 2. However, the only vertices with A-degree at least 2 are the l hub vertices. This leaves at least $2s - l$ vertices in B' which must therefore be non-free. $2s - l > s =$ number of components of B', therefore some component of B' must contain two or more of these non-free vertices. Therefore B', and hence B, is $\mathcal{G}$-hyperbolic.

Subcase 2: Each component of B' contains more than one spoke vertex

We divide B' into two subforests, C and D:

- Let $B'_1, \dots, B'_n$ be the components of B'. For each $i = 1, \dots, n$, let C_i be the unique minimal subtree of B'_i which contains all the spoke vertices in B'_i. Let $C := \bigcup_{i=1}^{n} C_i$. Since B' and the set of spoke vertices are $\langle\varphi\rangle$-invariant, C is also $\langle\varphi\rangle$-invariant.

- Define D to be the subforest of B' consisting of the edge set $E(B') - E(C)$, together with all vertices incident to this edge set. Since B' and C are $\langle\varphi\rangle$-invariant, D is also $\langle\varphi\rangle$-invariant.

By minimality of the C_i's, at least one spoke vertex has C-degree 1. Since C is $\langle\varphi\rangle$-invariant and the spoke vertices lie in the same $\langle\varphi\rangle$-orbit, this implies that all the spoke vertices have C-degree 1. However, recalling that the spoke vertices have B-degree at least 2, this tells us that the spoke vertices are all incident to an edge in D (and hence the spoke vertices themselves are all in D).

Claim (ii): A component of D cannot contain more than one point in $C \cap D$.

Proof of Claim (ii). Let $v, w \in C \cap D$, and suppose that v and w lie in the same component of D. Then they lie in the same component of B', and hence the same component of C. Therefore there exists a unique reduced path γ_D from v to w in D , and a unique reduced path γ_C from v to w in C.

However, B' is a forest, therefore $\gamma_D = \gamma_C$. By definition of D, there are no edges in $C \cap D$. Hence both paths are trivial, and $v = w$. This ends the proof of Claim (ii). □

In particular, Claim (ii) implies that each spoke vertex lies in a unique component of D. Let D' be defined as the subforest of D consisting only of the components which contain spoke vertices. Then, if we let s be the number of spoke vertices, D' will have s components. (Additionally, for any $v \in D'$, $\deg_D(v) = \deg_{D'}(v)$).

Each component of D' will contain at least two vertices of D'-degree 1. By Claim (ii), at least one of these will not lie in $C \cap D$, and hence it will also have B'-degree 1. Furthermore, since $\langle\varphi\rangle$ acts transitively on the spoke vertices, it will act transitively on the components of D'. Thus there exists a $\langle\varphi\rangle$-orbit of at least s vertices with B'-degree 1; at least one in each component of D'.

By Claim (i), if this orbit of vertices is free, then it must lie in A and have A-degree at least 2. However, the only $\langle\varphi\rangle$-orbit of vertices with A-degree at least 2 are the l hub vertices. $s \geqslant 2l$, therefore these cannot be the same orbit. Hence our orbit of B'-degree 1 vertices must be non-free.

Each component of B' contains more than one spoke vertex. Therefore each component of B' contains more than one component of D', and hence more than one of our non-free vertices. Thus B' is $\mathcal{G}$-hyperbolic. □

8 Main Theorem

For the duration of this section, let $\mathcal{G}$ be a free factor system for a group G, and let $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$.

Definition 8.1. The *covolume 1 slice of* $\mathcal{O}$, denoted $\mathcal{O}_1$, is defined to be the subspace of covolume 1 trees in $\mathcal{O}$.

Definition 8.2. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$. In a similar manner to $\mathcal{O}$, we define the *minimally displaced set* in $\mathcal{O}_1$ to be $\mathrm{Min}_1(\alpha) = \{T \in \mathcal{O}_1 \mid \Lambda_R(T, \alpha T) = \lambda_\alpha\}$ and we define $\mathrm{Fix}_1(\alpha) = \{T \in \mathcal{O}_1 \mid T \sim \alpha T\}$.

Remark 8.3. For $T \in \mathcal{O}$ and $\mu > 0$, let us write μT to denote the G-tree $(T, \mu d_T, \cdot)$. One can show that for all $T, S \in \mathcal{O}$, $\Lambda_R(T, S) = \Lambda_R(\mu T, \mu S)$ - that is to say, stretching factors are invariant under rescaling the volume of both G-trees. Additionally, T and αT have the same volume for all $T \in \mathcal{O}$, so by rescaling one, we rescale the other. Thus we observe the following:

$$\mathrm{Min}(\alpha) = \{\mu T \in \mathcal{O} \mid T \in \mathrm{Min}_1(\alpha), \mu > 0\},$$
$$\mathrm{Fix}(\alpha) = \{\mu T \in \mathcal{O} \mid T \in \mathrm{Min}_1(\alpha), \mu > 0\},$$

It then follows from Corollary 6.11 that $\mathrm{Min}_1(\alpha) = \mathrm{Fix}_1(\alpha)$.

Let $T \in \mathcal{O}_1$. The metric on T can be completely described by the length of one edge from each G-orbit - or equivalently, the lengths of the edges of T/G. Hence, if there are n edge orbits with lengths $x_1, \ldots, x_n$, then the open simplex $\{(x_1, \ldots, x_n) \mid x_1 + \ldots x_n = 1, x_i > 0 \;\; \forall i\}$ describes the set of all possible metrics on T. Repeating this for every tree in $\mathcal{O}_1$ allows us to think of $\mathcal{O}_1$ as a union of open simplices.

(Equivalently, the same structure can be thought of a simplicial complex with some missing faces. These missing faces are a result of edges of T which, were their lengths reduced to zero, would create new vertices whose stabilisers were not in $\mathcal{G}$, and hence the resulting tree could not lie in $\mathcal{O}$.)

Let Δ be an open simplex in $\mathcal{O}_1$. We write $\overline{\Delta}$ to denote the closure of Δ in $\mathcal{O}_1$. Note that this is not, in general, a closed simplex.

Definition 8.4. (Adapted from [10, p.19, Def. 5.1])

Let $T, S \in \mathcal{O}_1$. A *simplicial path* between T and S is given by:

(i) A finite sequence of points $T = T_0, T_1, \ldots, T_k = S \in \mathcal{O}_1$ such that $\forall i = 1 \ldots k$ there is a simplex Δ_i such that X_{i-1} and X_i both lie in $\overline{\Delta_i}$.

(ii) Euclidean segments $\overline{X_{i-i}X_i} \subseteq \overline{\Delta_i}$. (Here Euclidean segment refers to the coordinates $(x_1, \ldots, x_n)$ on $\overline{\Delta_i}$).

Definition 8.5. We say that a set $\chi \subseteq \mathcal{O}_1$ is *connected by simplicial paths* if for any $x, y \in \chi$, there is a simplicial path between x and y which is contained entirely in χ.

Lemma 8.6. *A simplicial path in $\mathcal{O}_1$ which only passes through the centres of simplices is a single point.*

Proof. Any such simplicial path must begin at the centre of an open simplex in $\mathcal{O}_1$. Observe that any nontrivial Euclidean segment which begins at the centre of a simplex must pass through a point which does not lie at the centre of a simplex. Thus the entire simplicial path is trivial. □

We can now state the main theorem of this paper.

Theorem 8.7. *Let $\alpha \in Out(G, \mathcal{G})$ be irreducible and displacement 1. Then $Min_1(\alpha) = Fix_1(\alpha)$ is a single point.*

Proof. By Theorem 6.8, $\mathrm{Min}(\alpha)$, and hence $\mathrm{Min}_1(\alpha)$, is non-empty.

Let $T \in \mathrm{Min}_1(\alpha)$. Then by Theorem 6.10 there exists an equivariant isometry $f : T \to \alpha T$. By Theorem 7.6, f cyclically permutes the edges of T, which means that all the edges in T must have the same length, and hence T must lie at the centre of an open simplex in $\mathcal{O}_1$. Thus $\mathrm{Min}_1(\alpha)$ is a subset of the set of simplex centres.

It is shown in [10, p.19, Cor 5.4] that $\mathrm{Min}_1(\alpha)$ is connected by simplicial paths. However, by Lemma 8.6, a simplicial path in $\mathcal{O}_1$ which only passes through the centres of simplices is a single point. It follows that $\mathrm{Min}_1(\alpha)$ is a single point. □

Corollary 8.8. *Let $\alpha \in Out(G, \mathcal{G})$ be irreducible and displacement 1. Then $Min(\alpha) = Fix(\alpha)$ is a single line.*

Proof. Follows directly from Theorem 8.7 and the fact that
$\mathrm{Min}(\alpha) = \{\mu T \in \mathcal{O} \mid T \in \mathrm{Min}_1(\alpha), \mu > 0\}$. □

There exists a space similar to $\mathcal{O}_1$ called the projectivized space $\mathcal{PO}$, where instead of taking the covolume one subspace of $\mathcal{O}$, one takes a quotient space of $\mathcal{O}$ by identifying all G-trees in the sets $\{(T, \lambda d_T, \cdot) \mid \lambda \in \mathbb{R}\}$ for each $T \in \mathcal{O}$.

When choosing a G-tree to represent a point in $\mathcal{PO}$, one usually takes the unique covolume one G-tree. In this way, we can construct a natural bijection between $\mathcal{O}_1$ and $\mathcal{PO}$.

The displacement λ_α of an automorphism is invariant under rescaling of the metrics d_T, hence we can define the minimally displaced set $\mathrm{Min}_{\mathcal{P}}(\alpha)$ in $\mathcal{PO}$. The set $\mathrm{Min}_1(\alpha)$ is a set of representatives for $\mathrm{Min}_{\mathcal{P}}(\alpha)$, thus Theorem 8.7 shows that $\mathrm{Min}_{\mathcal{P}}(\alpha)$ is also a single point.

References

[1] Mladen Bestvina, Mark Feighn, and Michael Handel. The Tits alternative for Out(F_n). I. Dynamics of exponentially-growing automorphisms. *Ann. of Math. (2)*, 151(2):517–623, 2000.

[2] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *The Annals of Mathematics*, 135(1):1–51, jan 1992.

[3] Daniel E. Cohen. *Combinatorial group theory: a topological approach*, volume 14 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1989.

[4] Marc Culler and John W. Morgan. Group actions on **R**-trees. *Proceedings of the London Mathematical Society*, s3-55(3):571–604, nov 1987.

[5] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones Mathematicae*, 84(1):91–119, feb 1986.

[6] Warren Dicks and Enric Ventura. Irreducible automorphisms of growth rate one. *Journal of Pure and Applied Algebra*, 88(1-3):51–62, 1993.

[7] Max Forester. Deformation and rigidity of simplicial group actions on trees. *Geometry and Topology*, 6:219–267, 2002.

[8] Stefano Francaviglia and Armando Martino. Stretching factors, metrics and train tracks for free products. *Illinois Journal of Mathematics*, 59(4):859–899, 2015.

[9] Stefano Francaviglia and Armando Martino. Displacements of automorphisms of free groups I: Displacement functions, minpoints and train tracks. *Transactions of the American Mathematical Society*, 374(5):3215–3264, 2021.

[10] Stefano Francaviglia and Armando Martino. Displacements of automorphisms of free groups II: Connectivity of level sets and decision problems. *Transactions of the American Mathematical Society*, 375(4):2511–2551, 2022.

[11] Stefano Francaviglia, Armando Martino, and Dionysios Syrigos. The minimally displaced set of an irreducible automorphism of F_N is co-compact. *Archiv der Mathematik*, 116(4):369–383, 2021.

Paper 2: **Growth and displacement of free product automorphisms**

Matthew Collins

ABSTRACT.
It is well known for an *irreducible* free group automorphism that its growth rate is equal to the minimal Lipschitz displacement of its action on Culler-Vogtmann space. This follows as a consequence of the existence of train track representatives for the automorphism.
We extend this result to the general - possibly reducible - case as well as to the free product situation where growth is replaced by 'relative growth'.

1 Introduction

This paper was prompted by a question which arose during one of the author's previous papers [6]. That work was a generalisation of the results of [7] from free groups to free products, and focused on the deformation space $\mathcal{O}(G, \mathcal{G})$ - a space of G-trees which can be thought of as a generalisation of Culler-Vogtmann space. In [7], the authors focused on irreducible, *growth rate 1* automorphisms of free groups, where the growth rate was defined with respect to word length. However, when generalising their results, we ultimately ended up using irreducible, *displacement 1* automorphisms of free products - the displacement being a value used to describe the action of an automorphism on $\mathcal{O}(G, \mathcal{G})$ (described below).

Our new results held regardless of this distinction, and the link between the growth rate and the displacement of an irreducible free group automorphism is reasonably intuitive given a solid understanding of Culler-Vogtmann space and train track maps. The goal of this paper is to turn that intuition into solid proof, and to extend that proof to a larger class of automorphisms by answering the following:

Questions.

(1) *In [7], the growth rate was defined for free groups. What is the "correct" generalisation of this definition to free products?*

(2) *When we have found this generalisation, does the growth rate of an* irreducible *free product automorphism equal its displacement in outer space?*

(3) *More generally, does the growth rate of* any *free product automorphism equal its displacement in outer space?*

(1)

First we address the definition of the growth rate. In the case of free groups, Dicks & Ventura [7] attribute the following definition of the growth rate in arbitrary groups to Thurston:

Definition 3.7. Let α be an automorphism of a finitely generated group G. Let E be a finite generating set of G, and let l_E denote the conjugacy length in the alphabet E. Then for $g \in G$ we define the *growth rate of α with respect to (the conjugacy class of) g* as

$$\mathrm{GR}(\alpha, l_E, g) = \limsup_{k\to\infty} \sqrt[k]{l_E(g\alpha^k)}$$

The *growth rate of α* is then

$$\mathrm{GR}(\alpha, l_E) = \sup\{\mathrm{GR}(\alpha, l_E, g) \mid g \in G\}$$

Using this definition, it can be fairly swiftly proved that the growth rate of an irreducible automorphism of a free group is equal to its displacement in Culler-Vogtmann space. We will say more about this proof when we talk about Question (2), but when we say that we are searching for the "correct" definition of growth rate in free products, it is the generalisation of this proof we have in mind.

We elect to use *relative generating sets*, which was inspired by Osin's paper on relatively hyperbolic groups [10].

Definition 3.9. We say that $E \subseteq G$ is a *relative generating set of G with respect to $\mathcal{G}$* if G is generated by the set

$$\left(\bigcup_{i=1}^{k} G_i\right) \cup E,$$

where $G = G_1 * \ldots * G_k * F_r$ is a free product decomposition corresponding to $\mathcal{G}$.

From this we can essentially follow the definition of the growth rate in free groups by defining *relative conjugacy length*, *relative Lipschitz equivalence*, and finally the *relative growth rate*.

Definition 3.17. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, and let l_E be a relative conjugacy length function. Then for $g \in G$ we define the *relative growth rate of α with respect to (the conjugacy class of) g* as

$$\mathrm{GR}_{\mathcal{G}}(\alpha, l_E, g) = \limsup_{k\to\infty} \sqrt[k]{l(g\alpha^k)}$$

The *relative growth rate of* α is then

$$\mathrm{GR}_{\mathcal{G}}(\alpha, l_E) = \sup\{\mathrm{GR}(\alpha, l_E, g) \mid g \in G\}$$

Remark 1.1. If we consider the case where G is a free group and $\mathcal{G}$ is the trivial free factor system (that is to say, we take the free product decomposition of G consisting of a single free factor - G itself), then all of these "relative" definitions restrict to their original counterparts, as they should.

(2)

Questions (2) and (3) are both true. In fact, (3) immediately implies (2), but the case of irreducible automorphisms is different enough to deserve separate consideration.

A free product decomposition of a group G determines a set of subgroups $\mathcal{G}$ called a free-factor system. We write $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$ to denote the space of equivalence classes of minimal, cocompact, edge-free G-trees whose set of elliptic subgroups is $\mathcal{G}$. Such a space is referred to as *outer space*. $\mathrm{Out}(G, \mathcal{G})$ - the group of outer automorphisms which preserve $\mathcal{G}$ - acts on outer space by "twisting" the action, and studying the effect of this twist gives us information about the original automorphisms.

For any two G-trees $T, S \in \mathcal{O}$, we write $\Lambda_R(T, S)$ to denote the asymmetric Lipschitz distance, or stretching factor, between them. For an automorphism $\alpha \in \mathrm{Out}(G, \mathcal{G})$, one can define the *displacement* of α as $\lambda_\alpha = \inf\{\Lambda(T, \alpha T) \mid T \in \mathcal{O}\}$. The *minimally displaced set* of α, $\mathrm{Min}(\alpha)$, is the set of G-trees T in outer space which realise this infimum.

It is well known for an irreducible *free group* automorphism that its growth rate is equal to the minimal Lipschitz displacement of its action on Culler-Vogtmann space. This follows as a consequence of the existence of train track representatives for the automorphism. It can be shown that the required properties of train track representatives in free groups also hold in free products, so the same method can be used.

More specifically, [8] proves that:

- The minimally displaced set $\mathrm{Min}(\alpha)$ is equal to the set of trees in $\mathcal{O}$ which support optimal train track maps $f : T \to \alpha T$.
- When α is irreducible, $\mathrm{Min}(\alpha)$ is non-empty.

From here the proof of (2) follows. We give a complete proof of this irreducible case in Appendix 2.

(3)

If we drop the irreducibility condition, a problem arises which prevents us from copying the train track proof outright: we cannot guarantee that $\mathrm{Min}(\alpha)$ will be non-empty, and hence we cannot guarantee the existence of an optimal train track map. Happily, however, we can guarantee the existence of a weaker set of maps known as *relative train tracks*.

A topological representative $f : T \mapsto \alpha T$ has an associated *transition matrix* $M = (m_{ij})$, where m_{ij} is the number of times the f-image of the j-th edge-orbit crosses the i-th edge-orbit in either direction. By relabelling edges appropriately, it is always possible to write M in block upper triangular form:

$$M = \begin{pmatrix} M_1 & ? & ? & ? \\ 0 & M_2 & ? & ? \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & M_n \end{pmatrix},$$

where the matrices $M_1, \ldots, M_n$ are either zero matrices or irreducible matrices.

Writing M in this form determines a partition of the edges of T: The *rth stratum* H_r of T is the subgraph of T given by closure of the union of the edge orbits corresponding to the rows/columns in M_r. This also determines a filtration $\varnothing = T_0 \subset \ldots \subset T_n = T$ of T, where $T_r = \bigcup_{i \leqslant r} H_r$. Observe that each T_i is f-invariant, but the H_i are not, in general. We say an edge path γ in T_r is *r-legal* if no component of $\gamma \cap H_r$ contains an illegal turn.

Definition 4.8. [Relative train track] Let $T \in \mathcal{O}(G, \mathcal{G})$, let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, and let $f : T \to \alpha T$ be a simplicial topological representative for α. Use this map to divide T into strata as described above. We say that f is a *relative train track map* if the following conditions hold:

(1) *f preserves r-germs*: For every edge $e \in H_r$, the path $f(e)$ begins and ends with edges in H_r.

(2) *f is injective on r-connecting paths*: For each nontrivial path $\gamma \in T_{r-1}$ joining points in $H_r \cap T_{r-1}$, the homotopy class $[f(\gamma)]$ is nontrivial.

(3) *f is r-legal*: If a path γ is r-legal, then $f(\gamma)$ is r-legal.

Theorem 4.9. *[5, Thm 2.12] For any automorphism $\alpha \in Out(G, \mathcal{G})$, there exists a relative train track map $f : T \to \alpha T$ on some $T \in \mathcal{O}$.*

We observe that M_r is the transition matrix of H_r, and each of these submatrices will have its own PF-eigenvalue μ_r. It can be shown that, even though a relative train track map will not, in general, satisfy the property of train track maps used in Question (2), it *will* satisfy a similar property on each stratum of T using these μ_r. This gives us the tools we require to prove that question (3) is true:

Theorem 5.1. *Let $\alpha \in Out(G, \mathcal{G})$. Then the following are equal:*

- *The relative growth rate of α, $GR_{\mathcal{G}}(\alpha)$.*
- *The largest PF-eigenvalue μ_R of any relative train track map $f : T \to \alpha T$, $T \in \mathcal{O}(G, \mathcal{G})$.*
- *The displacement λ_α of α in $\mathcal{O}$.*

We prove this by proving three inequalites:

A: $\mu_R \leqslant \mathrm{GR}_{\mathcal{G}}(\alpha)$,

B: $\mathrm{GR}_{\mathcal{G}}(\alpha) \leqslant \lambda_\alpha$,

C: $\lambda_\alpha \leqslant \mu_R$.

A follows from existing properties of relative train tracks. **B** can be proved explicitly using the definition of the right stretching factor. **C** requires more thought.

Recall the definition of the displacement: $\lambda_\alpha := \inf_{S \in \mathcal{O}} \Lambda_R(S, \alpha S)$. Ideally we would prove inequality **C** by emulating the proof of question (2) and finding a G-tree in $\mathcal{O}$ whose right stretching factor is exactly μ_R. However, unless α is irreducible, this is not always possible. Thus we instead find a *sequence* of G-trees whose right stretching factors *tend* towards μ_R.

The lengths of edges in T are determined by the PF-eigenvectors, but these are only determined up to scalar multiplication, so we are free to rescale the edges in each stratum by a constant of our choosing. We choose to rescale each H_r by N^r. As N tends to infinity, we observe that the growth in the stratum with the largest PF-eigenvalue becomes greater than that of all other strata. From this the result follows.

2 Bass-Serre Theory

2.1 G-trees

Definition 2.1. A *simplicial tree* is a non-empty, 1-dimensional simplicial complex in which every two points are joined by a unique arc. We call the 1-simplices *edges*, and the 0-simplices *vertices*.

A *metric simplicial tree* is a simplicial tree T together with a metric d_T such that the set of vertices is discrete in the topology induced by d_T.

Remark 2.2. The discreteness condition above is equivalent to saying that the lengths of the edges incident to a given vertex are bounded below - that is to say, $\forall v \in V, \exists C > 0$ such that for all vertices w adjacent to v, $d_T(v, w) \geqslant C$.

Let G be a group.

Definition 2.3. A *G-tree* is a triple $(T, d_T, \cdot)$, where T is a metric simplicial tree, d_T is the metric on T, and $\cdot$ is an isometric group action $T \times G \to T, (x, g) \mapsto x \cdot g$. For the sake of brevity of notation, if the specific metric and action are not required, we shall simply denote the triple $(T, d_T, \cdot)$ by T.

We say that two G-trees are *equivalent* if there exists an equivariant isometry between them.

We say that T is *minimal* if it does not contain a G-invariant subtree.

We say that T is *edge-free* if every edge has trivial stabiliser.

Remark 2.4. We say that an action on a G-tree T is *without inversions* if no edge is sent to its inverse by an element of G. If T does contain an inverted edge, then placing a new vertex at the midpoint will essentially remove this inversion. Furthemore, since the two trees before and after this operation are equivalent in the above sense, adding this vertex does not affect any relevant properties of T. Thus, to simplify our calculations, we shall assume that all our G-trees are without inversions.

Definition 2.5. Let $(T, d_T, \cdot)$ be a G-tree.

We say that an element $g \in G$ is *elliptic* with respect to T if $x \cdot g = x$ for some point $x \in G$. If g is not elliptic, we say it is *hyperbolic*.

We say that a subgroup $H \leqslant G$ is *elliptic* with respect to T if $x \cdot H = x$ for some point $x \in G$. If H is not elliptic, we say it is *hyperbolic*.

An elliptic subgroup will consist entirely of elliptic elements, but the converse is not necessarily true.

Definition 2.6. We write $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$ to denote the space of equivalence classes of minimal, cocompact, edge-free G-trees whose set of elliptic subgroups is $\mathcal{G}$. Such a space is referred to as *outer space*.

2.2 Graphs of groups

Definition 2.7. A *(Serre) graph* X consists of the following:

- a vertex set $V = V(X)$,

- an edge set $E = E(X)$,
- an initial vertex map $\iota : E \to V$,
- an edge reversal map $E \to E, e \mapsto \bar{e}$ such that $e \neq \bar{e}$ and $\bar{\bar{e}} = e$.

We call $\iota(\bar{e})$ the *terminal vertex* of e, and denote it $\tau(e)$.

Definition 2.8. [4, p.198] A *graph of groups* Γ consists of the following:

(i) a connected graph X,

(ii) a group G_v for each vertex v of X, and a group G_e for each edge e of X such that $G_{\bar{e}} = G_e$,

(iii) for each edge e of X, a monomorphism $\rho_e : G_e \to G_{\tau e}$.

If X has a metric, we say that Γ is a *metric graph of groups*.

Notation. We may write $V(\Gamma)$ and $E(\Gamma)$ to denote the sets $V(X)$ and $E(X)$ respectively.

Remark 2.9. All of the graphs of groups we shall be using in this paper will have trivial edge groups (and hence trivial monomorphisms ρ_e). Thus, for the sake of simplicity, we shall restrict our exposition on Bass-Serre Theory to graphs of groups with this property.

Let Γ be a graph of groups with trivial edge groups. Then the *path group* $\pi(\Gamma)$ is defined by

$$\pi(\Gamma) = \frac{\left(*_{v \in V(\Gamma)} G_v \right) * F(E(\Gamma))}{N},$$

where N is the normal closure of the set $\{e\bar{e} \mid e \in E(\Gamma)\}$.

A *path* in Γ is a sequence $g_0 e_1 g_1 \dots g_{n-1} e_n g_n$ where $g_{i-1} \in G_{\iota e_i}$ and $g_i \in G_{\tau e_i}$ for all i (so $e_1 \dots e_n$ is an edge path in the graph). If any $g_i = 1$, then we omit it from the notation.

We say a path is *reduced* if it does not contain any subpaths of the form $e\bar{e}$. We can think of $\pi(\Gamma)$ as the group of reduced paths.

Definition 2.10. Choose a base point $x \in V(\Gamma)$. The *fundamental group* $\pi_1(\Gamma, x)$ of Γ at x is the subgroup of $\pi(\Gamma)$ consisting of the reduced paths which start and end at x.

The isomorphism class of $\pi_1(\Gamma, x)$ does not depend on our choice of x. Thus we shall usually omit it from the notation, and simply write $\pi_1(\Gamma)$.

If X denotes the underlying graph of Γ, then it can be shown that

$$\pi_1(\Gamma) \cong \left(\ast_{v \in V(\Gamma)} G_v \right) * F_r,$$

where $F_r \cong \pi_1(X)$ is a free group of rank r.

Definition 2.11. Let G be a group, and let Γ be a metric graph of groups. A *marking* on Γ is an isomorphism $\phi : G \to \pi_1(\Gamma)$. The pair (Γ, ϕ) is called a *marked graph of groups*.

2.3 The Fundamental Theorem

Bass-Serre theory describes a process by which one may construct a marked graph of groups from a G-tree and, conversely, a G-tree from a marked graph of groups. The Fundamental Theorem of Bass-Serre Theory states that these two constructions are mutually inverse, up to isomorphism of the structures involved.

The details of Bass-Serre Theory have been thoroughly explored in the literature ([2] [4]), so we shall simply give a brief description of the two constructions.

Definition 2.12 (Bass-Serre Tree)**.** [1, p.7]

Let (Γ, ϕ) be a marked graph of groups with trivial edge groups and with marking $\phi : G \to \pi_1(\Gamma, x)$. We then define a graph T called the *universal cover*, or *Bass-Serre tree*, of Γ as follows:

- The vertex set $V(T)$ is the set of 'cosets' $G_v\gamma$, where $\gamma \in \pi(\Gamma)$ is a path from the vertex v to our base point x.
- Two vertices $G_{v_1}\gamma_1, G_{v_2}\gamma_2 \in V(T)$ are joined by an edge(-pair) in T if the vertices v_1 and v_2 are joined by an edge pair $(e, \bar{e})$ in X, and $\gamma_1 = eg_{v_2}\gamma_2$ or $\gamma_2 = \bar{e}g_{v_1}\gamma_1$.

It can be shown that this graph T is always a tree.

If $g \in G$, then $\phi(g) \in \pi_1(\Gamma, x)$ is a loop in Γ - that is to say, a path from x to x. Thus we can define a right action $\cdot$ of G on T as follows:

$$\forall G_v\gamma \in V(T), \qquad G_v\gamma \cdot \phi(g) := G_v(\gamma\phi(g))$$

This action respects adjacency, sending edges to edges. Together with this action, the Bass-Serre tree is an edge-free G-tree.

We say two marked graphs of groups are *equivalent* if their universal covers are equivalent as G-trees.

Remark 2.13. Observe that the Bass-Serre tree as we have defined it here is a Serre graph, whereas the definition of G-trees we have given views them as simplicial structures. This distinction ultimately matters little - we can reconcile the two viewpoints by thinking of each pair $(e, \bar{e})$ as denoting two orientations of a 1-simplex, rather than being individual edges.

Definition 2.14 (Quotient Graph of Groups)**.** Let T be an edge-free G-tree. Then we define the *quotient graph of groups* of T as follows:

- the underlying Serre graph is the quotient graph T/G,
- all edge groups are trivial,
- consider a connected fundamental domain in T. This will contain exactly one vertex from each orbit. We assign the stabilizers of these vertices to be the vertex groups of the corresponding vertices in T/G.

It is given by the Fundamental Theorem of Bass-Serre Theory that the fundamental group of this graph of groups is isomorphic to G. The action of G on T determines this isomorphism, giving us a marking.

2.4 Free factor systems

There exists a relation between G-trees and free product decompositions of G, which can be phrased as follows:

Consider the outer space $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$, and recall that trees in this space are cocompact. A G-tree is cocompact if and only if it has a finite number of edge orbits and vertex orbits. It follows that the quotient graph of groups Γ has a finite number of vertex groups, and hence we can write $G = \pi_1(\Gamma) \cong G_1 * \ldots * G_k * F_r$, where $G_1, \ldots, G_k$ are the vertex groups. (Since we have a choice of vertex groups when constructing Γ, this free product decomposition is only unique up to conjugates of the free factors).

Conversely, suppose we are given a free product $G = G_1 * \ldots * G_k * F_r$. Then one can easily construct graphs of groups with vertex groups $G_1, \ldots, G_k$ and fundamental group G (for example, see Figure 2.1). Taking the Bass-Serre trees will give us an outer space $\mathcal{O} = \mathcal{O}(G, \mathcal{G})$ of G-trees, where the elliptic subgroups are
$\mathcal{G} = \langle H \leqslant G_i^g \mid g \in G, i = 1, \ldots, k\rangle$ - the subgroups of the conjugates of the free factors.

To summarize, a space $\mathcal{O}(G, \mathcal{G})$ determines a family of free product decompositions of G, and a free product decomposition of G determines a space $\mathcal{O}(G, \mathcal{G})$. It follows from the Fundamental Theorem of Bass-Serre Theory that these two processes are mutually inverse (up to conjugacy of the free factors). Thus we shall refer to the set $\mathcal{G}$ as a *free factor system*.

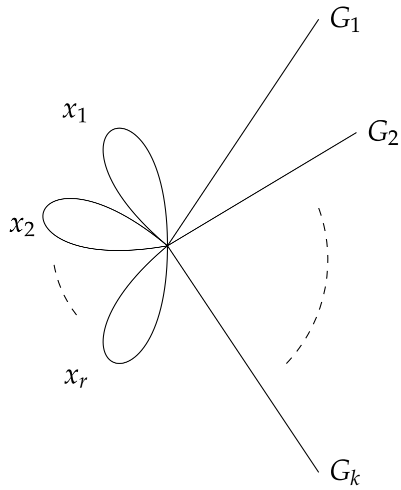

FIGURE 2.1: A graph of groups with fundamental group $G_1 * \ldots * G_k * \langle x_1, \ldots, x_r \rangle$

Definition 2.15. Let $\mathcal{G}$ be a free factor system for a group G, and let $g \in G$. We say that g is *$\mathcal{G}$-elliptic* if it lies in a subgroup in $\mathcal{G}$. Otherwise, we say it is *$\mathcal{G}$-hyperbolic*. We shall write $\mathrm{Hyp}(\mathcal{G})$ to denote the set of all $\mathcal{G}$-hyperbolic elements.

Remark 2.16. The notions of elliptic and hyperbolic given here align with those in Definition 2.5.

3 Length functions and growth

Definition 3.1. By a *length function* on a set X, we mean a map $l : X \to \mathbb{R}$ taking non-negative values.

3.1 Displacement in Outer Space

Definition 3.2. Let $g \in G$, and let $T \in \mathcal{O}(G, \mathcal{G})$. Then we write $l_T(g)$ to denote the *translation length* of g in T, given by

$$l_T(g) = \inf_{x \in T} \{d_T(x, x \cdot g)\}$$

Remark 3.3. $l_T(g^h) = l_T(g)$ for all $g, h \in G$. Thus we shall think of l_T as a length function on the conjugacy classes of single elements of G.

It can be shown that this infimum is achieved by some $x \in T$. If g is elliptic, then $l_T(g) = 0$ by definition.

If g is hyperbolic, then $l_T(g) > 0$. The set of points $x \in T$ which realise $l_T(g)$ will form a line in T called the *hyperbolic axis of g*. Points on the axis will be translated along the axis a distance of $l_T(g)$ by g.

Definition 3.4. [8, p.8] Let $T, S \in \mathcal{O}(G, \mathcal{G})$. We define the *right stretching factor* from T to S as

$$\Lambda_R(T, S) := \sup_{g \in \mathrm{Hyp}(\mathcal{G})} \frac{l_S(g)}{l_T(g)},$$

Remark 3.5. The right stretching factor is bounded above. In fact, it is realised by some $g \in \mathrm{Hyp}(\mathcal{G})$. This follows from Corollary 6.8 and Theorem 6.11 of [8].

Definition 3.6. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$. Then we define the *displacement* of α to be

$$l_\alpha := \inf_{T \in \mathcal{O}} \Lambda_R(T, \alpha T).$$

3.2 Growth rate

Notation. In this paper, the automophisms of a group G will act on G on the right.

Dicks & Ventura [7] attribute the following definition to Thurston.

Definition 3.7. Let α be an automorphism of a finitely generated group G. Let E be a finite generating set of G, and let l_E denote the conjugacy length in the alphabet E. Then for $g \in G$ we define the *growth rate of α with respect to (the conjugacy class of) g* as

$$\mathrm{GR}(\alpha, l_E, g) = \limsup_{k \to \infty} \sqrt[k]{l_E(g\alpha^k)}$$

The *growth rate of α* is then

$$\mathrm{GR}(\alpha, l_E) = \sup\{\mathrm{GR}(\alpha, l_E, g) \mid g \in G\}$$

Remark 3.8. It can be shown that any two finite generating sets of a finitely generated group G give Lipschitz equivalent conjugacy length functions, and from this it can be further shown that Definition 3.7 does not depend on the choice of E.

We shall now generalise this definition from finitely generated groups to finite free products - that is to say, free products of the form $G = G_1 * \ldots * G_k * F_r$, where k and r are finite, and F_r denotes a free group of rank r.

3.3 Relative length functions

Definition 3.9. We say that $E \subseteq G$ is a *relative generating set of G with respect to $\mathcal{G}$* if G is generated by the set

$$\left(\bigcup_{i=1}^{k} G_i \right) \cup E,$$

where $G = G_1 * \ldots * G_k * F_r$ is a free product decomposition corresponding to $\mathcal{G}$.

Let $\hat{G} := \bigsqcup_{i=1}^{k}(G_i \backslash \{1\})$. Then we observe that G is generated by the set $E \cup \hat{G}$.

Definition 3.10. Let E be a relative generating set of G with respect to $\mathcal{G}$. To each element $g \in G$, we assign its *relative length with respect to* $\mathcal{G}$, $|g|_{E \cup \hat{G}}$, to be the length of a shortest word in the alphabet $E \cup \hat{G}$ which represents g in G.

To each element $g \in G$, we assign the *relative conjugacy length with respect to* $\mathcal{G}$, $l_E(g)$, given by

$$l_E(g) = \min\{|h|_{E \cup \hat{G}} \mid h \in G \text{ is conjugate to } g\},$$

It is clear that this length depends upon the choice of E. However, when E is finite, the effect this choice has on the length is bounded in the following sense:

Definition 3.11. We say that two length functions l_1, l_2 on the same set X are *Lipschitz equivalent* if there exist constants $C, D > 0$ such that, for all $x \in X$,
$Cl_2(x) \leqslant l_1(x) \leqslant Dl_2(x)$. We write $l_1 \sim l_2$.

Let $G = G_1 * \ldots * G_k * F_r$. We say that two length functions l_1, l_2 on the conjugacy classes of G are *relatively Lipschitz equivalent* if they are Lipschitz equivalent when restricted to $\mathrm{Hyp}(\mathcal{G})$. We write $l_1 \sim_{\mathcal{G}} l_2$.

Proposition 3.12. *(Adapted from [10, p.13]) Let $G = G_1 * \ldots * G_k * F_r$ be a free product. Suppose that E_1 and E_2 are two finite relative generating sets of G with respect to $\{G_1, \ldots, G_k\}$. Then the corresponding length functions l_{E_1} and l_{E_2} are Lipschitz equivalent (and hence relatively Lipschitz equivalent).*

Theorem 3.13. *Let $\alpha \in Out(G, \mathcal{G})$, let E be a finite relative generating set for G, and let $T \in \mathcal{O}$. Then $l_E \sim_{\mathcal{G}} l_T$.*

Proof. There exists a free product decomposition $G = G_1 * \ldots * G_k * F_r$ corresponding to E. Let $E' = \{x_1, \ldots, x_r\}$ be a basis for F_r. Then E' is also a finite $\mathcal{G}$-relative generating set for G. Let Γ' denote the graph of groups depicted in Figure 2.1. The universal cover T' of Γ' lies in $\mathcal{O}$.

To define a metric on Γ' (and hence T'), we assign the loops on the left length 1, and the remaining edges on the right length $\frac{1}{2}$. By assigning these edge lengths, we ensure that $l_{E'}(g) = l_{T'}(g)$ for all hyperbolic $g \in G$, hence $l_{E'} \sim_{\mathcal{G}} l_{T'}$.

There exist Lipschitz continuous maps between any two G-trees in $\mathcal{O}$ [8, Lem 4.2]. This is sufficient to prove that $l_{T'} \sim_{\mathcal{G}} l_T$. Thus

$$l_E \overset{\text{Prop 3.12}}{\sim_{\mathcal{G}}} l_{E'} \sim_{\mathcal{G}} l_{T'} \sim_{\mathcal{G}} l_T$$

□

3.4 Outer Automorphisms

Definition 3.14. Let G be a group. The *outer automorphism group* of G is the quotient $\mathrm{Out}(G) := \mathrm{Aut}(G)/\mathrm{Inn}(G)$, where $\mathrm{Aut}(G)$ and $\mathrm{Inn}(G)$ denote the automorphism group and inner automorphism group respectively.

A free factor system $\mathcal{G}$ of a group G is a set of subgroups which is closed under conjugation. It follows that $\mathcal{G} = \mathcal{G}\phi = \{H\phi \mid H \in \mathcal{G}\}$ for all inner automorphisms ϕ. Therefore, when considering the effect of an automorphism on $\mathcal{G}$, it suffices to consider its outer automorphism class.

Notation. Since we need only consider outer automorphism classes, we use some slightly non-standard notation for the sake of brevity: We will not write out the usual square brackets $[\alpha]$ to denote the outer automorphism class of $\alpha \in \mathrm{Aut}(G)$. We will simply write $\alpha \in \mathrm{Out}(G)$.

Definition 3.15. Let $\mathcal{G}$ be a free factor system for a group G, and let $\alpha \in \mathrm{Out}(G)$. We say that $\mathcal{G}$ is α-invariant if $\mathcal{G} = \mathcal{G}\alpha$. We write $\mathrm{Out}(G, \mathcal{G})$ to denote the subgroup of $\mathrm{Out}(G)$ consisting of the elements α such that $\mathcal{G}$ is α-invariant.

3.5 The relative growth rate

Definition 3.16. Let $\mathcal{G}$ be a free factor system for a group G. We say a *$\mathcal{G}$-bounded* length function is a length function l on the conjugacy classes of single elements of G which satisfies the following:

- $\exists C > 0$ such that $\forall g \notin \mathrm{Hyp}(\mathcal{G})$, $l(g) \leqslant C$,
- $\exists \varepsilon > 0$ such that $\forall g \in \mathrm{Hyp}(\mathcal{G})$, $\varepsilon \leqslant l(g)$,
- $\forall g \in \mathrm{Hyp}(\mathcal{G})$, for all non-zero integers n, $l(g^n) = |n| l(g)$.

Definition 3.17. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, and let l_E be a relative conjugacy length function. Then for $g \in G$ we define the *relative growth rate of α with respect to (the conjugacy class of) g* as

$$\mathrm{GR}_{\mathcal{G}}(\alpha, l_E, g) = \limsup_{k \to \infty} \sqrt[k]{l(g\alpha^k)}$$

The *relative growth rate of α* is then

$$\mathrm{GR}_{\mathcal{G}}(\alpha, l_E) = \sup\{\mathrm{GR}(\alpha, l_E, g) \mid g \in G\}$$

Lemma 3.18. *Let $G, \mathcal{G}, \alpha$ and l be as in Definition 3.17. Then the relative growth rate of α is determined by the hyperbolic elements of G - that is to say,*

$$GR_{\mathcal{G}}(\alpha, l_E) = \sup\{GR(\alpha, l_E, g) \mid g \in Hyp(\mathcal{G})\}.$$

Proof. To prove this, it suffices to show there exists a hyperbolic element of G whose relative growth rate is greater than or equal to the relative growth rate of every elliptic element.

Let $g \in G$ be an elliptic element. $\mathcal{G}$ is α-invariant, hence $g\alpha^k$ is also elliptic. Furthermore, since l is $\mathcal{G}$-bounded, $\exists C > 0$ such that $l(g) \leqslant C$, so

$$\begin{aligned} \mathrm{GR}_{\mathcal{G}}(\alpha, l, g) &= \limsup_{k\to\infty} \sqrt[k]{l(g\alpha^k)} \\ &\leqslant \limsup_{k\to\infty} \sqrt[k]{C} && \text{(by properties of lim sup)} \\ &= 1 \end{aligned}$$

Thus we wish to find a hyperbolic element whose relative growth rate is at least 1. Since l is $\mathcal{G}$-bounded, $\exists \varepsilon > 0$ such that $\forall h \in \mathrm{Hyp}(\mathcal{G})$, $\varepsilon \leqslant l(g)$. Let $h \in \mathrm{Hyp}(\mathcal{G})$, and take $N \geqslant \frac{1}{\varepsilon}$. Then

$$l(h^N\alpha^k) = l((h\alpha^k)^N) = |N| \cdot l(g\alpha^k) \geqslant \frac{1}{\varepsilon} \cdot \varepsilon = 1.$$

h^N is also hyperbolic, hence this completes the proof. □

Note that we have only defined the growth rate for $\mathcal{G}$-bounded length functions. We can show that both of the length functions we have seen in this paper are $\mathcal{G}$-bounded:

Lemma 3.19. *Let $G = G_1 * \ldots * G_k * F_r$ be a free product with corresponding free factor system $\mathcal{G}$. Let E be a relative generating set with respect to $\mathcal{G}$, and let l_E be the corresponding relative conjugacy length. Then l_E is a $\mathcal{G}$-bounded length function.*

Proof. Let $g \in G$. If g is elliptic, then $l_E(g) = 1$, so take $C = 1$.

The only element of G with relative conjugacy length 0 is the identity element, which is elliptic. Thus we may take $\epsilon = 1$.

Let g be hyperbolic, and let $x_1 \ldots x_k$ be a shortest word in E representing a conjugate of g. If $x_1 = x_k^{-1}$, then conjugating by x_1^{-1} would shorten the word, resulting in a contradiction. Thus we may assume that this is not the case, and hence

$$l_E(g^n) = l_E((x_1 \ldots x_k)^n) = |n|k = |n|l_E(g).$$

□

Lemma 3.20. *Let $T \in \mathcal{O}(G, \mathcal{G})$. Then l_T is a $\mathcal{G}$-bounded length function.*

Proof. Elliptic elements are by definition elements which fix a point in T. Thus the translation length of elliptic elements is bounded above.

Hyperbolic elements g translate points along a hyperbolic axis - in particular, since the action is isometric, the points are translated by a whole number of edges. T is cocompact, hence it contains finitely many edge orbits - thus we may take ϵ to be the length of the shortest edge.

Additionally, points on the axis of g are translated along it by a distance of $l_T(g)$. g^n has the same hyperbolic axis, and it translates points on the axis a distance of
$l_T(g^n) = |n| l_T(g)$. □

We will now show that these two length functions give us the same relative growth rate:

Proposition 3.21. *Let l_1, l_2 be $\mathcal{G}$-bounded length functions, and suppose that $l_1 \sim_{\mathcal{G}} l_2$. Then $GR_{\mathcal{G}}(\alpha, l_1) = GR_{\mathcal{G}}(\alpha, l_2)$.*

Proof. By Lemma 3.18, it is sufficient to restrict our attention to the hyperbolic elements of G. Let $h \in \mathrm{Hyp}(\mathcal{G})$. $l_1 \sim_{\mathcal{G}} l_2$, hence $\exists D > 0$ such that $l_1(h) \leqslant D l_2(h)$. Thus

$$\begin{aligned}
\mathrm{GR}(\alpha, l_1, h) &= \limsup_{n\to\infty} \sqrt[n]{l_1(h\alpha^n)} \\
&\leqslant \limsup_{n\to\infty} \sqrt[n]{D}\sqrt[n]{l_2(h\alpha^n)} \\
&\leqslant \left(\limsup_{n\to\infty} \sqrt[n]{D}\right)\left(\limsup_{n\to\infty} \sqrt[n]{l_2(h\alpha^n)}\right) && \text{(by properties of limsup)} \\
&= \limsup_{n\to\infty} \sqrt[n]{l_2(h\alpha^n)} && (\limsup_{n\to\infty} \sqrt[n]{D} = \lim_{n\to\infty} \sqrt[n]{D} = 1) \\
&= \mathrm{GR}(\alpha, l_2, h).
\end{aligned}$$

The reverse inequality can be obtained in the analogous way, and hence
$\mathrm{GR}(\alpha, l_1, h) = \mathrm{GR}(\alpha, l_2, h)$. This holds for all $h \in \mathrm{Hyp}(\mathcal{G})$, hence $\mathrm{GR}(\alpha, l_1) = \mathrm{GR}(\alpha, l_2)$. □

Notation. In a similar manner to the free group case, Proposition 3.21 shows that, up to relative Lipschitz equivalence, the relative growth rate does not depend on our choice of l. Thus, unless the particular length function is required, we shall omit l from the notation, and simply write $\mathrm{GR}_{\mathcal{G}}(\alpha, g)$ and $\mathrm{GR}_{\mathcal{G}}(\alpha)$.

4 Relative Train Tracks and Perron-Frobenius

Definition 4.1. Let $T, S \in \mathcal{O}$. A map $f : T \to S$ is called an *$\mathcal{O}$-map* if it is Lipschitz continuous and G-equivariant. We write $\mathrm{Lip}(f)$ to denote the Lipschitz constant of f.

An $\mathcal{O}$-map $f : T \to S$ is *straight* if it has constant speed on edges - that is to say, the restriction of f to each edge is a linear map.

Remark 4.2. [9, Remark 3.4] $\mathcal{O}$-maps exist between any pair of G-trees in $\mathcal{O}$. Any $\mathcal{O}$-map f can be uniquely straightened - that is to say, there exists a unique $\mathcal{O}$-map $\mathrm{Str}(f)$ which is homotopic relative to vertices to f. We have $\mathrm{Lip}(\mathrm{Str}(f)) \leqslant \mathrm{Lip}(f)$.

Definition 4.3. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, and let $T \in \mathcal{O}(G, \mathcal{G})$. We call a straight $\mathcal{O}$-map $f : T \to \alpha T$ a *topological representative for α*. If f maps vertices to vertices, we say it is *simplicial*.

4.1 Stratification of G-trees for topological representatives

Let $f : T \to \alpha T$ be a simplicial topological representative for $\alpha \in \mathrm{Out}(G, \mathcal{G})$. Being simplicial, f will map edges in T to edge paths. This behaviour determines the associated *transition matrix* $M = (m_{ij})$ of f, where m_{ij} is the number of times the f-image of the j-th edge-orbit crosses the i-th edge-orbit in either direction.

Relabelling the edges (which equates to reordering the rows and columns of the matrix) allows us to write M in block upper triangular form

$$M = \begin{pmatrix} M_1 & ? & ? & ? \\ 0 & M_2 & ? & ? \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & M_n \end{pmatrix},$$

where the matrices $M_1, \dots, M_n$ are either zero matrices or irreducible matrices.

Writing M in this form determines a partition of the edges of T: The *rth stratum H_r* of T is the closure of the union of the edge orbits in T corresponding to the rows/columns in M_r. We observe that M_r is the transition matrix of H_r.

Remark 4.4. Dividing T into strata in this way also gives us a filtration
$\varnothing = T_0 \subset \ldots \subset T_n = T$ of T, where $T_r = \bigcup_{i \leqslant r} H_r$.

Examining M tells us that this filtration is f-invariant. The strata themselves are not; the f-image of an edge in one stratum may intersect the lower strata.

Theorem 4.5. *[Perron-Frobenius] Let A be a non-negative, irreducible, integer-valued square matrix. Then one of its eigenvalues, called the* PF-eigenvalue*, is a positive real number μ*

which is greater than or equal to the absolute value of all other eigenvalues. There is a positive real eigenvector corresponding to μ, unique up to scalar multiplication.

Definition 4.6. Let f and M be as above. We shall write μ_r to denote the PF-eigenvalue of M_r. We say that H_r is a *growing stratum* if M_r is not a zero matrix.

Observe that, since f is non-trivial, there will always exist at least one growing stratum.

4.2 Relative train tracks

Relative train track maps are a particular type of simplicial topological representative which were introduced by Bestvina & Handel in the case of free groups [3], and were generalised to free products by Collins & Turner [5]. Collins and Turner defined their maps on graphs of complexes - graphs with 2-complexes assigned to the vertices. However, this definition can be transferred to G-trees by replacing each 2-complex with its fundamental group to give a graph of groups and then lifting to the Bass-Serre tree.

Definition 4.7. Let $T \in \mathcal{O}$, and let v be a vertex in T. A *turn* at v is a pair of directed edges (e_1, e_2) such that $\tau(e_1) = v = \iota(e_2)$. We say the turn is *degenerate* if $e_2 = \overline{e_1}$.

A simplicial topological representative $f : T \to \alpha T$ induces a map Df on the turns of T: $Df(e_1, e_2)$ is the turn consisting of the first edges in the edge paths $f(e_1), f(e_2)$. A turn is *illegal* with respect to f if its image under some iterate of Df is degenerate. Otherwise, it is *legal*.

We say an edge path γ in T is *legal* if it does not contain any illegal turns.

We say an edge path γ in T_r is *r-legal* if $\gamma \cap H_r$ does not contain any illegal turns.

Definition 4.8. [Relative train track] Let $T \in \mathcal{O}(G, \mathcal{G})$, let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, and let $f : T \to \alpha T$ be a simplicial topological representative for α. Use this map to divide T into strata as described above. We say that f is a *relative train track map* if the following conditions hold:

(1) *f preserves r-germs*: For every edge $e \in H_r$, the path $f(e)$ begins and ends with edges in H_r.

(2) *f is injective on r-connecting paths*: For each nontrivial path $\gamma \in T_{r-1}$ joining points in $H_r \cap T_{r-1}$, the homotopy class $[f(\gamma)]$ is nontrivial.

(3) *f is r-legal*: If a path γ is r-legal, then $f(\gamma)$ is r-legal.

Theorem 4.9. *[5, Thm 2.12] For any automorphism $\alpha \in Out(G, \mathcal{G})$, there exists a relative train track map $f : T \to \alpha T$ on some $T \in \mathcal{O}$.*

Let $f : T \to \alpha T$ be a relative train track map. Then, by definition, f is a simplicial topological representative for α, so we may stratify T as described above. Let M be the transition matrix for f, with submatrices $M_1, \ldots, M_n$. Let $\emptyset = T_0 \subset \ldots \subset T_n = T$ be the corresponding filtration, and let $H_1, \ldots, H_n$ denote the strata.

For each r we shall write μ_r to denote the Perron-Frobenius eigenvalue for M_r. By Theorem 4.5, this corresponds to a positive real *row* eigenvector - the PF-eigenvector - which we shall denote $\mathbf{v}_r$. This eigenvector is unique up to scalar multiplication, but we will be taking $\mathbf{v}_r$ to be the normalised eigenvector. This choice determines *r-lengths* $L_r(e)$ which we can assign to the edges of H_r: we declare the r-length of the ith edge of H_r to b the ith entry in $\mathbf{v}_r$.

If γ is an edge path, then we define its r-length to be

$$L_r(\gamma) = \sum_{e \in \gamma \cap H_r} L_r(e)$$

Notation. We will adapt this notation slightly when considering elements of the group G:

Recall that a hyperbolic element $g \in \mathrm{Hyp}(\mathcal{G})$ corresponds to a hyperbolic axis in T, and that $l_T(g)$ is the distance points on the axis are translated by g. However, we can also think of $l_T(g)$ as the length of a path - specifically the length of a fundamental domian of the axis. We can always choose this fundamental domain such that it consists of whole edges (i.e. it is an edge path).

Let γ_g denote such a fundamental domain. We say that g is r-legal if γ_g is r-legal. This is independant of our choice of γ_g. We will write $L_r(g)$ to denote the r-length of γ_g.

Now, let us give the properties we shall require.

Lemma 4.10. *Let $f : T \to \alpha T$ be a relative train track map on some $T \in \mathcal{O}(G, \mathcal{G})$, and let H_r be a growing stratum in T. Then the following hold:*

(i) Every edge in H_r is r-legal.

(ii) If an edge path γ is r-legal, then $f(\gamma)$ is r-legal.

(iii) There exists an r-legal group element $g \in Hyp(\mathcal{G})$.

(iv) If an edge path γ is r-legal, then $L_r(f(\gamma)) = \mu_r L_r(\gamma)$.

(v) If g is r-legal, then $L_r(g\alpha^k) = \mu_r^k L_r(g)$.

Proof.

(i) Edges do not contain any turns - in particular, they do not contain any illegal turns. Thus they are r-legal.

(ii) This is one of the defining properties of a relative train track map (Definition 4.8 (iii)).

(iii) By (i), edges in H_r are r-legal. Therefore, by (ii), iterating f will give us longer and longer r-legal paths. Since T is cocompact, it contains finitely many edge orbits. Thus we will eventually reach an r-legal path $f^k(e)$ which crosses some edge orbit at least three times. A least two of these edges, which we can denote by ε and $\varepsilon \cdot g$, must point in the same direction - but this can only occur on the axis of g. It follows that $f^k(e)$ must contain some fundamental domain γ_g of this axis. Hence g is r-legal.

(iv) Follows from the definition of a relative train track map [5, p454].

(v) Since f is a topological representative for α, we have that $f^k(g) = g\alpha^k$ for all $k \geqslant 1$. Thus, by property (iv), $L_r(g\alpha) = \mu_r L_r(g)$, and property (ii) allows us to iterate f, giving us $L_r(g\alpha^k) = \mu_r^k L_r(g)$.

□

5 Main Theorem

Theorem 5.1. *Let $\alpha \in Out(G, \mathcal{G})$. Then the following are equal:*

- *The relative growth rate of α, $GR_{\mathcal{G}}(\alpha)$.*
- *The largest PF-eigenvalue μ_R of any relative train track map $f : T \to \alpha T$, $T \in \mathcal{O}(G, \mathcal{G})$.*
- *The displacement λ_α of α in $\mathcal{O}$.*

Proof. The existence of f is assured by Theorem 4.9. We observe that multiple strata in T may have PF-eigenvalue equal to μ_R. Therefore, when we write H_R, we are referring to the highest of these strata - that is to say, we are maximizing the size of T_R. H_R will always be a growing stratum.

We prove this theorem by proving the following three inequalities:

A: $\mu_R \leqslant \mathrm{GR}_{\mathcal{G}}(\alpha)$,

B: $\mathrm{GR}_{\mathcal{G}}(\alpha) \leqslant \lambda_\alpha$,

C: $\lambda_\alpha \leqslant \mu_R$.

A: By Lemma 4.10 (v), there exists an R-legal $g \in \mathrm{Hyp}(\mathcal{G})$ such that $L_R(g\alpha^k) = \mu_R^k L_R(g)$ for all $k \geqslant 1$. Hence,

$$\begin{aligned}
\mathrm{GR}_{\mathcal{G}}(\alpha, g) &= \limsup_{k\to\infty} \sqrt[k]{l_T(g\alpha^k)} \\
&\geqslant \limsup_{k\to\infty} \sqrt[k]{L_R(g\alpha^k)} \\
&= \limsup_{k\to\infty} \sqrt[k]{\mu_R^k L_R(g)} && \text{(by Lemma 4.10 (v))} \\
&= \mu_R
\end{aligned}$$

B: Let $S \in \mathcal{O}$, and for brevity of notation let Λ denote the right stretching factor $\Lambda_R(S, \alpha S)$. Then

$$\begin{aligned}
\Lambda_R(S, \alpha^k S) &\leqslant \Lambda_R(S, \alpha S)\Lambda_R(\alpha S, \alpha^2 S)\dots\Lambda_R(\alpha^{k-1}S, \alpha^k S) \\
&= \Lambda_R(S, \alpha S)^k && \text{(by triangle inequality)} \\
&= \Lambda^k \\
\Rightarrow \frac{l_S(g\alpha^k)}{l_S(g)} &\leqslant \Lambda^k \qquad \forall g \in G && \text{(by definition of } \Lambda_R) \\
\Rightarrow l_S(g\alpha^k) &\leqslant \Lambda^k l_S(g) \qquad \forall g \in G
\end{aligned}$$

$$\Rightarrow \mathrm{GR}_{\mathcal{G}}(\alpha, g) = \limsup_{k\to\infty} \sqrt[k]{l_S(g\alpha^k)} \leqslant \limsup_{k\to\infty} \sqrt[k]{\Lambda^k l_S(g)} = \Lambda \qquad \forall g \in G$$
$$\Rightarrow \mathrm{GR}_{\mathcal{G}}(\alpha) \leqslant \lambda_S$$

This holds for all $S \in \mathcal{O}$. Thus $\mathrm{GR}_{\mathcal{G}}(\alpha) \leqslant \inf_{S \in \mathcal{O}} \Lambda(S, \alpha S) = \lambda_\alpha$.

C: Recall the definition of the displacement: $\lambda_\alpha := \inf_{S\in\mathcal{O}} \Lambda_R(S, \alpha S)$. Ideally we would prove inequality **C** by finding a G-tree in $\mathcal{O}$ whose right stretching factor is exactly μ_R. However, unless α is irreducible, this is not always possible. Thus we shall instead find a *sequence* of G-trees whose right stretching factors *tend* towards μ_R.

Our proof will apply [8, Lemma 4.3], which states that the Lipschitz constant of an $\mathcal{O}$-map is equal to the right stretching factor between its endpoints. Recall that relative train tracks are straight $\mathcal{O}$-maps - that is to say, they have constant speed on edges. Since we only have finitely many edge orbits in T, it follows that $\mathrm{Lip}(f)$ is realised by some edge - i.e. $\mathrm{Lip}(f) = \max_{e\in T} \frac{l_T(f(e))}{l_T(e)}$.

Let H_r be a stratum in T. Recall that the lengths of the edges in T were determined by the Perron-Frobenius eigenvector corresponding to μ_r - or rather, by the one dimensional eigenspace containing this eigenvector. Since eigenvectors are only determined up to scalar multiplication, we are free to rescale the edge lengths in H_r by a positive constant without affecting any of the relevant properties.

For a positive number $N > 0$, let $T_N \in \mathcal{O}$ denote the G-tree acquired from G by rescaling the stratum H_r in T by N^r for every r. f induces a relative train track map on T_N, which we shall denote by f_N.

Since we are now working with multiple trees, we should make some adjustments to our r-length notation for this portion of the proof. L_r shall be used to denote the r-lengths of the original tree T. We introduce the notation $L_{r,N}$ for the r-lengths of the tree T_N.

Let e be an edge in the rth stratum of T_N. Then $l_{T_N}(e) = L_{r,N}(e) = N^r L_r(e)$.

$$\begin{aligned}
\Rightarrow l_{T_N}(f_N(e)) &= \sum_{i=1}^{r} L_{i,N}(f_N(e)) = \sum_{i=1}^{r} N^i L_i(f_N(e)) \\
\Rightarrow \frac{l_{T_N}(f(e))}{l_{T_N}(e)} &= \frac{\sum_{i=1}^{r} N^i L_i(f(e))}{N^r L_r(e)} \\
&= \frac{\sum_{i=1}^{r} N^{i-r} L_i(f(e))}{L_r(e)} \\
&\longrightarrow \frac{L_r(f(e))}{L_r(e)} \text{ as } N \to \infty
\end{aligned}$$

By Lemma 4.10 (i) & (iv), $\dfrac{L_r(f(e))}{L_r(e)} = \dfrac{\mu_r L_r(e)}{L_r(e)} = \mu_r$

$$\Rightarrow \lim_{N\to\infty}(\mathrm{Lip}(f_N)) = \lim_{N\to\infty}\left(\max_{e\in T_N} \frac{l_{T_N}(f(e))}{l_{T_N}(e)}\right) = \mu_R$$

Finally,

$$\lambda_\alpha = \inf_{S\in\mathcal{O}} \Lambda_R(S, \alpha S) \leqslant \lim_{N\to\infty} \Lambda_R(T_N, \alpha T_N) \overset{\text{by [8, Lem 4.3]}}{\leqslant} \lim_{N\to\infty}(\mathrm{Lip}(f_N)) = \mu_R$$

□

6 Appendix 1: A Bounding Function

Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$. As in the previous chapter, we consider a relative train track map $f : T \to \alpha T$ on a G-tree $T \in \mathcal{O}$. The purpose of this appendix is to find a function of k which acts as an upper bound on $l_T(g\alpha^k)$. If g is elliptic, then $l_T(g\alpha^k) = 0$, so we shall restrict our attention to hyperbolic elements.

As before, f gives a stratification of T. We begin by considering the effect α has on a single stratum:

Lemma 6.1. *Let $g \in Hyp(\mathcal{G})$. Then for each stratum H_r,*

$$L_r(g\alpha) \leqslant \sum_{i \geqslant r} A(r,i) L_i(g), \text{ where } A(r,i) = \max_{e \in H_i} \frac{L_r(f(e))}{l(e)}$$

Proof. Recall the notation we introduced earlier in the paper: When we write $L_r(g\alpha)$, we actually mean $L_r(f(\gamma_g))$, where γ_g is an edge path serving as a fundamental domain of the axis of g.

The filtration $\emptyset = T_0 \subset \ldots \subset T_n = T$ determined by f is f-invariant - therefore a point on the path $f(\gamma_g)$ can only lie in H_r if it was mapped from H_r itself or from a higher stratum. Thus we shall split γ_g into pieces $\gamma_g \cap H_i$ for $i \geqslant r$ and consider the effect f has on each piece.

(Note that $\gamma_g \cap H_i$ may not be connected. When we write $L_r(\gamma_g \cap H_i)$, we mean the sum of the r-lengths of its component paths.)

$$L_r(g\alpha) = L_r(f(\gamma_g)) = \sum_{i \geqslant r} L_r(f(\gamma_g \cap H_i))$$

Let $i \geqslant r$. Then

$$\begin{aligned}
L_R(f(\gamma_g \cap H_i)) &= \sum_{\text{edge orbits } e \in H_i} n_e L_R(f(e)) \qquad \text{where } n_e = \text{num. times } \gamma_g \text{ crosses } e \\
&= \sum_{\text{edge orbits } e \in H_i} n_e \frac{L_R(f(e))}{l(e)} l(e) \\
&\leqslant A(r,i) \sum_{\text{edge orbits } e \in H_i} n_e l(e) \\
&= A(r,i) L_i(\gamma_g) \\
&= A(r,i) L_i(g)
\end{aligned}$$

Thus

$$L_r(g\alpha) \leqslant \sum_{i \geqslant r} A(r,i) L_i(g)$$

□

We now iterate and consider the effect α^k has on a single stratum.

Lemma 6.2. *Let $g \in G$. Let m be the total number of strata in T. Then for each stratum H_r, and for all $k \geqslant 1$,*

$$\begin{aligned} L_r(g\alpha^k) &\leqslant \sum_{(i_1,\ldots,i_k)\in I_k[r,m]} A(r,i_1)A(i_1,i_2)\ldots A(i_{k-1},i_k)L_{i_k}(g), \textit{ where} \\ I_k[r,m] &= \{(i_1,\ldots,i_k) \in \mathbb{Z}^k \mid r \leqslant i_1 \leqslant \ldots \leqslant i_k \leqslant m\} \textit{ and} \\ A(i,j) &= \max_{e\in H_j} \frac{L_i(f(e))}{l(e)} \end{aligned}$$

Proof. We prove this by induction on k. The case of $k = 1$ is given by Lemma 6.1.

Assume that the statement holds for $k = n$. Then

$$\begin{aligned} L_r(g\alpha^{n+1}) &\leqslant \sum_{j\geqslant r} A(r,j)L_j(g\alpha^n) && \text{(by Lemma 6.1)} \\ &\leqslant \sum_{j\geqslant r} A(r,j)\left(\sum_{(i_1,\ldots,i_n)\in I_n[j,m]} A(j,i_1)A(i_1,i_2)\ldots A(i_{n-1},i_n)L_{i_n}(g)\right) && \text{(by inductive hypothesis)} \\ &= \sum_{j\geqslant r} \left(\sum_{(i_1,\ldots,i_n)\in I_n[j,m]} A(r,j)A(j,i_1)A(i_1,i_2)\ldots A(i_{n-1},i_n)L_{i_n}(g)\right) \\ &= \sum_{(j,i_1,\ldots,i_n)\in I_{n+1}[r,m]} A(r,j)A(j,i_1)A(i_1,i_2)\ldots A(i_{n-1},i_n)L_{i_n}(g) \\ &= \sum_{(i_1,\ldots,i_{n+1})\in I_{n+1}[r,m]} A(r,i_1)A(i_1,i_2)\ldots A(i_n,i_{n+1})L_{i_{n+1}}(g) && \text{(after relabelling indices)} \end{aligned}$$

Thus the statement holds for all $k \geqslant 1$.

□

Theorem 6.3. *Let $\alpha \in Out(G,\mathcal{G})$. Let $f : T \to \alpha T$ be a relative train track map on some $T \in \mathcal{O}(G,\mathcal{G})$. Let μ_R be the largest Perron-Frobenius eigenvalue of a stratum in T, and let m be the total number of strata in T. Then there exists a polynomial $P(k)$ of degree at most $m-1$ such that*

$$l_T(g\alpha^k) \leqslant P(k)\mu_R^k l_T(g)$$

for all $g \in G$.

Proof. It follows from Lemma 6.2 that for all k,

$$\begin{aligned} l_T(g\alpha^k) &= \sum_{r=1}^{m} L_r(g\alpha^k) \\ &\leqslant \sum_{r=1}^{m} \left(\sum_{(i_1,\ldots,i_k)\in I_k[r,m]} A(r,i_1)A(i_1,i_2)\ldots A(i_{k-1},i_k)L_{i_k}(g)\right) \end{aligned}$$

Thus we have an upper bound for $l_T(g\alpha^k)$. All that remains is to simplify it.

First we will find the size of the set $I_k[r,m]$, which will tell us how quickly the number of terms in our sum grows as $k \to \infty$. We can think of $I_k[r,m]$ as a set of combinations with repetition, where we choose k elements from the set $\{r,\ldots,m\}$ (which has size $m-r+1$). Therefore

$$\begin{aligned}|I_k[r,m]| &= C(m-r+1,k)\\ &= \frac{(k+m-r)!}{k!(m-r)!}\\ &= \frac{(k+m-r)(k+m-r-1)\ldots(k+1)}{(m-r)!}\end{aligned}$$

This is a polynomial of degree $m-r$ in k.

Secondly, consider the coefficients of the sum: $A(r,i_1)A(i_1,i_2)\ldots A(i_{k-1},i_k)L_{i_k}(g)$. If $i=j$, then $A(i,j)=\mu_i \leqslant \mu_R$; if $i \neq j$, then $A(i,j)$ can appear in each coefficient at most once. It follows that

$$A(r,i_1)A(i_1,i_2)\ldots A(i_{k-1},i_k) \leqslant \mathcal{A}\mu_R^k,$$

where

$$\mathcal{A} = \prod_{i\neq j} A(i,j).$$

Thirdly, we observe that $l_{T,i}(g) \leqslant l_T(g)$ for all i. Thus, combining all of the above:

$$\begin{aligned}l_T(g\alpha^k) &\leqslant \sum_{r=1}^{m}\left(\sum_{(i_1,\ldots,i_k)\in I_k[r,m]} A(r,i_1)A(i_1,i_2)\ldots A(i_{k-1},i_k)L_{i_k}(g)\right)\\ &\leqslant \sum_{r=1}^{m}\Big(|I_k[r,m]|\mathcal{A}\mu_R^k l_T(g)\Big)\\ &= P(k)\mu_R^k l_T(g)\end{aligned}$$

where $P(k) = \mathcal{A}\sum_{r=1}^{m}|I_k[r,m]|$ is a polynomial of degree at most $m-1$.

□

7 Appendix 2: Irreducible Automorphisms

In the case when $\alpha \in \mathrm{Out}(G,\mathcal{G})$ is irreducible, the proof that the relative growth rate and displacement are equal is simpler. This is because we can guarantee the existence of train track maps, not just relative train track maps. Furthermore, we can guarantee the existence of a point in $\mathcal{O}$ which realises the displacement, not just a sequence which tends towards it.

Definition 7.1. Let $\alpha \in \mathrm{Out}(G, \mathcal{G})$, $T \in \mathcal{O}(G, \mathcal{G})$. We say that a topological representative $f : T \to \alpha T$ is a *train track map* if every edge $e \in T$ is legal with respect to f.

When α is an irreducible automorphism, the transition matrix of any topological representative $f : T \to \alpha T$ will be an irreducible matrix - thus there is only one stratum of T (the whole tree), and only one PF-eigenvalue μ. As with the more general case, the positive row eigenvector corresponding to μ decides the metric on T. This ensures that the length of every edge is scaled exactly by μ, so $\forall e \in T$, $l(f(e)) = \mu l(e)$.

Lemma 7.2. *Let $f : T \to \alpha T$ be a train track map representing an irreducible automorphism α. Take the metric on T to be the one determined by the PF-eigenvector. Then $\exists g \in Hyp(\mathcal{G})$ such that $l_T(g\alpha^k) = \mu^k l_T(g)$ for all $k > 0$.*

Proof. This is a well-known property of train track maps, but it also follows from Lemma 4.10(v) since we can think of a train track map as a relative train track map with only a single stratum. In this specific case, legal and r-legal are equivalent conditions, and the r-length of a path is equal to its length. Thus the result follows.

□

Theorem 7.3. *Let $\mathcal{G}$ be a free factor system for a group G, let E be any relative generating set for G, and let $\alpha \in Out(G, \mathcal{G})$ be irreducible. Then $GR_{\mathcal{G}}(\alpha, l_E) = \lambda_\alpha$.*

Proof. $\mathrm{Min}(\alpha)$ is equal to the train-track bundle $\mathrm{TT}(\alpha)$ - the set of $T \in \mathcal{O}$ admitting train track representatives $f : T \to \alpha T$ with $\mathrm{Lip}(f) = \Lambda_R(T, \alpha T)$ [8, Thmm 8.19, Thm 6.11]. In addition, since α is irreducible, $\mathrm{Min}(\alpha)$ is non-empty [8, Theorem 8.4].

Thus we can guarantee the existence of a train track map $f : T \to \alpha T$ on some T such that:

$$\mu \overset{f \text{ is train track}}{=} \mathrm{Lip}(f) \overset{f \in \mathrm{TT}(\alpha)}{=} \Lambda_R(T, \alpha T) \overset{\text{by defn of } \mathrm{Min}(\alpha)}{=} \lambda_\alpha.$$

In addition, since $l_E \sim_{\mathcal{G}} l_T$, these length functions will produce the same growth rate.

Therefore in order to prove that $\mathrm{GR}_{\mathcal{G}}(\alpha, l_E) = \lambda_\alpha$ it suffices to prove that $\mathrm{GR}_{\mathcal{G}}(\alpha, l_T) = \mu$. The proof of this follows from the definition of the right stretching factor Λ_R. Recall,

$$\Lambda_R(T, \alpha^k T) := \sup_{g \in \mathrm{Hyp}(\mathcal{G})} \frac{l_T(g\alpha^k)}{l_T(g)}$$

$\Rightarrow$ For all $g \in G$, $l_T(g\alpha^k) \leqslant \mu^k l_T(g)$.

$\Rightarrow$ For all $g \in G$, $GR_{\mathcal{G}}(\alpha, g, l_T) = \limsup_{k\to\infty} \sqrt[k]{l_T(g\alpha^k)} \leqslant \limsup_{k\to\infty} \sqrt[k]{\mu^k l_T(g)} = \mu$.

Additionally, by Lemma 7.2, there exists $h \in G$ such that $\mu^k l_T(h) = l_T(h\alpha^k)$

$\Rightarrow GR_{\mathcal{G}}(\alpha, h, l_T) = \limsup_{k\to\infty} \sqrt[k]{l_T(h\alpha^k)} = \limsup_{k\to\infty} \sqrt[k]{\mu^k l_T(h)} = \mu$.

Thus $GR_{\mathcal{G}}(\alpha, l_T) := \sup_g GR_{\mathcal{G}}(\alpha, g, l_T) = \mu$, and we are done. $\square$

References

[1] Naomi Andrew. Serre's property (FA) for automorphism groups of free products. *Journal of Group Theory*, 24(2):385–414, 2021.

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

[3] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Annals of Mathematics. Second Series*, 135(1):1–51, 1992.

[4] Daniel E. Cohen. *Combinatorial group theory: a topological approach*, volume 14 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1989.

[5] D. J. Collins and E. C. Turner. Efficient representatives for automorphisms of free products. *Michigan Mathematical Journal*, 41(3):443–464, 1994.

[6] Matthew Collins. Fixed points of irreducible, displacement one automorphisms of free products. May 2023. https://arxiv.org/abs/2305.01451.

[7] Warren Dicks and Enric Ventura. Irreducible automorphisms of growth rate one. *Journal of Pure and Applied Algebra*, 88(1-3):51–62, 1993.

[8] Stefano Francaviglia and Armando Martino. Stretching factors, metrics and train tracks for free products. *Illinois Journal of Mathematics*, 59(4):859–899, 2015.

[9] Stefano Francaviglia and Armando Martino. Displacements of automorphisms of free groups I: Displacement functions, minpoints and train tracks. *Transactions of the American Mathematical Society*, 374(5):3215–3264, 2021.

[10] Denis V. Osin. Relatively hyperbolic groups: Intrinsic geometry, algebraic properties, and algorithmic problems. *Memoirs of the American Mathematical Society*, 179(843):vi+100, 2006.

Paper 3: **Length functions on groups and actions on graphs**

Matthew Collins and Armando Martino

ABSTRACT. We study generalisations of Chiswell's Theorem that 0-hyperbolic Lyndon length functions on groups always arise as based length functions of the the group acting isometrically on a tree. We produce counter-examples to show that this Theorem fails if one replaces 0-hyperbolicity with δ-hyperbolicity.
We then propose a set of axioms for the length function on a finitely generated group that ensures the function is bi-Lipschitz equivalent to a (or any) length function of the group acting on its Cayley graph.

1 Introduction

One of the key insights of geometric group theory is that one can obtain information on a group by viewing it as a metric space, via the word metric on its Cayley graph. More generally if a group, G, acts isometrically on a metric space (X, d) one can elucidate properties of the group from this action. For instance, the class of hyperbolic groups is precisely the class of those groups admitting a proper, co-compact isometric action on some locally compact, geodesic δ-hyperbolic space X.

Given a (right) isometric action of G on (X, d), and a point p in X, one can define a G-invariant pseudo-metric - which we denote by d_p - on G via $d_p(g, h) := d(pg, ph)$, which is a metric precisely when the stabiliser of p is trivial. In fact, this metric on G can be encoded via the *based length function*.

Definition 1.1. Let G act isometrically on the metric space (X, d). Then the *based length function* of G based at some point $p \in X$ is the function $l_p : G \to \mathbb{R}$, given by:

$$l_p(g) := d(p, pg)$$

It is straightforward to see that one can recover the invariant (pseudo) metric from the based length function via $d_p(g, h) = l_p(gh^{-1})$.

Of course, in order to obtain properties of the group it is helpful to impose conditions on the space and the action, just as for hyperbolicity above. A key area where one can recover a great deal of information about G is when X is a tree.

The source of inspiration for this paper is a striking result of Chiswell, that one can axiomatise the based length functions arising from actions on trees - sometimes called Lyndon length functions, following results from [6] - and, from the axioms, always recover an isometric action. Specifically,

Theorem 1.2 ([3]). *Suppose $l : G \to \mathbb{R}_{\geqslant 0}$ satisfies the following axioms:*

A1′: $l(g) = 0$ *if* $g = 1$,

A2 : $l(g^{-1}) = l(g)$,

A3 : $c(g, h) \geqq 0$,

H_0 : For all $g_1, g_2, g_3 \in G$,
$c(g_1, g_2) \geqslant m$, $c(g_2, g_3) \geqslant m$ *implies that* $c(g_1, g_3) \geqslant m$,

where

$$c(g, h) := \frac{1}{2}(l(g) + l(h) - l(gh^{-1})).$$

Then there exists an $\mathbb{R}$-tree, (X, d), admitting an isometric G-action and a point, $p \in X$, such that $l_p(g) = l(g)$. Moreover, if the images of l and c lie in $\mathbb{Z}$, then the tree will be simplicial.

Remark 1.3. As noted above a function $d : G \times G \to \mathbb{R}$ can be defined from l and, from this point of view, A1′ says that d vanishes on the diagonal, A2 says that it is symmetric and A3 says that it satisfies the triangle inequality.

The function $c(g, h)$ is then really the Gromov product and axiom H_0 should be thought of as a 0-hyperbolicity condition (see, for example, [1] for a discussion on hyperbolic groups, spaces and the Gromov product). Chiswell's Theorem can then be summarised as saying that a 0-hyperbolic Lyndon length function is always a based length function on a 0-hyperbolic space.

With this in mind, we are motivated to ask the following questions.

Questions.

- *Is there a generalisation of Chiswell's Theorem for isometric group actions on metric graphs?*
- *In particular, is there a generalisation of Chiswell's Theorem for isometric actions on δ-hyperbolic graphs?*

Remark 1.4. In the spirit of Chiswell's result, we will consider graphs whose edge lengths may not be integers. For instance, one could take the Cayley graph of a group with respect to some generating set, and then equivariantly assign positive real lengths to edges.

It turns out that these questions are somehow too broad in their scope. Given a (strictly positive) length function on G (see Definition 2.2 for the definition of a length function) there is always a metric graph whose based length function is equal to this function:

take the complete graph on G where the edge between g and h has length $l(hg^{-1})$ - Lemma 3.1. The based length function on this graph, with respect to the basepoint 1, is equal to l. However, this action is not particularly useful.

In order to rule out this kind of example we will add some restrictions.

Questions. *Let us suppose that G is finitely generated and let us restrict ourselves to isometric, co-compact actions on locally compact graphs, X.*

- *Given a (strictly positive) length function, l, does G admit an isometric, co-compact action on a locally compact metric graph, X, such that $l = l_p$ for some $p \in X$?*
- *What if we add the hypothesis that l is δ-hyperbolic (see Definition 2.2 for the definition of hyperbolicity)?*

It turns out that the answer to both of these questions is **no**. By Proposition 3.4, there exists a δ-hyperbolic length function which cannot arise as the based length function associated to any isometric, co-compact action on a locally compact graph.

However, that example is bi-Lipschitz equivalent to a length function on a Cayley graph. (Note that, for a finitely generated groups, all based length functions on Cayley graphs with respect to finite generating sets are bi-Lipschitz equivalent). But one can also produce examples of δ-hyperbolic length functions which are not bi-Lipschitz equivalent to any based length function on a Cayley graph, as in Proposition 3.8. In fact, *every* finitely generated group admits a hyperbolic length function.

Theorem 3.9. *There exists a finitely generated group, G, with a hyperbolic length function, $l : G \to \mathbb{R}_{\geqslant 0}$ such that $l \neq l_p$ for any co-compact, metric G-graph.*

Moreover, any *finitely generated group admits a (free) hyperbolic length function. In particular, we can find an example of a group G with a hyperbolic length function, l, which is not quasi-isometric to any based length function arising from an isometric action of G on a geodesic and proper δ-hyperbolic metric space.*

This leads us to the following.

Questions. *Suppose that G is finitely generated.*

- *Can one axiomatise those length functions which are bi-Lipschitz equivalent to some (and hence all) based length functions on a Cayley graph for G (with respect to a finite generating set)?*
- *Can we make these axioms apply to - for instance - any free F_n action on a simplicial tree as well as Cayley graphs?*

- *Does this axiomatisation define a connected/contractible/finite dimensional subspace of* $\mathbb{R}^G$ *on which* $\mathrm{Aut}(G)$ *acts?*

Remark 1.5. We do come up with an axiom scheme, below, and we observe that these axioms hold for all sufficiently well behaved actions - see Proposition 4.1 and Corollary 4.2 - and in particular to all points of Culler-Vogtmann space.

The third question here arises from the fact that one key use of Chiswell's Theorem is in the study of group actions on trees, and the definition of the space of such actions which are then encoded via functions (usually the translation length function, which is related to the Lyndon length function). See [5] for the seminal paper on the 'Outer Space' of free actions on trees, encoded by length functions (amongst other things).

It is clear that the space of all length functions which are bi-Lipscitiz equivalent to one arising from a Cayley graph is a contractible space (because a linear combination of such functions is another such function). Therefore, this provides a contractible space on which $\mathrm{Aut}(G)$ acts. However, it is far too large and so one might hope that an axiomatisation could provide a more reasonable subspace.

With these questions in mind, we propose the following axioms for our length functions:

Definition 4.3. Let G be a group. We say that $l : G \to \mathbb{R}_{\geqslant 0}$ is a *graph-like length function* if it satisfies the following axioms:

A1: $l(g) = 0$ if and only if $g = 1$,

A2: $l(g^{-1}) = l(g)$,

A3: $c(g, h) \geqq 0$,

A4: for all $R \geqslant 0$, the closed ball $B_R := \{g \in G \mid l(g) \leqslant R\}$ is finite,

A5: there exists $0 \leqslant \epsilon < 1$ and $K > 0$ such that, for any $g \in G$, if $l(g) > K$ then there exists an $x \in G$ with:

(i) $0 < l(x) \leqslant K$, and

(ii) $c(gx^{-1}, x^{-1}) \leqslant \frac{\epsilon l(x)}{2}$.

Remark 1.6. Here, the mysterious looking axiom A5 is encoding the fact that if one had a reasonable action on a graph, then one could approximate geodesics in the graph with uniform quasi-geodesics built from the translates of finitely many paths; it is morally a co-compactness condition expressed solely in terms of the length function. In fact, we prove that this axiom holds for a fairly wide class of actions in Proposition 4.1 and Corollary 4.2.

We also note that if G acts on its Cayley graph then one easily gets that the based length function satisfies these axioms with $K = 1$ and $\epsilon = 0$. However, if once considers actions on graphs with more than one orbit of vertices, then one quickly discovers that the correct condition is A5(ii) with $\epsilon \neq 0$. Moreover, scaling the graph by a constant clearly changes the value of K. For these reasons, to allow these kinds of deformations, we consider these axioms for more general K and ϵ.

It turns out that this is indeed sufficient to prove the following:

Theorem 4.9. *Let $l : G \to \mathbb{R}_{\geqslant 0}$ be a graph-like length function on a group G. Then l is bi-Lipschitz equivalent to some (and hence to all) based length function l_p arising from a locally compact, co-compact, metric G-graph and with $Stab(p) = 1$.*

Note that in view of Theorem 3.9, since any finitely generated group admits a hyperbolic length function, the extra axioms are clearly necessary.

Remark 1.7. We should note that another length function one can extract from an action is the *translation length function*, which has the advantage of not relying on a basepoint. This is the point of view of [4]. An important result here, building on the work of [4], is that of [7] which states that a translation length function (which is 0-hyperbolic) always arises from an action on a tree. However, this builds crucially on Chiswell's Theorem 1.2 so it seems reasonable to start with Lyndon length functions.

2 Preliminaries

We begin with some preliminary definitions and notation. Let G be a group.

Definition 2.1. Given a metric, $d : G \times G \to \mathbb{R}_{\geqslant 0}$, on a group G we say that d is *right-invariant* if $d(g_1h, g_2h) = d(g_1, g_2)$ for all $g_1, g_2, h \in G$.

Definition 2.2. A map $l : G \to \mathbb{R}_{\geqslant 0}$ which satisfies the following axioms is called a *length function*:

A1: $l(g) = 0$ if and only if $g = 1$,

A2: $l(g^{-1}) = l(g)$,

A3: $c(g, h) \geqq 0$ where

$$c(g, h) := \frac{1}{2}(l(g) + l(h) - l(gh^{-1}))$$

is the *Gromov product* of $g, h \in G$.

If, in addition, l satisfies

H_δ: $c(g_1, g_2) \geqslant m$, $c(g_2, g_3) \geqslant m$ implies that $c(g_1, g_3) \geqslant m - \delta$

for some $\delta \geqslant 0$, we say it is a *δ-hyperbolic* length function. The condition H

$isreferredtoas\delta$-*hyperbolicity*.

Remark 2.3. Given a length function, it is easy to verify that $d(g, h) := l(gh^{-1})$ is a right-invariant metric on G. In particular, A3 is equivalent to the triangle inequality, which can be written as

$$l(gh^{-1}) \leqq l(g) + l(h).$$

Also note that here we write axiom A1: $l(g) = 0$ if and only if $g = 1$ rather than A1′: $l(g) = 0$ if $g = 1$. This is largely because we end up wanting to characterise those length functions which are bi-Lipschitz equivalent (or quasi-isometric) to those arising from Cayley graphs. We will sometimes emphasise this by saying that the length function is *free*.

Definition 2.4. A *metric graph* is a 1-dimensional CW-complex with a metric structure. A *metric tree* is a metric graph in which any two vertices are connected by exactly one simple path. We always equip metric graphs with the path metric.

Definition 2.5. By a *metric G-graph*, we mean a metric graph Γ together with an isometric right action of G on Γ, sending vertices to vertices and edges to edges.

Since we think of our graphs as metric spaces, given a point p in Γ, we may invoke Definition 1.1; $l_p(g) = d_\Gamma(p, p.g)$ is the based length function on Γ, based at p.

3 Hyperbolicity, Length Functions and Counter-Examples

Given a length function, l, as in Definition 2.2 - that is to say, given a metric on G - one can always construct *some* metric graph on which G acts isometrically and such that $l = l_p$:

Lemma 3.1. *Let l be a length function on the group, G, as in Definition 2.2.*

Let Γ be the complete graph on vertex set G, where the length of the edge between g and h is set to $l(gh^{-1}) = l(hg^{-1})$. Then G acts isometrically on Γ and l is equal to the based length function on Γ - Definition 1.1 - based at the vertex 1.

However, this is not a very useful object and we will want to insist on some finiteness conditions; namely, co-compactness and (usually) local compactness.

Since this work arose as an attempt to generalise the celebrated result of Chiswell, Theorem 1.2, it is a natural way to try to generalise that result by weakening 0-hyperbolicity to δ-hyperbolicity and instead only expecting the action to be on a (hyperbolic) graph. It turns out that this doesn't work and we present two counter-examples, in Proposition 3.4 and Proposition 3.8.

Before presenting the first example, it is worthwhile observing some examples of hyperbolic length functions which *do* arise as the length function of a co-compact action on a graph. In these examples, we can take an existing length function and deform it slightly, but the following examples show that in doing so one might still end up with a length function arising from an action on a graph.

Examples 3.2. For both of these examples, our group is the infinite cyclic group, $\mathbb{Z}$.

(i) Given $0 \leqslant \epsilon < 1$, define:

$$l(n) = \begin{cases} 1+\epsilon & \text{if } n = \pm 1 \\ |n| & \text{otherwise} \end{cases}$$

One can verify that this is the length function of the Cayley Graph of $\mathbb{Z}$, with respect to the generating set $\{1, 2, 3\}$ where 1 is given length $1 + \epsilon$, 2 is given length 2 and 3 is given length 3.

(ii) Again, given $0 \leqslant \epsilon < 1$, define:

$$l(n) = \begin{cases} 0 & \text{if } n = 0 \\ |n| + \epsilon & \text{otherwise} \end{cases}$$

This is actually 0-hyperbolic, and arises from a non-minimal action on a tree. More precisely, take a graph with two vertices, u and v, and two edges, one of which is a loop of length 1 at v and the other is an edge of length $\epsilon/2$ joining u to v. The fundamental group of that graph is $\mathbb{Z}$ and the action on the universal cover gives our length function (with respect to any lift of u).

Next we show how to deform the standard length function on $\mathbb{Z}$ so as to end up with something that does not arise from an action.

In order to proceed, we need the following observation:

Lemma 3.3. *Let Γ be a co-compact, metric G-graph and $p \in \Gamma$. Let $l_p(g) = d_\Gamma(p, p.g)$ denote the based length function. Then there exist finitely many positive real numbers, $\alpha_1, \ldots, \alpha_k$ such that, for any $g \in G$, $l_p(g)$ belongs to the submonoid of the (additive) real numbers generated by the α_i.*

That is, for every g, there exist non-negative integers n_i such that $l_p(g) = \sum n_i \alpha_i$.

Proof. We simply let the α_i be the lengths of the edges in Γ. Since the action is isometric and there are finitely many edge-orbits, it suffices to take only finitely many of them. □

Now we are ready to show that a δ-hyperbolic length function need not come from an action on a graph.

Proposition 3.4. *For any $0 \leqslant \epsilon < 1$, the function $l_\epsilon : \mathbb{Z} \to \mathbb{R}_{\geqslant 0}$ defined by $l_\epsilon(n) = |n| + \epsilon^{|n|}$, for $n \neq 0$ and $l(0) = 0$ is a hyperbolic length function.*

For $\epsilon = 1/2$, this cannot be equal to any based length function arising from a co-compact, isometric action of $\mathbb{Z}$ on a metric graph.

Proof. First we verify axioms A1 to A3 and H_δ from Definition 2.2. Note that for $\epsilon = 0$, this is just the standard length function of $\mathbb{Z}$ acting on the line (which is 0-hyperbolic). For each l_ϵ we define c_ϵ to be the corresponding Gromov product. Note that both l_0 and c_0 take values in $\mathbb{Z}$.

Observe that A1 and A2 are clear for all ϵ directly from the definition. To verify A3, notice that

$$l_0(n) \leqslant l_\epsilon(n) \leqslant l_0(n) + \epsilon,$$

and hence that for any $n, m \in \mathbb{Z}$,

$$c_\epsilon(n, m) \geqslant c_0(n, m) - \epsilon/2.$$

Therefore, since c_0 takes values in $\mathbb{Z}$, and $\epsilon < 1$, the only values for which $c_\epsilon(n, m)$ could be negative would be those where $c_0(n, m) = 0$. Since, for positive integers n, k we have that $c_0(n, n + k) = c_0(-n, -n - k) = n$, we see that $c_0(n, m)$ can only be zero if one of n, m is zero or if one is positive and one is negative. We calculate: if n, m are positive then

$$c_\epsilon(0, n) = c_\epsilon(0, -n) = 0$$

and,

$$c_\epsilon(n, -m) = \epsilon^n + \epsilon^m - \epsilon^{n+m} > 0.$$

This verifies A3. To verify H_δ, note that the inequality $l_0(n) \leqslant l_\epsilon(n) \leqslant l_0(n) + \epsilon$ also gives us that $\epsilon + c_0(n, m) \geqslant c_\epsilon(n, m)$. Hence we get, for all n, m,

$$\epsilon + c_0(n, m) \geqslant c_\epsilon(n, m) \geqslant c_0(n, m) - \epsilon/2.$$

But since l_0 is 0-hyperbolic, this implies that l_ϵ is $\frac{3\epsilon}{2}$-hyperbolic.

To see that $l_{1/2}$ cannot arise as the length function coming from a co-compact metric $\mathbb{Z}$-graph, we invoke Lemma 3.3 and argue by contradiction. That is, suppose that $l_{1/2}$ arises from the action of $\mathbb{Z}$ on a co-compact metric graph, Γ. Then, by Lemma 3.3, we have $\alpha_1, \ldots, \alpha_k$ such that for any $g \in G$, there exist positive integers, $n_1, \ldots, n_k$ with $l_{1/2}(g) = \sum_{i=1}^{k} n_i \alpha_i$. We now show that this is not possible.

Without loss of generality, by enlarging the set, we may assume that $\alpha_1 = 1$. Further, again without loss of generality, we may assume that $\alpha_1, \ldots, \alpha_r$ is a maximal, $\mathbb{Q}$-linearly independent subset of the α_i. Thus for any $j > r$, α_j is a $\mathbb{Q}$-linear sum of $\alpha_1, \ldots, \alpha_r$. Fix such an expression for each j (in fact, it is unique) and notice that the denominators in the coefficients of these expressions are bounded. In particular, this means that any expression $\sum_{i=1}^{k} n_i \alpha_i$, where the n_i are integers can be re-written as an expression $\sum_{i=1}^{r} q_i \alpha_i$, where the q_i are now rational, but with bounded denominator. In particular, this means that there exists an integer, M, such that for any $g \in \mathbb{Z}$, there exists integers m_i such that,

$$l_{1/2}(g) = \frac{1}{M} \sum_{i=1}^{r} m_i \alpha_i.$$

However, notice that $l_{1/2}(g)$ are rational for every g, and the set $\alpha_1, \ldots, \alpha_r$ are $\mathbb{Q}$-linearly independent. Hence the $\mathbb{Q}$-linear independence forces $m_i = 0$ for $i \geqslant 2$, and therefore,

$$l_{1/2}(g) = \frac{1}{M} m_1 \alpha_1 = \frac{1}{M} m_1.$$

This is clearly impossible, since the values of $l_{1/2}$ do not belong to the additive cyclic subgroup generated by a rational number. $\square$

Remark 3.5. Note that the same proof shows that l_ϵ cannot be equal to any length function arising from a co-compact, isometric action of $\mathbb{Z}$ on a metric graph for any rational ϵ.

The idea of Proposition 3.4 is that we started with a 0-hyperbolic length function (which is the standard length function of $\mathbb{Z}$ acting on its Cayley graph) and deformed it slightly to obtain a length function that is δ-hyperbolic but is not equal to any based length function coming from a co-compact graph. Naturally, since this is a small deformation we obtain a length function which is bi-Lipschitz equivalent to the original length function. We could also consider quasi-isometry.

Definition 3.6. We say that two length functions l_1, l_2 on a group G are *quasi-isometric* if there exists $A \geqslant 1, B \geqslant 0$ such that, $\forall g \in G$,

$$\frac{1}{A} l_1(g) - B \leqslant l_2(g) \leqslant A l_1(g) + B.$$

If, in addition, we can take $B = 0$, we say that l_1, l_2 are bi-Lipschitz equivalent.

We record a standard consequence of the Svarc-Milnor Lemma (see, for example, [2], I.8.19):

Lemma 3.7. *Let X, Y be co-compact, locally compact metric G-graphs. Then for all points $p \in X$, $q \in Y$ such that $Stab(p) = Stab(q) = 1$, the based length functions l_p and l_q are bi-Lipschitz equivalent.*

Instead of seeking length functions on G which are equal to the based length function of a suitable G-graph, we can instead seek $l : G \to \mathbb{R}_{\geqslant 0}$ which lies in the *quasi-isometry class* of a suitable G-graph, ideally a Cayley graph for G. Our aim is then to produce axioms for a length function that make it quasi-isometric, or bi-Lipschitz equivalent to a based length function on a Cayley graph.

Even here it turns out that hyperbolicity is not sufficient.

Proposition 3.8. *Let G be a finitely generated group and let $|.| : G \to \mathbb{R}$ be the word metric with respect to some finite generating set. Define a function, $l : G \to \mathbb{R}$ by $l(g) := \log(|g|+1)$. Then this is a δ-hyperbolic length function, for a uniform $\delta = \frac{1}{2}\log 32$.*

When $G = \mathbb{Z}$ then l is not quasi-isometric (and hence not bi-Lipschitz equivalent) to any *based length function on a geodesic and proper hyperbolic space - for an isometric action of $\mathbb{Z}$.*

Proof. First we verify the axioms from Definition 2.2. We immediately see that l satisfies axioms A1 and A2. To see that A3 holds we observe that for all $g, h \in G$,
$|g| + |h| \geqslant |gh^{-1}|$. Thus,

$$\begin{aligned}
\log(|g|+1) + \log(|h|+1) &= \log((|g|+1)(|h|+1)) \\
&= \log(|g||h| + |g| + |h| + 1) \\
&\geqslant \log(|g| + |h| + 1) \\
&\geqslant \log(|gh^{-1}| + 1) \\
\Rightarrow c(g,h) = \frac{1}{2}(l(g) + l(h) - l(gh^{-1})) \geqslant 0.
\end{aligned}$$

Thus l is a length function.

To see that the length function is δ-hyperbolic, consider the function,

$$d(g,h) := e^{2c(g,h)} = \frac{(|g|+1)(|h|+1)}{|gh^{-1}|+1}, g, h \in G.$$

It will be sufficient to show that there exists a $\delta \geqslant 0$ such that for any three group elements, g, h, k, and any $R \geqslant 0$,

$$d(g,h) \geqslant e^{2R} \text{ and } d(h,k) \geqslant e^{2R} \implies d(g,k) \geqslant e^{2(R-\delta)}.$$

To do this, first observe the following two inequalities:

$$|g| \geqslant 2|h| \implies 2(|h|+1) \geqslant d(g,h) \tag{1}$$
$$d(g,h) \geqslant \min\{\frac{|g|+1}{2}, \frac{|h|+1}{2}\}. \tag{2}$$

To see that (1) is true, simply observe that if $|g| \geqslant 2|h|$ then,

$$|gh^{-1}|+1 \geqslant |g|-|h|+1 \geqslant \frac{|g|}{2}+1 \geqslant \frac{|g|+1}{2},$$

from which it follows that $2(|h|+1) \geqslant d(g,h)$.

To see that (2) is true, observe that if $|g| \geqslant |h|$ then, $|gh^{-1}|+1 \leqslant |g|+|h|+1 \leqslant 2(|g|+1)$, from which the desired inequality follows.

To verify that our length function is hyperbolic, let us suppose that we have a triple of group elements, g,h,k, and a real number, $R \geqslant 0$ such that $d(g,h) \geqslant e^{2R}$ and $d(h,k) \geqslant e^{2R}$.

Our aim is to find a (uniform) $\delta > 0$ such that $d(g,k) \geqslant e^{2(R-\delta)}$.

We set $\Lambda = \max\{|g|,|h|,|k|\}$ and $\lambda = \min\{|g|,|h|,|k|\}$ the argument breaks into two cases now, depending on whether $\Lambda \geqslant 4\lambda$, or $\Lambda < 4\lambda$.

case(i): $\Lambda \geqslant 4\lambda$:

Without loss of generality, we will assume that $|g| \geqslant |k|$. In particular this implies, from Equation (2), that $d(g,k) \geqslant \frac{|k|+1}{2}$.

Suppose first that $|h| \geqslant 2|k|$. Then from Equation (1), $2(|k|+1) \geqslant d(h,k)$. Therefore,

$$d(g,k) \geqslant \frac{|k|+1}{2} \geqslant \frac{d(h,k)}{4} \geqslant \frac{e^{2R}}{4},$$

as required (with $\delta = \log(2)$). (We haven't used the fact that $\Lambda \geqslant 4\lambda$ yet).

If instead we have that, $|h| < 2|k|$ then we must get that $|g| > 2|h|$, since $\Lambda \geqslant 4\lambda$.

Hence equations (1), (2) give us that

$$d(g,k) \geqslant \frac{|k|+1}{2} > \frac{|h|+1}{4} \geqslant \frac{d(g,h)}{8} \geqslant \frac{e^{2R}}{8},$$

as required (here with $\delta = \frac{1}{2}\log(8)$).

case(ii): $\Lambda < 4\lambda$.

Here, we invoke the triangle inequality to get that:

$$|gk^{-1}| + 1 \leqslant 2\max\{|gh^{-1}| + 1, |hk^{-1}| + 1\}.$$

Without loss of generality, we assume that $|gh^{-1}| \geqslant |hk^{-1}|$. Then,

$$d(g,k) \geqslant \frac{(\lambda+1)^2}{2(|gh^{-1}|+1)} > \frac{(\Lambda+1)^2}{32(|gh^{-1}|+1)} \geqslant \frac{d(g,h)}{32} \geqslant \frac{e^{2R}}{32}.$$

This completes the proof that our length function is δ-hyperbolic (with $\delta = \frac{1}{2}\log(32)$ as the final and maximal estimate).

To finish, note that if $\mathbb{Z}$ were to act isometrically on a locally compact hyperbolic space, X, with based length function l_p, then either l_p would have to be bounded, or quasi-isometric to a linear function. Since $\log(|n|+1)$ is neither, it is not quasi-isometric to such an l_p. □

Theorem 3.9. *There exists a finitely generated group, G, with a hyperbolic length function, $l : G \to \mathbb{R}_{\geqslant 0}$ such that $l \neq l_p$ for any co-compact, metric G-graph.*

Moreover, any *finitely generated group admits a (free) hyperbolic length function. In particular, we can find an example of a group G with a hyperbolic length function, l, which is not quasi-isometric to any based length function arising from an isometric action of G on a geodesic and proper δ-hyperbolic metric space.*

Proof. This is simply the content of Propositions 3.4 and 3.8. □

4 Axioms for graph-like length functions

We finally turn to positive results and produce a set of axioms that do result in length functions which are bi-Lipschitz equivalent to the based length function on a (or any) Cayley graph.

Before introducing our axioms, we would like to demonstrate that they are *reasonable*, to the extent that they arise naturally from group actions on fairly general yet well behaved spaces. So we consider the following, noting that the hypotheses on X are satisfied by a locally finite metric graph equipped with the path metric and a co-compact group action.

Proposition 4.1. *Let X be a geodesic metric space with a given basepoint, p. Suppose a group, G, acts on X isometrically, and co-boundedly. Then there exist constants, $K > 0$ and $0 < \epsilon_0 \leqslant 1$ such that for any $g \in G$ with $d(p, pg) \geqslant K$, there exists an $x \in G$ such that:*

- $0 < d(p, px) \leqslant K$ *and,*
- $\epsilon_0 d(p, px) + d(px, pg) \leqslant d(p, pg)$

Proof. Recall that,

- X geodesic means that for any two points in X there exists an isometry from a closed real interval to X where the images of the endpoints are our given two points of X.
- The action is co-bounded means there is a closed ball whose G translates cover X.

Since the action is co-bounded, there exists a closed ball centered at p, of radius $K/3$, say, whose G translates cover X. Set $B = B_{K/3}(p)$ to be this ball.

Given $g \in G$ with $d(p, pg) \geqslant K$, let $q \in X$ be the point on a geodesic from p to pg such that $d(p, q) = K/2$. Since q is on a geodesic we also have $d(p, pg) = d(p, q) + d(q, pg)$.

Now, since the translates of B cover X, there exists some $x \in G$ such that $q \in Bx$. This implies that $d(q, px) \leqslant K/3$.

First note that $d(p, px) > 0$ since

$$d(p, px) \geqslant d(p, q) - d(px, q) \geqslant K/2 - K/3 = K/6 > 0.$$

Next note that,

$$d(p, px) \leqslant d(p, q) + d(q, px) \leqslant K/2 + K/3 = 5K/6.$$

and also,

$$d(px, pg) \leqslant d(px, q) + d(q, pg) \leqslant K/3 + (d(p, pg) - K/2) = d(p, pg) - K/6.$$

Putting these together we get that,

$$\frac{1}{5} d(p, px) + d(px, pg) \leqslant K/6 + (d(p, pg) - K/6) = d(p, pg).$$

Hence we are done, with $\epsilon_0 = 1/5$. □

Corollary 4.2. *With the same hypotheses as above, set:*

- $l_p(g) = d(p, pg)$ *and*
- $c_p(g, h) = \frac{1}{2}(l_p(g) + l_p(h) - l_p(gh^{-1}))$.

Then there exists $0 \leqslant \epsilon < 1$ *and* $K > 0$ *such that, for any* $g \in G$*, if* $l_p(g) > K$ *then there exists an* $x \in G$ *with:*

(i) $0 < l_p(x) \leqslant K$*, and*

(ii) $c_p(gx^{-1}, x^{-1}) \leqslant \frac{\epsilon l_p(x)}{2}$.

Proof. Just set $\epsilon = 1 - \epsilon_0$ from Proposition 4.1 since

$$\begin{aligned} c_p(gx^{-1}, x^{-1}) &\leqslant \frac{\epsilon l_p(x)}{2} \quad \Longleftrightarrow \\ l_p(gx^{-1}) + l_p(x) - l_p(g) &\leqslant \epsilon l_p(x) \quad \Longleftrightarrow \\ (1-\epsilon) l_p(x) + l_p(gx^{-1}) &\leqslant l_p(g) \end{aligned}$$

and the last line is equivalent to the conclusion of Proposition 4.1 (where we have also used the fact that $l_p(w) = l_p(w^{-1})$ for all w which is just a consequence of the symmetry of the metric). □

The idea is that a fairly general class of spaces and actions satisfy the equation given by Corollary 4.2 and hence we will add this as an axiom for our length functions. Therefore, we propose the following.

Definition 4.3. Let G be a group. We say that $l : G \to \mathbb{R}_{\geqslant 0}$ is a *graph-like length function* if it satisfies the following axioms:

A1: $l(g) = 0$ if and only if $g = 1$,

A2: $l(g^{-1}) = l(g)$,

A3: $c(g, h) \geqq 0$,

A4: for all $R \geqslant 0$, the closed ball $B_R := \{g \in G \mid l(g) \leqslant R\}$ is finite,

A5: there exists $0 \leqslant \epsilon < 1$ and $K > 0$ such that, for any $g \in G$, if $l(g) > K$ then there exists an $x \in G$ with:

(i) $0 < l(x) \leqslant K$, and

(ii) $c(gx^{-1}, x^{-1}) \leqslant \frac{\epsilon l(x)}{2}$.

Remark 4.4. We note that A4 is really a statement about the action being properly discontinuous, especially in view of Proposition 4.8, which says that in the presence of A5, A4 is equivalent to the statement that B_K is finite.

In view of Proposition 4.1 and Corollary 4.2 one should view A5 as a co-compactness condition; the challenge here was writing an axiom down which could be stated purely in terms of the length function.

As noted in the introduction, for a standard Cayley graph, its based length function will satisfy these axioms with $K = 1$ and $\epsilon = 0$. The A5 condition with $\epsilon = 0$ is effectively saying that for every $g \in G$, there is an x of length at most K, such that xp lies on a geodesic from p to pg.

For an example of a group acting on a graph where $\epsilon \neq 0$, consider the free group of rank 2, F_2, realised as the fundamental group of a graph with two vertices, u, v, and three edges: an edge-loop at u, E_u, an edge-loop at v, E_v, and an edge from u to v, E_{uv}. The action of F_2 on the universal cover, T, of this graph will induced a based length function which is graph-like, but not with $\epsilon = 0$.

Namely, take a lift $\overline{u}$ of u, as the basepoint of T and consider the orbit of $\overline{u}$ under the group elements corresponding to elements of the fundamental group of the form $g_n = E_{uv}E_v^n {E_{uv}}^{-1}$, for $n \in \mathbb{Z}$. Then the geodesic from $\overline{u}$ to $\overline{u}g_n$ only meets the orbit of $\overline{u}$ at its endpoints. Hence this cannot satisfy the A5 condition with $\epsilon = 0$ for any K.

In fact, it is straightforward to see that any free F_n action on a metric tree - that is, any point in Culler Vogtmann space - satisfies the axioms above, with $K = 1$ but not necessarily with $\epsilon = 0$.

Let us start with the following preparatory results:

Lemma 4.5. *A length function satisfying A4 is discrete.*

Proof. Recall that we say a length function $l : G \to \mathbb{R}_{\geqslant 0}$ is discrete if there exists $\mu > 0$ such that, for all non-trivial $g \in G$, $l(g) \geqslant \mu$.

If $G = 1$, then l is immediately discrete. Otherwise, take $\mu = \min\{R > 0 \mid B_R \neq 1\}$. Since A4 holds, this minimum will be realised by some $R > 0$. □

Lemma 4.6. *Given l satisfying A5, set $\lambda = \frac{1}{1-\epsilon}$. Then for the g, x listed in A5, we have that:*

$$l(gx^{-1}) \leqslant l(g) - \frac{1}{\lambda} l(x).$$

Proof. By A5,

$$\begin{aligned} c(gx^{-1}, x^{-1}) &\leqslant \frac{\epsilon l(x)}{2} \\ \Rightarrow \frac{1}{2}(l(gx^{-1}) + l(x) - l(g)) &\leqslant \frac{\epsilon l(x)}{2} \\ \Rightarrow l(gx^{-1}) + l(x) - l(g) &\leqslant \epsilon l(x) \\ \Rightarrow l(gx^{-1}) &\leqslant l(g) - (1-\epsilon) l(x) \\ \Rightarrow l(gx^{-1}) &\leqslant l(g) - \frac{1}{\lambda} l(x) \end{aligned}$$

□

Lemma 4.7. *The ball $B_K = \{g \in G \mid l(g) \leqslant K\}$ is a generating set for G. In particular, by A4, G is finitely generated.*

Proof. We will show, by induction on n, that $\langle B_K \rangle$ contains all group elements g with

$$l(g) \leqslant K + \frac{n\mu}{\lambda}$$

(taking λ from Lemma 4.6 and μ from Lemma 4.5) and hence contains all of G.

Firstly, take $g \in G$ such that $l(g) \leqslant K$. Then $g \in B_K \in \langle B_K \rangle$, and we are done.

Now assume that, for all $g \in G$ satisfying $l(g) \leqslant K + \frac{(n-1)\mu}{\lambda}$, g lies in $\langle B_K \rangle$.

Take g such that $l(g) \leqslant K + \frac{n\mu}{\lambda}$. Then, by Lemma 4.6, there exists $x \in B_k$ such that

$$\begin{aligned} l(gx^{-1}) &\leqslant l(g) - \frac{1}{\lambda} l(x) \\ &\leqslant K + \frac{n\mu}{\lambda} - \frac{\mu}{\lambda} \qquad\qquad \text{(by properties of } g \text{ and Lemma 4.5)} \\ &= K + \frac{(n-1)\mu}{\lambda} \end{aligned}$$

Thus $gx^{-1} \in \langle B_K \rangle$, and since $x \in B_K$, this means that $g = gx^{-1}x \in \langle B_K \rangle$. □

Proposition 4.8. *Let $l : G \to \mathbb{R}_{\geqslant 0}$ satisfy A1, A2, A3 and A5. Let K, ϵ be as in A5, let $\lambda = \frac{1}{1-\epsilon}$, and suppose that the ball $B_K = \{g \in G \mid l(g) \leqslant K\}$ is finite. Then,*

(a) For any nontrivial $g \in G$, there exists a finite sequence, $x_0, \ldots, x_k$ such that:

(i) each $0 < l(x_i) \leqslant K$ (i.e. each $x_i \in B_K \backslash \{1\}$),

(ii) $l(g{x_0}^{-1}{x_1}^{-1} \ldots {x_k}^{-1}) = 0$, and

(iii) $\frac{1}{\lambda} \sum_{i=0}^{k} l(x_i) \leqslant l(g) \leqslant \sum_{i=0}^{k} l(x_i)$.

(b) Axiom A4 holds - that is to say, the ball B_R is finite for all $R \geqslant 0$.

Proof. Part *(a)* is clearly true if $l(g) \leqslant K$, since we can just take $x_0 = g$. To prove it in general, we use the discreteness of the length function to argue by induction. More precisely, we let P_n be the statement that *(a)* holds for all g with $l(g) \leqslant K + \frac{n\mu}{\lambda}$. Thus our initial observation is that P_0 holds. We also observe that, since B_K is finite, there exists a minimum length, $\mu > 0$, for elements in B_K.

Suppose then that P_{n-1} holds and consider a $g \in G$ with $l(g) \leqslant K + \frac{n\mu}{\lambda}$. If $l(g) \leqslant K$ then we are done, as above. Otherwise, by Lemma 4.6 and the existence of μ, there exists an $x \in G$ with $0 < l(x) \leqslant K$ and

$$l(gx^{-1}) \leqslant l(g) - \frac{1}{\lambda} l(x) \leqslant K + \frac{(n-1)\mu}{\lambda}. \tag{3}$$

Now by the induction hypothesis applied to $g_0 = gx^{-1}$ we can find $x_1, \ldots, x_k \in G$ such that

(i) each $0 < l(x_i) \leqslant K$

(ii) $l(g_0 {x_1}^{-1} {x_2}^{-1} \ldots {x_k}^{-1}) = 0$, and

(iii) $\frac{1}{\lambda} \sum_{i=1}^{k} l(x_i) \leqslant l(g_0) \leqslant \sum_{i=1}^{k} l(x_i)$.

Now set $x = x_0$. Then, by Equation 3, we have that

$$\frac{1}{\lambda} \sum_{i=0}^{k} l(x_i) = \frac{1}{\lambda} l(x) + \frac{1}{\lambda} \sum_{i=1}^{k} l(x_i) \leqslant \frac{1}{\lambda} l(x) + l(g_0) \leqslant l(g).$$

Moreover, by A3,

$$l(g) \leqslant l(x) + l(g_0) \leqslant \sum_{i=0}^{k} l(x_i).$$

Hence, by induction, *(a)* result holds for all g.

We now prove *(b)*. Let $M \geqslant 0$, and let S_M denote the following set of finite sequences:

$$S_M = \left\{ x_0, \ldots x_k \in B_K \backslash \{1\} \mid \sum_{i=1}^{k} l(x_i) \leqslant M \right\}$$

Let $R \geqslant 0$, and let $g \in B_R$. By *(a)*, there exists a sequence $x_0, \ldots, x_k \in S_{R\lambda}$ such that $g = x_k \ldots x_0$. Therefore, if we can prove that $S_{R\lambda}$ is finite, we will prove that B_R is finite.

Recall that, for all $x \in B_K$, $l(x) \geqslant \mu$. Therefore, for all sequences in $S_{R\lambda}$,

$$R\lambda \geqslant \sum_{i=1}^{k} l(x_i) \geqslant k\mu \Rightarrow k \leqslant \frac{R\lambda}{\mu}.$$

Thus $S_{R\lambda}$ is a set of sequences of elements from the finite set B_K, and the length k of these sequences has an upper bound. Thus $S_{R\lambda}$ is finite, and we are done. □

Theorem 4.9. *Let $l : G \to \mathbb{R}_{\geqslant 0}$ be a graph-like length function on a group G. Then l is bi-Lipschitz equivalent to some (and hence to all) based length function l_p arising from a locally compact, co-compact, metric G-graph and with $Stab(p) = 1$.*

Proof. We take Γ to be the Cayley graph on the set $B_K = \{g \in G \mid l(g) \leqslant K\}$, but instead of assigning every edge length 1, we assign it the length of the corresponding generating element under l. That is, the vertex set of Γ is G, and we join two vertices, g, h by an edge if and only if $gh^{-1} = y \in B_K$; in that case we assign that edge a length of

$l(y)$. (Note that $hg^{-1} = y^{-1}$ will also be in B_K in that case and have the same length). Γ is then equipped with the path metric.

We then take the base point p to be the vertex corresponding to the identity. By Lemma 4.7, B_K is a finite generating set for G; hence Γ is well-defined and the action of G is co-compact.

We can immediately see that $0 = l(1) = l_p(1)$, so we shall restrict our attention to nontrivial $g \in G$. For $g \in G$ we write, as always, $l_p(g) = d_\Gamma(p, p \cdot g)$ to denote the based length induced by Γ. The metric on Γ is the path metric, and so the distance from p to pg is the infimum of the lengths of all edge paths from p to $p \cdot g$. Thus for all $g \neq 1$,

$$\begin{aligned} l_p(g) &= \inf\Big\{ \sum_{i=0}^{k} l_p(y_i) \mid y_0, \ldots, y_k \in B_K, y_k \ldots y_0 = g \Big\} \\ &= \inf\Big\{ \sum_{i=0}^{k} l(y_i) \mid y_0, \ldots, y_k \in B_K, y_k \ldots y_0 = g \Big\}, \end{aligned}$$

where the second equality arises from the fact that $l_p(y) = l(y)$ for all $y \in B_K$, as these are the edges of Γ.

By Proposition 4.8 *(a)*, there exists a sequence $x_0, \ldots, x_k \in B_K$ such that $x_k \ldots x_0 = g$ and $\frac{1}{\lambda} \sum_{i=0}^{k} l(x_i) \leqslant l(g)$, where $\lambda = \frac{1}{1-\epsilon}$. Thus $\frac{1}{\lambda} l_p(g) \leqslant l(g)$.

Conversely, by inductively applying A3, the triangle inequality, $l(g) \leqslant \sum_{i=0}^{k} l(y_i) = \sum_{i=0}^{k} l_p(y_i)$ for all sequences $y_0, \ldots, y_k \in B_K$ with $y_k \ldots y_0 = g$. Thus $l(g) \leqslant l_p(g)$.

We have $\frac{1}{\lambda} l_p(g) \leqslant l(g) \leqslant l_p(g)$, hence l_p is bi-Lipschitz equivalent to l with bi-Lipschitz constant $\lambda = \frac{1}{1-\epsilon}$.

Therefore, by Corollary 3.7, l lies in the bi-Lipschitz equivalence class of all based length functions arising from free, locally compact, co-compact metric G-graphs. □

Remark 4.10. A hyperbolic graph-like length function is a length function that satisfies the axioms from Definition 4.3 as well as the H_δ axiom from Definition 2.2 for some $\delta > 0$.

References

[1] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short. Notes on word hyperbolic groups. In *Group theory from a geometrical viewpoint. Proceedings of a workshop, held at the International Centre for Theoretical Physics in Trieste, Italy, 26 March to 6 April 1990*, pages 3–63. Singapore: World Scientific, 1991.

[2] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[3] I. M. Chiswell. Abstract length functions in groups. *Mathematical Proceedings of the Cambridge Philosophical Society*, 80:451–463, 1976.

[4] Marc Culler and John W. Morgan. Group actions on $\mathbb{R}$-trees. *Proceedings of the London Mathematical Society. Third Series*, 55:571–604, 1987.

[5] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones Mathematicae*, 84:91–119, 1986.

[6] R. C. Lyndon. Length functions in groups. *Mathematica Scandinavica*, 12:209–234, 1963.

[7] Walter Parry. Axioms for translation length functions. In *Arboreal group theory. Proceedings of a workshop, held September 13-16, 1988, in Berkeley, CA (USA)*, pages 295–330. New York etc.: Springer-Verlag, 1991.