

Train track maps on graphs of groups

Rylee Alanza Lyman

Abstract. In this paper we develop the theory of train track maps on graphs of groups. Expanding a definition of Bass, we define a notion of a map of a graph of groups, and of a homotopy equivalence. We prove that under one of two technical hypotheses, any homotopy equivalence of a graph of groups may be represented by a relative train track map. The first applies in particular to graphs of groups with finite edge groups, while the second applies in particular to certain generalized Baumslag–Solitar groups.

1. Introduction

A homotopy equivalence $f: G \rightarrow G$ of a connected graph G is a *train track map* when f maps vertices to vertices and the restriction of any iterate of f to an edge of G yields an immersion. (Relative) train track maps were introduced in [2]; they are perhaps the main tool for studying outer automorphisms of free groups.

A train track map $f: G \rightarrow G$ induces a well-defined outer automorphism of $\pi_1(G)$, a free group. The train track condition simplifies the analysis of the action of f on paths and loops in G . Choosing a basepoint \star in G and a path from \star to $f(\star)$ determines an automorphism $f_{\#}: \pi_1(G, \star) \rightarrow \pi_1(G, \star)$ and a lift of f to the universal covering tree Γ of G .

The map $\tilde{f}: \Gamma \rightarrow \Gamma$ is $f_{\#}$ -*twisted equivariant* in the sense that for $g \in \pi_1(G, \star)$ and $x \in \Gamma$, we have

$$\tilde{f}(g.x) = f_{\#}(g).\tilde{f}(x).$$

The lift $\tilde{f}: \Gamma \rightarrow \Gamma$ also satisfies the definition of a train track map. This formulation of train track maps as twisted equivariant maps of trees can be adapted in a straightforward way to automorphisms of groups acting on trees.

Theorem A (Tree version). *Suppose a group F acts cocompactly on a simplicial tree T , that $\Phi: F \rightarrow F$ is an automorphism, and that $\tilde{f}: T \rightarrow T$ is a Φ -twisted equivariant map. Suppose that one of the following conditions holds. Then there exists a Φ' -twisted equivariant relative train track map $\tilde{f}': T' \rightarrow T'$, where Φ and Φ' represent the same outer automorphism φ and where T and T' belong to the same deformation space \mathcal{D} in the sense of Forester [7, 10].*

- (1) Let T'' be a reduced tree in \mathcal{D} . Assume that edge stabilizers of T'' are finitely generated and that no iterate of Φ maps a generalized edge group of T'' properly into a conjugate of itself.
- (2) Assume that F is finitely generated, T is locally finite, and that the subgroup $\text{Mod}(\mathcal{D})$ of $\text{Out}(F)$ leaving \mathcal{D} invariant acts with finitely many orbits of cells on the deformation retract $P\mathcal{G}$ of \mathcal{D} considered in [10, Theorem 7.6] and [3].

If φ is irreducible, then the relative train track map constructed is a train track map.

Two F -trees T and T' belong to the same deformation space if and only if there exist F -equivariant maps $T \rightarrow T'$ and $T' \rightarrow T$ (see [10, Theorem 3.8]). A tree T is *reduced* if collapsing any orbit of edges of T yields a tree not in the same deformation space. A *generalized edge stabilizer* is a subgroup $H \leq F$ with the property that H contains the stabilizer of some edge \tilde{e} of T and is contained in the stabilizer of another edge \tilde{e}' of T . Each relative train track map $\tilde{f}: T \rightarrow T$ is a *morphism*: after F -equivariantly subdividing edges in the domain tree into finitely many edges, the map \tilde{f} becomes *simplicial*, in the sense that it maps edges to edges. Twisted equivariant morphisms of trees map edge stabilizers to generalized edge stabilizers.

Relative train track maps are defined in Section 4. This is not the first construction of relative train track maps on graphs of groups, see [4, 9, 15], but it is the most general, allowing in particular for infinite edge groups (but see [12] in the irreducible case). The proof of Theorem A relies on an algorithm of Bestvina–Handel [2], which requires a bound on the number of edges of G , the underlying graph of \mathcal{G} . Without further assumptions on \mathcal{G} , it appears at least possible that certain problematic valence-two vertices could proliferate in G , destroying any guarantee that the algorithm will terminate. It would be very interesting to have an example where this proliferation actually occurs. Perhaps an example could be found in considering the group

$$\langle a, s, t : tat^{-1} = sas^{-1} = a^2 \rangle.$$

Since one is primarily interested in using train track maps to study *outer* automorphisms, the choice of automorphism Φ in the statement of Theorem A is inconvenient. It would be more convenient to be able to work directly in the quotient graph of groups. This is the purpose of this paper.

Bass [1] defines a notion of a morphism of a graph of groups and proves that his morphisms of graphs of groups induce twisted equivariant simplicial maps of trees and vice versa. In Section 2, we offer an expanded definition of a *map* of a graph of groups and prove that our maps induce twisted equivariant maps of trees sending vertices to vertices and vice versa. We define *homotopy* of maps and when a map is a *homotopy equivalence*. The Bass–Serre trees of homotopy equivalent graphs of groups belong to the same deformation space and conversely.

Theorem A (Graph of groups version). *Let \mathcal{G} be a finite, connected graph of groups, φ be an outer automorphism of $\pi_1(\mathcal{G})$, and suppose that φ is induced by a map $f: \mathcal{G} \rightarrow \mathcal{G}$ satisfying one of the following conditions. Then there exists a relative train track map $f': \mathcal{G}' \rightarrow \mathcal{G}'$ representing φ on a graph of groups \mathcal{G}' homotopy equivalent to \mathcal{G} .*

- (1) *Let \mathcal{G}'' be a reduced graph of groups homotopy equivalent to \mathcal{G} . Assume that edge groups of \mathcal{G}'' are finitely generated and for some and hence every Φ representing φ , no iterate of the map Φ induces a proper inclusion of a generalized edge group of \mathcal{G} into itself.*
- (2) *Assume that vertex groups of \mathcal{G} are finitely generated and edge groups of \mathcal{G} have finite index in their incident vertex groups. Assume further that there are only finitely many isomorphism types of graphs of groups \mathcal{G}' homotopy equivalent to \mathcal{G} with the property that each edge e of \mathcal{G}' is surviving, in the sense that there is some reduced collapse of \mathcal{G}' in which the edge e is not collapsed.*

If φ is irreducible, then the relative train track map constructed is a train track map.

Item (1) of Theorem A applies in particular whenever edge groups are finite (or more generally when generalized edge groups are co-Hopfian), and thus to all accessible groups with infinitely many ends. Item (2) applies in particular to certain generalized Baumslag–Solitar groups.

The strategy of the proofs in this paper is to find the correct equivariant perspective so that the original arguments in [2] and [6] can be adapted without too much extra effort.

Here is the organization of the paper. We build up the aforementioned equivariant perspective in Section 2. The proof of Theorem A follows the outline in [2]; it occupies Sections 3 and 4.

2. Maps of graphs of groups

The purpose of this section is to define *maps* of graphs of groups and discuss their relationship with twisted equivariant maps of trees. *Morphisms* of graphs of groups were originally defined by Bass [1]. Our definition differs from his in two main respects: first, while his morphisms send edges to edges, our maps may send edges to edge paths (which may contain no edges), and second, we require our maps to respect basepoints. Assuming a map does not collapse edges to vertices, we call it a morphism, and it becomes a morphism in the sense of Bass after subdividing edges in the domain graph of groups into finitely many edges.

Let us take up the discussion from the introduction. Let F and F' be groups acting on simplicial trees Γ and Γ' , let $\Phi: F \rightarrow F'$ be a homomorphism, and suppose there is a (continuous) map $\tilde{f}: \Gamma \rightarrow \Gamma'$ which is Φ -twisted equivariant in the sense that for all $\tilde{x} \in \Gamma$ and all $g \in F$, we have

$$\tilde{f}(g.\tilde{x}) = \Phi(g).\tilde{f}(\tilde{x}).$$

An *equivariant homotopy* between two Φ -twisted equivariant maps \tilde{f} and \tilde{f}' is a homotopy $\tilde{f}_t: \Gamma \rightarrow \Gamma'$ with $\tilde{f}_0 = \tilde{f}$ and $\tilde{f}_1 = \tilde{f}'$ such that each map \tilde{f}_t is Φ -twisted equivariant. A Φ -twisted equivariant map $\tilde{f}: \Gamma \rightarrow \Gamma'$ is a *homotopy equivalence* if Φ is an isomorphism and there exists a Φ^{-1} -twisted equivariant map $\tilde{g}: \Gamma' \rightarrow \Gamma$ such that $\tilde{f}\tilde{g}$ and $\tilde{g}\tilde{f}$ are equivariantly homotopic to the identity. If we use Φ to identify F with F' , this says that Γ and Γ' belong to the same *deformation space* in the sense of [7, 10].

Each Φ -twisted equivariant map $\tilde{f}: \Gamma \rightarrow \Gamma'$ is equivariantly homotopic to a map $\tilde{f}': \Gamma \rightarrow \Gamma'$ which sends vertices to vertices and which has the property that for each edge \tilde{e} of Γ , either $\tilde{f}'(\tilde{e})$ is a vertex or after subdividing \tilde{e} into finitely many edges, the map \tilde{f}' restricted to the newly created edges is simplicial. We will work exclusively with such maps. We say the image of $\tilde{f}'(\tilde{e})$ is an *edge path* $\tilde{e}'_1 \dots \tilde{e}'_k$ in Γ' . In other words, if \tilde{f}' does not collapse edges, it is a Φ -twisted equivariant *morphism* of trees.

Quotient graph of groups. A *graph of groups* \mathcal{G} is a graph G (i.e. a 1-dimensional CW complex), which we usually assume to be connected, together with, for each edge e and vertex v of G , an assignment of groups \mathcal{G}_e and \mathcal{G}_v . For an oriented edge e with initial vertex v and terminal vertex w , there are injective homomorphisms $\iota_e: \mathcal{G}_e \rightarrow \mathcal{G}_v$ and $\iota_e: \mathcal{G}_e \rightarrow \mathcal{G}_w$, respectively. (We have $\mathcal{G}_{\bar{e}} = \mathcal{G}_e$.) The reader is referred to [1, 11, 13, 14] for additional background on graphs of groups, although we give a reasonably self-contained exposition of the aspects of the theory we will use.

Suppose F is a group acting on a simplicial tree Γ by simplicial automorphisms, and suppose that the action is *without inversions in edges*, i.e. that no group element sends an edge \tilde{e} to itself reversing orientation. This can always be arranged by passing to the barycentric subdivision of Γ . There is a *quotient graph of groups* \mathcal{G} , which we now describe.

Since the action of F on Γ is without inversions in edges, the quotient $F \backslash \Gamma$ naturally inherits the structure of a graph from Γ , call this graph G . The graph of groups structure on G depends on a choice of *fundamental domain* for the action of F on Γ ; we now describe how to choose the fundamental domain. Choose a spanning tree $S \subset G$ and lift S to $\tilde{S} \subset \Gamma$. For each edge $e \notin S$, choose a lift \tilde{e} in Γ such that one vertex of \tilde{e} belongs to \tilde{S} . Write T for the union of \tilde{S} with the (closed) edges \tilde{e} for $e \notin S$. For v a vertex of G , set \mathcal{G}_v to be the stabilizer of the unique preimage \tilde{v} of v in \tilde{S} . For e an edge of G , set \mathcal{G}_e to be the stabilizer of the unique preimage \tilde{e} of e in T . Let e be an oriented edge with terminal vertex v and write \tilde{x} for the terminal vertex of \tilde{e} . By definition, there is some group element $g_e \in F$ such that $g_e \cdot \tilde{x} = \tilde{v}$. If $\tilde{x} = \tilde{v}$, choose $g_e = 1 \in F$. If $h \in F$ stabilizes \tilde{e} , then $g_e h g_e^{-1}$ stabilizes \tilde{v} , so define $\iota_e: \mathcal{G}_e \rightarrow \mathcal{G}_v$ to be the map $h \mapsto g_e h g_e^{-1}$. This defines the graph of groups structure \mathcal{G} on the quotient graph G .

Graphs of spaces. Associated to a graph of groups \mathcal{G} with underlying graph G , we can build a *graph of spaces* $X_{\mathcal{G}}$. See [13] for more details. For a vertex v of G , take a connected CW complex X_v with one vertex \star_v which is a $K(\mathcal{G}_v, 1)$ and fix an identification $\pi_1(X_v, \star_v) = \mathcal{G}_v$. Do the same for each edge e of G , producing a CW complex X_e with one vertex \star_e and $\pi_1(X_e, \star_e) = \mathcal{G}_e$. Suppose the oriented edge e has initial

vertex v and terminal vertex w . Associated to the homomorphisms $\iota_e: \mathcal{G}_e \rightarrow \mathcal{G}_w$ and $\iota_{\bar{e}}: \mathcal{G}_e \rightarrow \mathcal{G}_v$, there are continuous, skeleta-preserving maps $i_e: (X_e, \star_e) \rightarrow (X_w, \star_w)$ and $i_{\bar{e}}: (X_e, \star_e) \rightarrow (X_v, \star_v)$ such that the induced maps on fundamental groups satisfy $(i_e)_\# = \iota_e$ and $(i_{\bar{e}})_\# = \iota_{\bar{e}}$. If V is the set of vertices of G and E is the set of oriented edges, the graph of spaces $X_{\mathcal{G}}$ is the quotient of the disjoint union

$$\coprod_{v \in V} X_v \amalg \coprod_{e \in E} (X_e \times [0, 1])$$

by the equivalence relation identifying $(x, 1) \in X_e \times [0, 1]$ with $i_e(x) \in X_w$, where w is the terminal vertex of the oriented edge e , and identifying $(x, t) \in X_e \times [0, 1]$ with $(x, 1 - t) \in X_{\bar{e}} \times [0, 1]$. Thus $X_{\mathcal{G}}$ is a CW complex. Note that after identifying each (open) edge of G with $(0, 1)$, there is a surjection $X_{\mathcal{G}} \rightarrow G$ whose fibers are naturally identified with the spaces X_v and X_e . By identifying G with the subspace of $X_{\mathcal{G}}$ comprising the points \star_v and (\star_e, t) for $t \in (0, 1)$, we can view the map $X_{\mathcal{G}} \rightarrow G$ as a retraction.

The fundamental group of a graph of groups. The *fundamental group of the graph of groups* $\pi_1(\mathcal{G})$ is the fundamental group of the graph of spaces $X_{\mathcal{G}}$. For convenience, choose a basepoint $p \in X_{\mathcal{G}}$ in the image of the retraction $X_{\mathcal{G}} \rightarrow G$. Each loop in $\pi_1(X_{\mathcal{G}}, p)$ is homotopic into the 1-skeleton of $X_{\mathcal{G}}$, and thus may be represented as an *edge path* γ of the form

$$\gamma = e'_1 g_1 e_2 g_2 \dots e_k g_k e'_{k+1}$$

where e_2, \dots, e_k are edges of G , e'_1 and e'_{k+1} terminal and initial segments of edges e_1 and e_{k+1} of G respectively, where the g_i for $1 \leq i \leq k$ are elements of $\pi_1(X_{v_i}, \star_{v_i}) = \mathcal{G}_{v_i}$, and where $v_i = \tau(e_i) = \tau(\bar{e}_{i+1})$. We allow the case where e'_1 and e'_{k+1} are empty, in which case they will be dropped from the notation. A path is *nontrivial* if it contains (a segment of) an edge.

Notice that under our identifications of $\pi_1(X_v, \star_v)$ with \mathcal{G}_v , the notion of an edge path in \mathcal{G} makes sense without reference to $X_{\mathcal{G}}$. Homotopy rel endpoints of paths in $X_{\mathcal{G}}$ yields a corresponding notion of homotopy rel endpoints for edge paths in \mathcal{G} . It is generated by replacing a segment of the form $e \iota_e(h)$ with $\iota_{\bar{e}}(h)e$, where e is an edge and $h \in \mathcal{G}_e$ is an element of the edge group, and by adding or removing segments of the form $e\bar{e}$ for an edge e . An edge path γ is *tight* if the number of edges in γ cannot be lessened by a homotopy.

Maps of graphs of groups. A *map of graphs of groups* is a pair of maps $f: G \rightarrow G'$ and $f_X: X_{\mathcal{G}} \rightarrow X_{\mathcal{G}'}$ such that the diagram

$$\begin{array}{ccc} X_{\mathcal{G}} & \xrightarrow{f_X} & X_{\mathcal{G}'} \\ \downarrow r & & \downarrow r' \\ G & \xrightarrow{f} & G' \end{array}$$

commutes, where r and r' are the retractions. A *homotopy of maps* is a pair of homotopies $f_{X,t}: X_{\mathcal{G}} \rightarrow X_{\mathcal{G}'}$ and $f_t: G \rightarrow G'$ such that for each t , the diagram of the form

groups $f': \mathcal{G} \rightarrow \mathcal{G}'$ to be the data of a map of graphs $f': G \rightarrow G'$ which either collapses edges to vertices or maps edges to edge paths, edge paths $f'(e) = g_0 e'_1 g_2 \dots e'_k g_k \in \mathcal{G}'$, and the homomorphisms $f'_v: \mathcal{G}_v \rightarrow \mathcal{G}'_{f(v)}$ and $f'_{e,e'_i}: \mathcal{G}_e \rightarrow \mathcal{G}'_{e'_i}$ subject to the *compatibility conditions* given by the above commutative diagrams. A map of graphs of groups is a *collapse map* if the map of graphs $f: G \rightarrow G'$ either sends edges to edges or collapses edges to vertices. It is a *morphism* if the map of graphs $f: G \rightarrow G'$ does not collapse edges. If $f: \mathcal{G} \rightarrow \mathcal{G}'$ is a homotopy equivalence, a morphism, and for each edge e of G , the edge path $f(e) = g_0 e_1 g_1 \dots e_k g_k$ is tight, then we say f is a *topological representative* of the induced outer automorphism $\varphi \in \text{Out}(\pi_1(\mathcal{G}))$.

Given a map of graphs of groups $f: \mathcal{G} \rightarrow \mathcal{G}'$, the rule

$$g \in \mathcal{G}_v \mapsto f_v(g) \in \mathcal{G}'_{f(v)} \quad \text{and} \quad e \mapsto f(e) = g_0 e'_1 g_1 \dots e'_k g_k$$

defines the action of f on paths in \mathcal{G} and a well-defined homomorphism $f_{\#}: \pi_1(X_{\mathcal{G}}, \star_w) \rightarrow \pi_1(X_{\mathcal{G}'}, \star_{f(w)})$. Since $X_{\mathcal{G}}$ is a $K(\pi_1(X_{\mathcal{G}}), 1)$, the homomorphism $f_{\#}$ is induced by a map $f_X: X_{\mathcal{G}} \rightarrow X_{\mathcal{G}'}$. We claim that we can choose f_X and f within their homotopy classes so that the pair (f_X, f) defines a map of graphs of groups. Since we have little further need of $X_{\mathcal{G}}$, we leave the details to the interested reader.

The following operations preserve the homomorphism $f_{\#}: \pi_1(\mathcal{G}, w) \rightarrow \pi_1(\mathcal{G}', f(w))$, and we claim that it follows that the resulting maps of graphs of groups are homotopic, but again we leave the demonstration to the reader.

- (1) Replace the edge path $f(e) \in \mathcal{G}'$ with a path which is homotopic to it rel endpoints and replace the homomorphisms f_{e,e'_i} with new homomorphisms so that the compatibility conditions above are still satisfied. Note that since f_v and f_w are unchanged under this operation, the images of the new edge group homomorphisms must be abstractly isomorphic to the original f_{e,e'_i} . This restricts the possible homotopies we can perform, but does not prevent us from tightening $f(e)$ if it is not already tight.
- (2) Suppose that $v \neq w$ is a vertex of G and that $g \in \mathcal{G}'_{f(v)}$. Replace f_v with $\text{ad}(g)f_v$ and for each oriented edge e of G with initial vertex v , replace the edge path $f(e)$ with $gf(e)$.
- (3) Suppose that $v \neq w$ is a vertex of G , that e' is an edge of G' with terminal vertex $f(v)$ and initial vertex v' such that $\iota_{e'}(\mathcal{G}'_{e'})$ contains the image of f_v . Change the map f of graphs by a homotopy supported in a neighborhood of v by pulling the image of v across \bar{e}' so that the new map called f' satisfies $f'(v) = v'$. The new map f'_v is $\iota_{\bar{e}'} \iota_{e'}^{-1} f_v$. That is to say, the map f'_v is accomplished by viewing $f_v(\mathcal{G}_v)$ as a subgroup of $\mathcal{G}'_{e'}$, then mapping it to $\mathcal{G}'_{v'}$. For each oriented edge e of G with initial vertex v , we have the new edge path $f'(e) = e' f(e)$. Add the homomorphism $f'_{e,e'}: \mathcal{G}_e \rightarrow \mathcal{G}'_{e'}$ defined as $f'_{e,e'}(h) = \iota_{e'}^{-1} f_v \iota_{\bar{e}'}(h)$. Observe that the compatibility conditions are satisfied.

If, ignoring the stipulations around the basepoint w , the map $f: \mathcal{G} \rightarrow \mathcal{G}'$ can be transformed to a map $f': \mathcal{G} \rightarrow \mathcal{G}'$ by a finite number of the above operations, we will say that f and $f': \mathcal{G} \rightarrow \mathcal{G}'$ are *homotopic*.

The Bass–Serre tree. Let \mathcal{G} be a graph of groups and let $p \in G$ be a basepoint. For v a vertex of G , write $[p, v]$ for the set of homotopy classes of paths in \mathcal{G} from p to v . The group \mathcal{G}_v acts on $[p, v]$ on the right: an element g sends the homotopy class $[\gamma]$ of the path γ to the homotopy class of the composite path $[\gamma g]$. Let V denote the set of vertices of G . The set

$$\coprod_{v \in V} [p, v] / \mathcal{G}_v$$

forms the vertex set of a graph, where two elements $[\gamma] \mathcal{G}_v$ and $[\gamma'] \mathcal{G}_w$ are adjacent if the path $\bar{\gamma} \gamma'$ is homotopic to a path of length one. The fundamental group $\pi_1(\mathcal{G}, p)$ acts naturally on this graph; for example the homotopy class of a loop $[\sigma]$ sends the vertex $[\gamma] \mathcal{G}_v$ to the vertex $[\sigma \gamma] \mathcal{G}_v$. The fundamental theorem of Bass–Serre theory asserts that this graph is a tree (see [1, Theorem 1.17]) and that the quotient graph of groups of the action of $\pi_1(\mathcal{G}, p)$ on this tree may be identified with \mathcal{G} (see [1, Corollary 3.7] and [14, Chapter I, Theorem 13]). (Let us remark that non-vertex points of the Bass–Serre tree may be identified with homotopy classes of paths in \mathcal{G} that do not end at vertices.)

Indeed, suppose we began with a group F acting on a tree Γ with quotient graph of groups \mathcal{G} defined relative to the choice of spanning tree S containing p , of fundamental domain T and of group elements g_e in F . For v a vertex of G , let γ_v denote the unique tight path in S as an ordinary graph from p to v . Then for e an edge of G and $g \in \mathcal{G}_v$ an element of a vertex group, thought of as the stabilizer of \tilde{v} in F , the map

$$[\gamma_{\tau(\tilde{e})} e \bar{\gamma}_{\tau(\tilde{e})}] \mapsto g_e \quad \text{and} \quad [\gamma_v g \bar{\gamma}_v] \mapsto g$$

defines an isomorphism $\Phi: \pi_1(\mathcal{G}, p) \rightarrow F$. For v a vertex of \mathcal{G} , let \tilde{v} be the unique pre-image of v in \tilde{S} . For \tilde{e} an edge of $T \setminus \tilde{S}$, orient \tilde{e} so that its initial vertex \tilde{w} is in \tilde{S} and its terminal vertex \tilde{w}' is not. The smallest subtree of the Bass–Serre tree containing the vertices

$$[\gamma_v] \mathcal{G}_v \quad \text{and} \quad [\gamma_w e] \mathcal{G}_{w'}$$

is a fundamental domain for the action of $\pi_1(\mathcal{G}, p)$ on its Bass–Serre tree, and the map

$$[\gamma_v] \mathcal{G}_v \mapsto \tilde{v} \quad \text{and} \quad [\gamma_w e] \mathcal{G}_{w'} \mapsto \tilde{w}'$$

extends to a Φ -twisted equivariant simplicial isomorphism between the Bass–Serre tree of \mathcal{G} and Γ taking this fundamental domain to T .

We will use Φ to identify $\pi_1(\mathcal{G}, p)$ with F and this Φ -twisted equivariant simplicial isomorphism to identify Γ with the Bass–Serre tree of \mathcal{G} . Write $\pi: \Gamma \rightarrow G$ for the natural projection.

Normal forms, projecting and lifting paths. Suppose γ is a tight path in \mathcal{G} from a point x in the interior of an edge to a point y , and that $\tilde{x} = [\sigma]$ is a lift of x to Γ . Then the unique tight path $\tilde{\gamma}$ from $\tilde{x} = [\sigma]$ to $[\sigma\gamma]$ (or $[\sigma\gamma]\mathcal{G}_y$ if y is a vertex) is a tight path in Γ which *lifts* γ in the sense that $\pi(\tilde{\gamma})$ has the same underlying edge path in G as γ . If x is a vertex, then for each $g \in \mathcal{G}_x$, there is a unique tight path $\tilde{\gamma}$ from $\tilde{x} = [\sigma]\mathcal{G}_x$ to $[\sigma g\gamma]$ (or $[\sigma g\gamma]\mathcal{G}_y$ if y is a vertex) which lifts γ . If γ is *not* a tight path, we may decompose it into a concatenation of tight paths and successively lift those paths to obtain a lift of γ . We would like to define a notion of *projecting* paths from Γ to \mathcal{G} so that lifting and projecting are (nearly) inverse operations.

Recall [1, §1.12] that given a choice, for each oriented edge e with terminal vertex v , of a set S_e of left coset representatives for $\mathcal{G}_v/\iota_e(\mathcal{G}_e)$ containing $1 \in \mathcal{G}_v$, there is a *normal form* for edge paths in \mathcal{G} , where an edge path

$$\gamma = e'_1 g_1 e_2 \dots g_{k-1} e_k g_k e'_{k+1}$$

is in normal form if it is tight and each g_i for $1 \leq i \leq k - 1$ belongs to $S_{\tilde{e}_{i+1}}$. A path can be inductively put in normal form by performing homotopies from “left to right.”

Suppose \tilde{x} is a point $[\sigma]$ (or $[\sigma]\mathcal{G}_x$ if \tilde{x} is a vertex) of Γ and that \tilde{y} is a point $[\eta]$ (or $[\eta]\mathcal{G}_y$ if \tilde{y} is a vertex) of Γ , and let $\tilde{\gamma}$ be the unique tight path from \tilde{x} to \tilde{y} . If \tilde{y} is a vertex, choose η in its \mathcal{G}_y -orbit so that its final vertex group element is 1. Let $\gamma = \pi(\tilde{\gamma})$ be the unique path in normal form homotopic to $\tilde{\sigma}\eta$, thus γ has the form

$$\gamma = e'_1 g_1 e_2 \dots g_{k-1} e_k 1$$

where each g_i for $1 \leq i \leq k - 1$ belongs to $S_{\tilde{e}_{i+1}}$. One checks that if \tilde{x} is not a vertex, then the lift of γ defined above is $\tilde{\gamma}$. If $\tilde{x} = [\sigma]\mathcal{G}_x$ is a vertex and we chose σ in its \mathcal{G}_x -orbit so that its final vertex group element is 1, then the lift of γ defined above corresponding to the choice of $1 \in \mathcal{G}_x$ is $\tilde{\gamma}$.

Conversely, if we first lift the normal-form path γ to $\tilde{\gamma}$ and then project to $\pi(\tilde{\gamma})$ as in the previous paragraph, then $\pi(\tilde{\gamma})$ and γ agree except for the final vertex group element if y is a vertex.

Proposition 2.1 (Lifting maps, cf. [1, Proposition 2.4]). *Suppose that $f : (\mathcal{G}, v) \rightarrow (\mathcal{G}', v')$ is a map of graphs of groups and that (Γ, \tilde{v}) and (Γ', \tilde{v}') are the Bass–Serre trees of \mathcal{G} and \mathcal{G}' respectively. There exists an $f_{\#}$ -twisted equivariant map $\tilde{f} : (\Gamma, \tilde{v}) \rightarrow (\Gamma', \tilde{v}')$ such that the following diagram commutes as maps of underlying graphs:*

$$\begin{array}{ccc} (\Gamma, \tilde{v}) & \xrightarrow{\tilde{f}} & (\Gamma', \tilde{v}') \\ \downarrow \pi & & \downarrow \pi' \\ (\mathcal{G}, v) & \xrightarrow{f} & (\mathcal{G}', v'), \end{array}$$

where π and π' are the natural projections. Furthermore, the normal form of the path $f\pi(\tilde{\gamma})$ agrees with the path $\pi'f(\tilde{\gamma})$ except for the final vertex group element if $\tilde{\gamma}$ ends at a vertex. If the maps f and f' are homotopic fixing v , then the lifted maps \tilde{f} and \tilde{f}' are equivariantly homotopic.

It follows from the final claim that if f is a homotopy equivalence, then \tilde{f} is too.

Proof. A point $\tilde{p} \in \Gamma$ corresponds to a homotopy class $[\gamma]$ of paths from v to $p = \pi(\tilde{p})$ (or to $[\gamma]\mathcal{G}_p$ if p is a vertex). Define $\tilde{f}(\tilde{p}) = [f(\gamma)]$ (or $[f(\gamma)]\mathcal{G}'_{f(p)}$ if $f(p)$ is a vertex). A simple calculation shows that \tilde{f} is $f_{\#}$ -twisted equivariant and that the claimed properties hold. ■

Proposition 2.2 (Projecting maps, cf. [1, §§4.1–4.5]). *Suppose F and F' are groups acting on trees Γ and Γ' respectively with quotient graphs of groups \mathcal{G} and \mathcal{G}' . Let $\Phi: F \rightarrow F'$ be a homomorphism and let $\tilde{f}: (\Gamma, \tilde{v}) \rightarrow (\Gamma', \tilde{v}')$ be a Φ -twisted equivariant map of trees. There is a map $f: (\mathcal{G}, v) \rightarrow (\mathcal{G}', v')$ of graphs of groups such that the following diagram commutes as maps of underlying graphs:*

$$\begin{CD} (\Gamma, \tilde{v}) @>\tilde{f}>> (\Gamma', \tilde{v}') \\ @VV\pi V @VV\pi' V \\ (\mathcal{G}, v) @.-.-f.-.-> (\mathcal{G}', v'), \end{CD}$$

where π and π' are the natural projections. Furthermore, the normal form of the path $f\pi(\tilde{\gamma})$ agrees with the path $\pi'\tilde{f}(\tilde{\gamma})$ except for the final vertex group element if $\tilde{\gamma}$ ends at a vertex. If two maps \tilde{f} and \tilde{f}' are Φ -twisted equivariantly homotopic, then f and f' are homotopic.

Proof. By Φ -twisted equivariance, the map $\pi'\tilde{f}$ yields a well-defined map on $\pi_1(\mathcal{G}, v)$ -orbits; this is the map $f: G \rightarrow G'$ as a map of graphs.

Let $T \subset \Gamma$ and $T' \subset \Gamma'$ be fundamental domains containing \tilde{v} and \tilde{v}' respectively and let \mathcal{G} and \mathcal{G}' be the quotient graphs of groups associated to T and T' respectively. Each edge $e \in G$ and $e' \in G'$ has a single preimage $\tilde{e} \in T$ and $\tilde{e}' \in T'$ respectively. The groups $\mathcal{G}_{\tilde{e}}$ and $\mathcal{G}'_{\tilde{e}'}$ are the stabilizers of \tilde{e} and \tilde{e}' respectively. Each vertex $v \in G$ and $v' \in G'$ has a preferred preimage $\tilde{v} \in T$ and $\tilde{v}' \in T'$ respectively. The groups $\mathcal{G}_{\tilde{v}}$ and $\mathcal{G}'_{\tilde{v}'}$ are the stabilizers of \tilde{v} and \tilde{v}' respectively. For each oriented edge e of G and e' of G' there are elements $g_e \in F$ and $g'_{e'} \in F'$ such that the monomorphisms ι_e and $\iota_{e'}$ are the restrictions of the maps

$$h \mapsto g_e h g_e^{-1} \quad \text{and} \quad h' \mapsto g'_{e'} h' g'_{e'}^{-1}$$

to \mathcal{G}_e and $\mathcal{G}'_{e'}$ respectively.

Let v be a vertex of G and write $w = f(v)$. To define $f_v: \mathcal{G}_v \rightarrow \mathcal{G}'_w$, note that the stabilizers of $\tilde{f}(\tilde{v})$ and \tilde{w} are conjugate in F' by some element t_w such that $t_w \cdot \tilde{f}(\tilde{v}) = \tilde{w}$. The restriction of

$$h \mapsto t_w \Phi(h) t_w^{-1}$$

to \mathcal{G}_v defines a homomorphism $f_v: \mathcal{G}_v \rightarrow \mathcal{G}'_w$.

Let e be an oriented edge of G with initial vertex v and terminal vertex w and suppose first that $f(e)$ is a vertex. Then for $h \in \mathcal{G}_e$, we have

$$\begin{aligned} f_v \iota_{\bar{e}}(h) &= t_{f(v)} \Phi(g_{\bar{e}} h g_{\bar{e}}^{-1}) t_{f(v)}^{-1}, \\ f_w \iota_e(h) &= t_{f(w)} \Phi(g_e h g_e^{-1}) t_{f(w)}^{-1}. \end{aligned}$$

Therefore the compatibility conditions force us to define $f(e) = g_0$, where

$$g_0 = t_{f(v)} \Phi(g_{\bar{e}} g_e^{-1}) t_{f(w)}^{-1}.$$

One checks that g_0 belongs to $\mathcal{G}_{f(v)} = \mathcal{G}_{f(w)}$.

Now suppose that

$$\tilde{f}(\tilde{e}) = \tilde{E}'_1 \dots \tilde{E}'_k.$$

Φ -twisted equivariance implies that if $h \in \mathcal{G}_e$, then $\Phi(h)$ stabilizes \tilde{E}'_i for $1 \leq i \leq k$. Suppose $\pi'(\tilde{E}'_i) = e'_i$ in G' . These stabilizers are conjugate to $\mathcal{G}'_{e'_i}$ by elements $t_{e'_i} \in F'$ such that $t_{e'_i} \cdot \tilde{E}'_i = \tilde{e}'_i$. The restriction of

$$h \mapsto t_{e'_i} \Phi(h) t_{e'_i}^{-1}$$

to \mathcal{G}_e defines a homomorphism $f_{e,e'_i}: \mathcal{G}_e \rightarrow \mathcal{G}'_{e'_i}$. Given $h \in \mathcal{G}_e$, we have

$$f_v \iota_{\bar{e}}(h) = t_{f(v)} \Phi(g_{\bar{e}} h g_{\bar{e}}^{-1}) t_{f(v)}^{-1}.$$

On the other hand, we have

$$\iota_{\tilde{e}'_1} f_{e,e'_1}(h) = g'_{\tilde{e}'_1} t_{e'_1} \Phi(h) t_{e'_1}^{-1} g_{\tilde{e}'_1}^{-1}.$$

The compatibility conditions force us to define

$$g_0 = t_{f(v)} \Phi(g_{\bar{e}}) t_{e'_1}^{-1} g_{\tilde{e}'_1}^{-1}.$$

One checks that we have $g_0 \in \mathcal{G}_{f(v)}$. The derivations of the other vertex group elements are similar. The additional claims are straightforward to check; we leave them to the reader. This completes the proof. ■

Given a vertex $v \in G$, let $\text{st}(v)$ denote the set of oriented edges e of G with initial vertex v . Recall that for each lift $\tilde{v} \in \pi^{-1}(v)$ there is a \mathcal{G}_v -equivariant bijection

$$\text{st}(\tilde{v}) = \coprod_{e \in \text{st}(v)} \mathcal{G}_v / \iota_e(\mathcal{G}_e) \times \{e\}.$$

Suppose $f: \mathcal{G} \rightarrow \mathcal{G}$ is a homotopy equivalence. It defines an outer automorphism of $\pi_1(\mathcal{G})$. Choosing a basepoint p in G and an edge path σ in \mathcal{G} from p to $q = f(p)$ defines an automorphism $f_{\#}: \pi_1(\mathcal{G}, p) \rightarrow \pi_1(\mathcal{G}, p)$ defined as

$$f_{\#}([\gamma]) = [\sigma f(\gamma) \bar{\sigma}].$$

Let (Γ, \tilde{p}) be the Bass–Serre tree covering (\mathcal{G}, p) , and let \tilde{q} be the endpoint of the lift of σ to Γ beginning at \tilde{p} . By Proposition 2.1, there is an $f_{\#}$ -twisted equivariant map

$\tilde{f}: (\Gamma, \tilde{p}) \rightarrow (\Gamma, \tilde{q})$ taking the point $\tilde{x} = [\gamma]$ (or $\tilde{x} = [\gamma]\mathcal{G}_x$ if x is a vertex) to the point $\tilde{f}(\tilde{x}) = [\sigma f(\gamma)]$ (or $\tilde{f}(\tilde{x}) = [\sigma f(\tilde{\gamma})]\mathcal{G}_{f(x)}$ if $f(x)$ is a vertex).

Conversely, we saw earlier in the definition of a map of graphs of groups and Proposition 2.2 that any Φ -twisted equivariant map $\tilde{f}: \Gamma \rightarrow \Gamma'$ is equivariantly homotopic to a map which projects to a map of graphs of groups. It is not quite true that if Φ is an automorphism, then \tilde{f} projects to a homotopy equivalence: for that one needs the existence of a Φ -twisted equivariant homotopy inverse for \tilde{f} .

Topological representatives. Recall that we defined a *topological representative* to be a homotopy equivalence $f: \mathcal{G} \rightarrow \mathcal{G}$ which is a morphism and which sends edges to non-trivial tight edge paths. In the following sections, we prove Theorem A by performing a number of operations on topological representatives.

Given a graph of groups \mathcal{G} , the collection of outer automorphisms of $F = \pi_1(\mathcal{G})$ that admit a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ forms a subgroup of $\text{Out}(F)$. It is precisely the subgroup of $\text{Out}(F)$ leaving invariant the deformation space \mathcal{D} to which the Bass–Serre tree of \mathcal{G} belongs. We suggest the name *modular group* or *mapping class group* of \mathcal{G} for this group, and in this paper we will write it as $\text{Mod}(\mathcal{D})$ or $\text{Mod}(\mathcal{G})$. An outer automorphism $\varphi \in \text{Out}(F)$ belongs to the “modular group” of \mathcal{G} if for any conjugacy class $[g]$ in F , the conjugacy class $\varphi([g])$ is elliptic in the Bass–Serre tree for \mathcal{G} if and only if $[g]$ is elliptic. In some cases this subgroup is all of $\text{Out}(F)$. One case where this happens is when F is virtually free and vertex groups of \mathcal{G} are finite. Another is when F is a *generalized Baumslag–Solitar group* and vertex and edge groups of \mathcal{G} are infinite cyclic.

Another case where this “modular group” is all of $\text{Out}(F)$ is when $F = A_1 * \dots * A_n * F_k$ is the Grushko decomposition of a finitely generated group. That is, the A_i are freely indecomposable and not infinite cyclic and F_k is a free group of rank k . The graph of groups \mathcal{G} can be any graph of groups with trivial edge groups, vertex groups the A_i and ordinary fundamental group free of rank k . For example, \mathcal{G} may be the *thistle with n prickles and k petals*. This is a graph of groups with one vertex \star with trivial vertex group, n vertices with vertex group each of the A_i , and $n + k$ edges. The first n edges connect vertices with nontrivial vertex group to \star , and the remaining k edges form loops based at \star .

Example 2.3. Consider

$$F = C_2 * C_2 * C_2 * C_2 = \langle a, b, c, d \mid a^2 = b^2 = c^2 = d^2 = 1 \rangle,$$

the free product of four copies of the cyclic group of order two. Let $\Phi: F \rightarrow F$ be the automorphism

$$\Phi \begin{cases} a \mapsto b \\ b \mapsto c \\ c \mapsto d \\ d \mapsto cbdadbc. \end{cases}$$




Figure 1. The topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$.

(Notice that, e.g., $c^{-1} = c$.) A topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ of Φ on the thistle with four prickles is depicted in Figure 1. The maps on vertex groups are the unique isomorphisms.

3. Train track maps

The purpose of this section is to prove the irreducible case of Theorem A. The strategy is a straightforward adaptation of the arguments of [2, Section 1] to graphs of groups. At the end of the section we prove a proposition characterizing irreducibility for outer automorphisms of free products.

Fix once and for all a graph of groups \mathbb{G} . A *marked graph of groups* is a graph of groups \mathcal{G} together with a homotopy equivalence $\sigma: \mathbb{G} \rightarrow \mathcal{G}$. In the language of [10], the marking keeps track not only of the fundamental group of \mathcal{G} , but constrains the *deformation space* to which it belongs. We will assume that \mathbb{G} is *reduced* in the sense of [7], i.e. there is no homotopy equivalence that collapses an edge of \mathbb{G} .

Given a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ and an ordering e_1, \dots, e_m of the edges of G , there is an associated $m \times m$ *transition matrix* M with ij th entry counting the number of times the f -image of the j th edge crosses the i th edge in either direction. The map f is *irreducible* if the matrix M is irreducible. Recall that a matrix is *irreducible* if for $1 \leq i, j \leq m$, there is an integer ℓ such that the ij th entry of M^ℓ is positive. Associated to every irreducible matrix is its *Perron–Frobenius eigenvalue* $\lambda \geq 1$. An irreducible matrix with Perron–Frobenius eigenvalue $\lambda = 1$ is a transitive permutation matrix. The transition matrix of Example 2.3 is

$$\begin{pmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 2 \\ 0 & 1 & 0 & 2 \\ 0 & 0 & 1 & 2 \end{pmatrix},$$

which is irreducible and for which the Perron–Frobenius eigenvalue is the largest real root of the polynomial $x^4 - 2x^3 - 2x^2 - 2x - 1$ and satisfies $\lambda \approx 2.948$.

Call a vertex v of \mathcal{G} *inessential* if for some oriented edge e with terminal vertex v the homomorphism $\iota_e: \mathcal{G}_e \rightarrow \mathcal{G}_v$ is surjective.

A subgraph G_0 of G is *invariant* with respect to a map $f: \mathcal{G} \rightarrow \mathcal{G}$ if $f(G_0) \subset G_0$. It is a *forest* if each component C of G_0 is a tree and in the induced graph of groups structure we have that $\pi_1(\mathcal{G}|_C)$ acts with global fixed point on its Bass–Serre tree. In C , this means there is a choice of vertex v in C and an orientation of each edge e of C toward this vertex such that each homomorphism ι_e away from v is surjective. A forest is *nontrivial* if it contains at least one edge. An outer automorphism $\varphi \in \text{Out}(\pi_1(\mathcal{G}))$ is *irreducible* if it admits a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ (i.e. we have $\varphi \in \text{Mod}(\mathcal{G})$) and if whenever \mathcal{G} has no inessential valence-one vertices and no nontrivial invariant forests, then the topological representative f is irreducible.

A homotopy equivalence $f: \mathcal{G} \rightarrow \mathcal{G}$ (taking vertices to vertices) is *tight* if for each edge e , either $f(e)$ is a tight edge path, or $f(e)$ is a vertex. A homotopy equivalence may be *tightened* to a tight homotopy equivalence by a homotopy relative to the vertices of G . In the language of the previous section, the homotopy only involves the first operation in the definition of homotopy of maps of graphs of groups.

Suppose $f: \mathcal{G} \rightarrow \mathcal{G}$ is a tight homotopy equivalence. A forest $G_0 \subset \mathcal{G}$ is *pretrivial* if each edge in the forest is eventually mapped to a point. *Maximal* (with respect to inclusion) pretrivial forests are in particular invariant.

Lemma 3.1 ([2, p. 7]). *If $f: \mathcal{G} \rightarrow \mathcal{G}$ is a tight homotopy equivalence, collapsing a maximal pretrivial forest in \mathcal{G} produces a topological representative $f': \mathcal{G}' \rightarrow \mathcal{G}'$. If instead $f: \mathcal{G} \rightarrow \mathcal{G}$ is a topological representative of an irreducible outer automorphism and G has no inessential valence-one vertices, collapsing a maximal invariant forest yields an irreducible topological representative $f': \mathcal{G}' \rightarrow \mathcal{G}'$.*

Proof. We describe how to collapse invariant forests.

If $f: \mathcal{G} \rightarrow \mathcal{G}$ is a tight homotopy equivalence and $G_0 \subset G$ is an invariant forest, define $\mathcal{G}_1 = \mathcal{G}/G_0$ to be the quotient graph of groups obtained by collapsing each component C of G_0 to a vertex. The vertex group of the vertex determined by C is the fundamental group $\pi_1(\mathcal{G}|_C, p_C)$ with respect to some basepoint $p_C \in C$. Since G_0 is a forest, this fundamental group is equal to some vertex group in C . Choose p_C equal to that vertex; choose arbitrarily if there are multiple choices. Given a vertex v of C , let γ_v be the unique tight path without vertex group elements from v to p_C . If v does not belong to any component of G_0 , let γ_v be the trivial path (without vertex group elements). Let $\pi: \mathcal{G} \rightarrow \mathcal{G}_1$ be the quotient map. It is a collapse map of graphs of groups.

Each edge e of G_1 has a unique preimage in G ; abusing notation, call it e as well. Define $f_1(e) = \pi f(\bar{\gamma}_{\tau(\bar{e})} e \gamma_{\tau(e)})$. If v in G_1 is a vertex, then v either corresponds to a unique vertex of G , call it v as well, or to a component C of G_0 . In the former case define $(f_1)_v = f_v$. Suppose in the latter case that f maps the component C to C' . Let $\gamma_{p'_C}$ be the unique tight path in C' without vertex group elements from $f(p_C)$ to p'_C ; this determines a map $f_{\sharp}: \pi_1(\mathcal{G}|_C, p_C) \rightarrow \pi_1(\mathcal{G}|_{C'}, p_{C'})$; this is the map $(f_1)_v$ in this case.

If $e \subset G$ is an edge not in G_0 , then the edge path for $f_1(e)$ is obtained from $f(e)$ by deleting all occurrences of edges in G_0 and possibly adding vertex group elements at the

ends. Since f was tight, if $e\sigma\bar{e}$ is a subpath of the f -image of some edge e' not in G_0 , where σ is a nontrivial path in G_0 , then σ must be homotopic to a path of the form $\sigma'g\bar{\sigma}'$ for some path σ' in G_0 without vertex group elements and g an element of some vertex group. In $f_1(e')$, the path $e\sigma\bar{e}$ is replaced by $eg\bar{e}$. This implies that $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is tight. Moreover, if G_0 was a maximal pretrivial forest, then f_1 is a topological representative.

If instead f was a topological representative of an irreducible outer automorphism, then the transition matrix for $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is obtained from the transition matrix for $f: \mathcal{G} \rightarrow \mathcal{G}$ by deleting the rows and columns associated to the edges of G_0 . If in this latter case G had no inessential valence-one vertices, then neither does G_1 , and since G_0 was a maximal invariant forest, G_1 has no invariant forests; therefore $f: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ must be irreducible. ■

Recall we write $\text{st}(v)$ for the set of oriented edges e with initial vertex v . A *direction* at v is an element of the set

$$\coprod_{e \in \text{st}(v)} \mathcal{G}_v / \iota_e(\mathcal{G}_e) \times \{e\}.$$

A *turn* at v is a pair of directions at v . If $f: \mathcal{G} \rightarrow \mathcal{G}$ is a topological representative, then f determines a map Df on directions sending a direction based at v to a direction based at $f(v)$ via the rule

$$([g], e) \mapsto ([f_v(g)g_0], e_1),$$

where the edge path $f(e)$ begins with g_0e_1 . The compatibility condition ensures that this map is well-defined; we have

$$f_v(g\iota_e(h))g_0 = f_v(g)f_v(\iota_e(h))g_0 = f_v(g)g_0\iota_{\bar{e}_1}(f_{e,e_1}(h)).$$

The map Df induces a map on turns, which we also denote by Df . In Example 2.3, the vertex \star is mapped to itself by f ; the restriction of Df to \star is determined by the dynamical system $\bar{e}_1 \mapsto \bar{e}_2 \mapsto \bar{e}_3 \leftrightarrow \bar{e}_4$.

A turn is *degenerate* if it consists of a pair of identical elements and is *nondegenerate* otherwise. A turn is *illegal* with respect to a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ if its image under some iterate of Df is degenerate and is *legal* otherwise. In Example 2.3, a turn $\{\bar{e}_i, \bar{e}_j\}$ based at \star is illegal if i and j are equal mod 2, and is legal otherwise.

Consider the edge path

$$\gamma = g_1e_1g_2e_2 \dots e_kg_{k+1}.$$

We say γ takes the turns $\{([1], \bar{e}_i), ([g_{i+1}], e_{i+1})\}$. The path γ is *legal* if it takes only legal turns.

A topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ is a *train track map* if $f(e)$ is a legal path for each edge e of Γ . Equivalently, f is a train track map if for each $k \geq 1$ and each edge e of Γ , we have that $f^k(e)$ is a tight edge path. In Example 2.3, f is not a train track map because the image of e_4 takes the illegal turn $\{\bar{e}_4, \bar{e}_2\}$.

Example 2.3 continued. Let us fold f at the illegal turn $\{\bar{e}_2, \bar{e}_4\}$. To do this, first subdivide e_4 at the preimage of the vertex with vertex group $\langle c \rangle$ so e_4 becomes the edge path $e'_4 e''_4$ and identify e'_4 with e_2 . The action of the resulting map $f': \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is obtained from f by replacing instances of e_4 with $e'_4 e_2$. Thus we have

$$f'(e_4) = e_1 \bar{e}_2 \bar{e}'_4 d e'_4 e_2 \bar{e}_2 b e_2 \bar{e}_3 c.$$

Tighten f' by a homotopy with support on e'_4 to remove $e_2 \bar{e}_2$, yielding an irreducible topological representative $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$. See Figure 2.




Figure 2. The topological representative $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$.

The Perron–Frobenius eigenvalue λ_1 for $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is the largest real root of the polynomial $x^4 - 2x^3 - 2x^2 + x - 1$ and satisfies $\lambda_1 \approx 2.663$; thus $\lambda_1 < \lambda$. However, f_1 is still not a train track map: Df_1 sends the turn $\{(1, \bar{e}'_4), (b, e_2)\}$, which is crossed by $f_1(e'_4)$, to $\{(c, e_3), (c, e_3)\}$; thus this turn is illegal. We cannot quite fold e_2 and the end of \bar{e}'_4 because the f_1 -image of the latter ends with $\bar{e}_3 c$. Lifting to the Bass–Serre tree $\tilde{f}_1: \Gamma_1 \rightarrow \Gamma_1$, it is not the edge \bar{e}'_4 which is folded with \bar{e}_2 but $b \cdot \bar{e}'_4$. We may remedy the situation by changing the fundamental domain in Γ_1 , or equivalently by changing the marking on \mathcal{G}_1 by twisting the edge e'_4 by $b^{-1} = b$. This replaces $\langle d \rangle$ with $\langle bdb \rangle$, replaces $f_1(e_3)$ with $e'_4 b e_2$ and replaces $f_1(e'_4)$ with $e_1 \bar{e}_2 b \bar{e}'_4 d e'_4 e_2 \bar{e}_3$. Then we fold e'_4 and \bar{e}_2 . The resulting graph of groups \mathcal{G}_2 is abstractly isomorphic to our original graph of groups \mathcal{G} , but the marking differs. The action of the resulting map $f'': \mathcal{G}_2 \rightarrow \mathcal{G}_2$ on edges is obtained by replacing instances of e'_4 with $e''_4 \bar{e}_2$. Thus we have

$$f''(e''_4) = e_1 \bar{e}_2 b e_2 e''_4 b d b e''_4 \bar{e}_2 e_2,$$

and we may tighten to produce an irreducible topological representative $f_2: \mathcal{G}_2 \rightarrow \mathcal{G}_2$. See Figure 3. The Perron–Frobenius eigenvalue λ_2 is the largest real root of $x^4 - 2x^3 - 2x^2 + 2x - 1$ and satisfies $\lambda_2 \approx 2.539$; thus $\lambda_2 < \lambda_1$. The restriction of Df_2 to turns incident to \star is determined by the dynamical system $\bar{e}_1 \mapsto \bar{e}_2 \leftrightarrow \bar{e}_3, \bar{e}_4 \mapsto \bar{e}_4$. The only illegal turn in \mathcal{G}_2 is $\{\bar{e}_1, \bar{e}_3\}$, which is not crossed by the f_2 -image of any edge, so $f_2: \mathcal{G}_2 \rightarrow \mathcal{G}_2$ is a train track map.

A subgroup H of $\pi_1(\mathbb{G}, p)$ is a *generalized edge group* if there exist edges e and e' of G with terminal vertices v and v' and paths σ from p to v and σ' from p to v' such




Figure 3. The topological representative $f_2: \mathcal{G}_2 \rightarrow \mathcal{G}_2$.

that each element of H may be represented by a loop of the form $\sigma_{\iota_e}(h)\bar{\sigma}$ for $h \in \mathbb{G}_e$ and if H contains all elements of the form $\sigma'_{\iota_{e'}}(h')\bar{\sigma}'$ for $h' \in \mathbb{G}_{e'}$. In the language of the action of $\pi_1(\mathbb{G}, p)$ on the Bass–Serre tree Γ , the first condition says that there is some edge \tilde{e} lifting e such that H is contained in the stabilizer of \tilde{e} , while the second says that there is some edge \tilde{e}' lifting e' such that H contains the stabilizer of \tilde{e}' . All reduced graphs of groups homotopy equivalent to \mathbb{G} have the same generalized edge groups by [10, Proposition 4.6], and by its proof, we have that any element $\varphi \in \text{Mod}(\mathbb{G})$ permutes the conjugacy classes of generalized edge groups in $\pi_1(\mathbb{G})$.

An edge e of a graph of groups \mathcal{G} is *surviving* if there is a collapse map $\mathcal{G} \rightarrow \mathcal{G}'$ with \mathcal{G}' reduced such that e is not collapsed. Suppose the Bass–Serre tree of \mathcal{G} belongs to a deformation space \mathcal{D} . Guirardel–Levitt consider a space $P\mathcal{G}$ of trees all of whose edges are surviving and prove that if \mathcal{D} is what is called *non-ascending*, then $P\mathcal{G}$ is a finite-dimensional deformation retract of the projectivized deformation space $P\mathcal{D}$ (see [10, Theorem 7.6]). Clay proves that if we work in the *weak topology* and \mathcal{G} is *irreducible*, then there is still a deformation retraction from $P\mathcal{D}$ to the *spine* of $P\mathcal{G}$ (see the discussion after [3, Lemma 1.11]). $\text{Mod}(\mathcal{D}) = \text{Mod}(\mathcal{G})$ acts on $P\mathcal{G}$ and its spine. There are finitely many isomorphism types of graphs of groups \mathcal{G}' homotopy equivalent to \mathcal{G} all of whose edges are surviving precisely when this action has finitely many orbits of cells. For *generalized Baumslag–Solitar groups* where all vertex and edge groups are infinite cyclic, Forester proved in [8, Theorem 8.2] that this happens when there is no *nontrivial integer modulus*.

The main result of this section is the following theorem.

Theorem 3.2. *Suppose $\varphi \in \text{Mod}(\mathbb{G})$ is irreducible and that one of the following conditions holds.*

- (1) *No generalized edge group of \mathbb{G} is mapped properly into a conjugate of itself by some and hence any automorphism Φ representing φ .*
- (2) *Edge groups of \mathbb{G} have finite index in their incident vertex groups, i.e. the Bass–Serre tree Γ is locally finite and there are only finitely many isomorphism types of graphs of groups \mathcal{G}' homotopy equivalent to \mathbb{G} all of whose edges are surviving.*

Then there exists a train track map $f: \mathcal{G} \rightarrow \mathcal{G}$ representing φ on a \mathbb{G} -marked graph of groups \mathcal{G} with the appropriate property above.

(The finite generation assumption is unnecessary in this section.)

The broad-strokes outline of the proof of Theorem 3.2 is much the same as the previous example. By folding at illegal turns, we often produce nontrivial tightening, which decreases the Perron–Frobenius eigenvalue. By controlling the presence of valence-one and valence-two vertices, we may argue that the transition matrix lies in a finite set of matrices, thus the Perron–Frobenius eigenvalue may only be decreased finitely many times. In the remainder of this section, we make this precise by recalling Bestvina and Handel’s original analysis. The proofs are largely identical to the original, so we omit them.

Subdivision. Given a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$, if p is a point in the interior of an edge e such that $f(p)$ is a vertex, we may give \mathcal{G} a new graph of groups structure by declaring p to be a vertex with vertex group equal to $\mathcal{G}_{f(p)}$. If $f(e) = \gamma_1 g \gamma_2$ is the subdivision of the graph of groups edge path $f(e)$ at the image of the point p , where $g \in \mathcal{G}_{f(p)}$, and the new edges incident to p are e_1 and e_2 , define for definiteness $f(e_1) = \gamma_1 g$ and $f(e_2) = \gamma_2$. The map $f_p: \mathcal{G}_p \rightarrow \mathcal{G}_{f(p)}$ is given by the commutative diagram

$$\begin{array}{ccccc}
 & & \mathcal{G}_e & & \\
 & \swarrow f_{e'} & & \searrow f_{e''} & \\
 \mathcal{G}_{e'} & \xrightarrow{\iota_{e'}} & \mathcal{G}_{f(p)} & \xleftarrow{\text{ad}(g)\iota_{e''}} & \mathcal{G}_{e''}
 \end{array}$$

where e' and e'' are the last edge of γ_1 and first edge of γ_2 , respectively. (Note that up to homotopy, we may factor $g \in \mathcal{G}_{f(p)}$ arbitrarily as $g'g''$ and define $f(e_1) = \gamma_1 g'$ and $f(e_2) = g''\gamma_2$.)

Lemma 3.3 ([2, Lemma 1.10]). *If $f: \mathcal{G} \rightarrow \mathcal{G}$ is a topological representative and $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is obtained by subdivision, then f_1 is a topological representative. If f is irreducible, then f_1 is too, and the associated Perron–Frobenius eigenvalues are equal. ■*

Valence-one homotopy. Recall that a valence-one vertex v with incident edge e is *inessential* if the monomorphism $\iota_e: \mathcal{G}_e \rightarrow \mathcal{G}_v$ is an isomorphism.

If v is an inessential valence-one vertex with incident edge e , let \mathcal{G}_1 denote the subgraph of groups determined by $G \setminus \{e, v\}$, and let $\pi: \mathcal{G} \rightarrow \mathcal{G}_1$ be the map collapsing e . Let $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ be the topological representative obtained from $\pi f|_{\mathcal{G}_1}$ by tightening and collapsing a maximal pretrivial forest. We say that $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is obtained from $f: \mathcal{G} \rightarrow \mathcal{G}$ by a *valence-one homotopy*.

Lemma 3.4 ([2, Lemma 1.11]). *If $f: \mathcal{G} \rightarrow \mathcal{G}$ is an irreducible topological representative with Perron–Frobenius eigenvalue λ and $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is obtained from $f: \mathcal{G} \rightarrow \mathcal{G}$ by performing valence-one homotopies on all inessential valence-one vertices of \mathcal{G} followed by the collapse of a maximal invariant forest, then $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is irreducible, and the associated Perron–Frobenius eigenvalue λ_1 satisfies $\lambda_1 < \lambda$.*

Valence-two homotopy. We likewise distinguish two kinds of valence-two vertices. A valence-two vertex v with incident edges e_i and e_j is *inessential* if at least one of the monomorphisms $\iota_{e_i}: \mathcal{G}_{e_i} \rightarrow \mathcal{G}_v$ and $\iota_{e_j}: \mathcal{G}_{e_j} \rightarrow \mathcal{G}_v$ is an isomorphism, say $\iota_{e_j}: \mathcal{G}_{e_j} \rightarrow \mathcal{G}_v$. Let π be the map that collapses e_j to a point and expands e_i over e_j . Define a map $f': \mathcal{G} \rightarrow \mathcal{G}$ by tightening πf . Observe that no vertex of \mathcal{G} is mapped to v . Thus we may define a new graph of groups structure \mathcal{G}' by removing v from the set of vertices. Thus the edge path $e_i \bar{e}_j$ is now an edge, which we will call e_i with edge group \mathcal{G}_{e_i} . Let $f'': \mathcal{G}' \rightarrow \mathcal{G}'$ be the map obtained by tightening $f''(e_i) = f'(e_i \bar{e}_j)$. Finally, let $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ be the topological realization obtained by collapsing a maximal pretrivial forest. We say that $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is obtained by a *valence-two homotopy of v across e_j* .

Lemma 3.5 ([2, Lemma 1.13]). *Let $f: \mathcal{G} \rightarrow \mathcal{G}$ be an irreducible topological representative, and suppose \mathcal{G} has no inessential valence-one vertices. Suppose $f_2: \mathcal{G}_2 \rightarrow \mathcal{G}_2$ is the irreducible topological representative obtained by performing a valence-two homotopy of v across e_j followed by the collapse of a maximal invariant forest. Let M be the transition matrix of f and choose a positive eigenvector \vec{w} with $M\vec{w} = \lambda\vec{w}$. If $w_i \leq w_j$, then $\lambda_2 \leq \lambda$; if $w_i < w_j$, then $\lambda_2 < \lambda$.*

Remark 3.6. The statement of the lemma hides a problem: if we cannot freely choose which edge incident to an inessential valence-two vertex to collapse via a valence-two homotopy, we may be forced to *increase* λ . Since we aim always to decrease λ , we cannot perform such valence-two homotopies. These are the *problematic valence-two vertices* mentioned in the introduction. Our assumptions are designed to limit their proliferation.

Folding. Suppose some pair of edges e_1, e_2 in \mathcal{G} sharing a common initial vertex has the same f -image (as graph-of-groups edge paths). Define a new graph of groups \mathcal{G}_1 by identifying e_1 and e_2 to a single edge e . The map $f: \mathcal{G} \rightarrow \mathcal{G}$ descends to a well-defined homotopy equivalence $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$. This is an *elementary fold*. More generally if e'_1 and e'_2 are maximal initial segments of e_1 and e_2 with equal f -images and endpoints sent to a vertex by f , we first subdivide at the endpoints of e'_1 and e'_2 if they are not already vertices and then perform an elementary fold on the resulting edges.

Let us remark that when lifting the map f to the Bass–Serre tree, it is possible that the lifted map may identify a pair of edges \tilde{e} and $g.\tilde{e}$ in the same orbit e and sharing a common initial vertex. Suppose $\tau(e) = v$ and that e' is the last edge in the edge path $f(e)$. This happens when \mathcal{G}_v contains an element g such that g is not in the image $\iota_e(\mathcal{G}_e)$ but $f_v(g)$ is in the image of $\iota_{e'}(\mathcal{G}_{e'})$. To perform the fold, we may need to subdivide at the preimage of $\tau(\tilde{e}')$. In \mathcal{G} , the “fold” merely changes the edge group, increasing it to $f_v^{-1}(f_v(\mathcal{G}_v) \cap \iota_{e'}(\mathcal{G}_{e'}))$. This fold has no effect on the transition matrix for f unless nontrivial tightening occurs, in which case the Perron–Frobenius eigenvalue decreases. (See [2, Remark 1.6].) It is these folds which may introduce problematic valence-two vertices, so we shall have to be careful about when to perform them.

Lemma 3.7 ([2, Lemma 1.15]). *Suppose $f: \mathcal{G} \rightarrow \mathcal{G}$ is an irreducible topological representative and that $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ is obtained by folding a pair of edges. If f_1 is a topological representative, then it is irreducible, and the associated Perron–Frobenius eigenvalues satisfy $\lambda_1 = \lambda$. Otherwise, let $f_2: \mathcal{G}_2 \rightarrow \mathcal{G}_2$ be the irreducible topological representative obtained by tightening, collapsing a maximal pretrivial forest, and collapsing a maximal invariant forest. Then the associated Perron–Frobenius eigenvalues satisfy $\lambda_2 < \lambda$.*

Lemma 3.8. *Suppose $f: \mathcal{G} \rightarrow \mathcal{G}$ is an irreducible topological representative of the outer automorphism $\varphi \in \text{Out}(\pi_1(\mathcal{G}))$ and no generalized edge group of \mathcal{G} is mapped properly into a conjugate of itself by some and hence any automorphism Φ representing φ . Then the edge groups of \mathcal{G} are all isomorphic, and the injective maps f_{e,e_i} are isomorphisms. If $f(e) = f(e')$ and e and e' share a common vertex so that we may fold e and e' , then $\iota_e(\mathcal{G}_e) = \iota_{e'}(\mathcal{G}_{e'})$ as subgroups of $\mathcal{G}_{\tau(e)}$, and the resulting map $f': \mathcal{G}' \rightarrow \mathcal{G}'$ again has the property above.*

Proof. It is clear that for each edge e and each edge e_i of the edge path $f(e)$, the map f_{e,e_i} is injective, for if it were not, then the map on the fundamental group $f_{\#}$ would fail to be injective. Since f is irreducible, there exists a sequence $e = e^0, e^1, \dots, e^k = e$ such that e^i appears in the edge path $f(e^{i-1})$ in either orientation for i satisfying $1 \leq i \leq k$. Thus we have a sequence of injective homomorphisms

$$\mathcal{G}_e = \mathcal{G}_{e^0} \xrightarrow{f_{e^0,e^1}} \dots \xrightarrow{f_{e^{k-1},e^k}} \mathcal{G}_{e^k} = \mathcal{G}_e.$$

The composition of these maps is an isomorphism, for otherwise the generalized edge group corresponding to \mathcal{G}_e would be mapped properly into a conjugate of itself, from which we conclude that each composing map is an isomorphism. By irreducibility of f , we may choose e^1 to be any edge of the edge path $f(e)$, so we see that each map f_{e,e_i} is an isomorphism.

Suppose that $f(e) = f(e') = g_0 e_1 g_1 \dots e_k g_k$. We have the following pair of commutative diagrams:

$$\begin{array}{ccc} \mathcal{G}_e & \xrightarrow{\iota_e} & \mathcal{G}_v \\ f_{e,e_k} \downarrow & & \downarrow \text{ad}(g_k) f_v \\ \mathcal{G}_{e_k} & \xrightarrow{\iota_{e_k}} & \mathcal{G}_{f(v)}, \end{array} \qquad \begin{array}{ccc} \mathcal{G}_{e'} & \xrightarrow{\iota_{e'}} & \mathcal{G}_v \\ f_{e',e_k} \downarrow & & \downarrow \text{ad}(g_k) f_v \\ \mathcal{G}_{e_k} & \xrightarrow{\iota_{e_k}} & \mathcal{G}_{f(v)}. \end{array}$$

Since f_{e,e_k} and f_{e',e_k} are isomorphisms, we conclude that

$$\text{ad}(g_k) f_v \iota_e(\mathcal{G}_e) = \iota_{e_k}(\mathcal{G}_{e_k}) = \text{ad}(g_k) f_v \iota_{e'}(\mathcal{G}_{e'})$$

and since $\text{ad}(g_k) f_v$ is injective, it follows that $\iota_e(\mathcal{G}_e) = \iota_{e'}(\mathcal{G}_{e'})$. Observe that if Γ is the original Bass–Serre tree and Γ' is the new Bass–Serre tree after folding, then every edge stabilizer in Γ is an edge stabilizer in Γ' and conversely every edge stabilizer in Γ' comes from an edge stabilizer in Γ . It follows that \mathcal{G} and \mathcal{G}' have the same generalized edge groups, so the resulting map $f': \mathcal{G}' \rightarrow \mathcal{G}'$ has the desired property. ■

From Lemma 3.8, we deduce that in fact, if \mathcal{G} satisfies our first standing assumption, then all folding takes place between *distinct* edges of \mathcal{G} , and that the edge group of the newly folded edge is isomorphic to \mathcal{G}_e and $\mathcal{G}_{e'}$.

Lemma 3.9. *The number of edges of a graph of groups homotopy equivalent to \mathbb{G} without inessential valence-one or valence-two vertices is bounded.*

If the Bass–Serre tree of \mathbb{G} is locally finite and there are finitely many isomorphism types of graphs of groups homotopy equivalent to \mathbb{G} all of whose edges are surviving, then there is a bound to the number of edges of a graph of groups homotopy equivalent to \mathbb{G} without inessential valence-one vertices and for which every inessential valence-two vertex is problematic.

Here we call an inessential valence-two vertex of a graph of groups *problematic* if exactly one of its edge-to-vertex group inclusions is surjective.

Proof. Call a vertex of a graph of groups \mathcal{G} *essential* if for all oriented edges $e \in \text{st}(v)$, the monomorphism $\iota_e: \mathcal{G}_e \rightarrow \mathcal{G}_v$ is not surjective. Because the graph of groups \mathbb{G} is assumed to be reduced, every vertex of \mathbb{G} is either essential or incident to an edge e which forms a loop and for which one of the monomorphisms ι_e or $\iota_{\bar{e}}$ is surjective. Let $\eta(\mathbb{G})$ be the number of essential vertices of \mathbb{G} and let $\beta(\mathbb{G})$ be the first Betti number of \mathbb{G} .

We claim that any graph of groups \mathcal{G} homotopy equivalent to \mathbb{G} has at most $\eta(\mathbb{G})$ essential vertices, but it may have fewer. Each essential vertex of \mathbb{G} corresponds to the conjugacy class of a *maximal elliptic* subgroup H ; here elliptic means H stabilizes some vertex of the relevant Bass–Serre tree and maximal means that H is not conjugate to a proper subgroup of an elliptic subgroup. Let T be the Bass–Serre tree of \mathbb{G} . We claim that each maximal elliptic subgroup H of F with the additional property that the fixed-point set of H is bounded is represented by the vertex group of some essential vertex of \mathbb{G} . Indeed, if H is maximal elliptic, then it is contained in and hence equal to some vertex stabilizer in T . The corresponding vertex v of \mathbb{G} must be essential: by maximality, if v is incident to an edge which forms a loop and for which one of the monomorphisms ι_e or $\iota_{\bar{e}}$ is surjective, then both are surjective, which would contradict boundedness of the fixed-point set of H . The properties of being maximal elliptic and having bounded fixed-point set are invariant under homotopy equivalence by [10, Theorem 3.8], from which it follows that \mathcal{G} has at most $\eta(\mathbb{G})$ essential vertices.

Now, \mathcal{G} has at most $2\eta(\mathbb{G}) + 3\beta(\mathbb{G}) - 3$ edges. To see this, form a new graph G' from G by cyclically ordering the essential vertices of \mathcal{G} and attaching an edge from each essential vertex to its neighbors in the cyclic ordering. The graph G' has no valence-one or valence-two vertices and first Betti number at most $\eta(\mathbb{G}) + \beta(\mathbb{G})$. An Euler characteristic argument reveals that G' has at most $3(\eta(\mathbb{G}) + \beta(\mathbb{G})) - 3$ edges, from which the stated bound for \mathcal{G} follows.

Now suppose that the Bass–Serre tree of \mathbb{G} is locally finite and that there are finitely many isomorphism types of graphs of groups homotopy equivalent to \mathbb{G} all of whose edges are surviving. The Bass–Serre tree of each of these graphs of groups is locally

finite with finitely many orbits of vertices, so there is a maximum valence M of any vertex in any such Bass–Serre tree. Let \mathcal{G} be a graph of groups homotopy equivalent to \mathbb{G} without inessential valence-one vertices and for which every inessential valence-two vertex is problematic. It is not quite true that the graph of groups obtained from \mathcal{G} by performing a maximal collapse of edges incident to inessential valence-two vertices has the property that every edge is surviving, so we further collapse all non-surviving edges to obtain a graph of groups \mathcal{G}' without inessential valence-one or valence-two vertices, every edge of which is surviving. Let Γ and Γ' be the Bass–Serre trees of \mathcal{G} and \mathcal{G}' . The collapse map $p: \Gamma \rightarrow \Gamma'$ has compact fibers, so each vertex \tilde{v} of Γ' corresponds to a finite subtree $T_{\tilde{v}}$ of Γ . The valence of \tilde{v} is the sum of the edges incident to $T_{\tilde{v}}$ but not contained in it and is bounded by M . Since each valence-two vertex of \mathcal{G} is problematic, each vertex of Γ has valence at least three, and it follows that the size of each tree $T_{\tilde{v}}$ is bounded depending only on M , from which we conclude that there is a bound on the number of problematic valence-two vertices of \mathcal{G} and thus a bound on the number of edges of \mathcal{G} . ■

Proof of Theorem 3.2. Let $f: \mathbb{G} \rightarrow \mathbb{G}$ be an irreducible topological representative of φ , where we recall that \mathbb{G} is *reduced* in the sense of [7] and satisfies one of our assumptions.

Suppose the Perron–Frobenius eigenvalue λ satisfies $\lambda = 1$. Then f transitively permutes the edges of G and is thus a train track map. So assume $\lambda > 1$. Recall our standing assumptions:

- (1) No generalized edge group of \mathbb{G} is mapped properly into a conjugate of itself by some and hence any automorphism Φ representing φ .
- (2) Edge groups of \mathbb{G} have finite index in the incident vertex groups.

If \mathbb{G} satisfies the first assumption, we will show that Lemma 3.8 implies that any graph of groups \mathcal{G}' obtained from \mathbb{G} has no problematic valence-two vertices and by Lemma 3.9 we conclude that there is a uniform bound L to the number of edges of \mathcal{G}' . If \mathbb{G} satisfies the second assumption, then we cannot prevent \mathcal{G}' from having problematic valence-two vertices, but again by Lemma 3.9 there is a uniform bound L to the number of edges.

We will show that if $f: \mathbb{G} \rightarrow \mathbb{G}$ is not a train track map, then there is an irreducible topological representative $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ without inessential valence-one vertices such that the associated Perron–Frobenius eigenvalues satisfy $\lambda_1 < \lambda$. If \mathbb{G} satisfies the first assumption, we show that \mathcal{G}_1 has no inessential valence-two vertices. If instead \mathbb{G} satisfies the second assumption, we show that inessential valence-two vertices of G_1 are problematic. It follows that the size of the transition matrix of f_1 is uniformly bounded.

Furthermore, if M is an irreducible matrix, its Perron–Frobenius eigenvalue λ is bounded below by the minimum sum of the entries of a row of M . To see this, let \vec{w} be a positive eigenvector. If w_j is the smallest entry of \vec{w} , then $\lambda w_j = (M \vec{w})_j$ is greater than w_j times the sum of the entries of the j th row of M .

Thus if we iterate this argument reducing the Perron–Frobenius eigenvalue, there are only finitely many irreducible transition matrices that can occur, so at some finite stage the Perron–Frobenius eigenvalue will reach a minimum. At this point, we must have a train track map.

To complete the proof, we turn to the question of decreasing λ . Suppose $f: \mathbb{G} \rightarrow \mathbb{G}$ is not a train track map. Then there exists a point p in the interior of an edge such that $f(p)$ is a vertex, and f^k is not locally injective (as a map of graphs of groups) at p for some $k > 1$. We assume that topological representatives act linearly on edges with respect to some metric on G . Since $\lambda > 1$, this means the set of points of G eventually mapped to a vertex is dense. Thus we can choose a neighborhood U of p so small that it satisfies the following conditions.

- (1) The boundary ∂U is a two-point set $\{s, t\}$, where $f^\ell(s)$ and $f^\ell(t)$ are vertices for some $\ell \geq 1$.
- (2) $f^i|_U$ is injective (as a map of graphs of groups) for $1 \leq i \leq k - 1$.
- (3) f^k is two-to-one on $U \setminus \{p\}$, and $f^k(U)$ is contained within a single edge.
- (4) $p \notin f^i(U)$ for $1 \leq i \leq k$.

Note that a priori the map $f^i|_U$ could fail to be injective as a map of graphs sooner than it fails to be injective as a map of graphs of groups. Suppose that there exists $1 \leq j \leq k$ such that $f^i|_U$ is injective as a map of graphs for $1 \leq i \leq j - 1$, and $f^j|_U$ is two-to-one on $U \setminus \{p\}$ as a map of graphs. We will show that Lemma 3.8 implies that in fact $j = k$. If instead edge groups of \mathbb{G} merely have finite index in the incident vertex groups, we allow the same-orbit edge fold, in view of our bound on the number of problematic valence-two vertices that can arise.

First we subdivide at p . Then we subdivide at $f^i(s)$ and $f^i(t)$ for $0 \leq i \leq \ell - 1$ (in reverse order so that subdivision is allowed). The vertex p has valence two; denote the incident edges by e and e' . Observe that in the Bass–Serre tree Γ , there are lifts \tilde{e} and \tilde{e}' such that $\tilde{f}^{k-1}(\tilde{e})$ and $\tilde{f}^{k-1}(\tilde{e}')$ are single edges sharing a common initial vertex that are identified by \tilde{f} . Thus we may fold. Suppose that $j < k$. Then this fold increases the edge group $\mathcal{G}_{f^{k-1}(e)} = \mathcal{G}_{f^{k-1}(e')}$, so the map of edge groups $\mathcal{G}_{f^{k-1}(e)} \rightarrow \mathcal{G}_{f^k(e)}$ is injective (for otherwise $f_\#$ could not be injective) and not surjective. This contradicts Lemma 3.8. Therefore $j = k$ if \mathbb{G} satisfies the first assumption, and in fact $f^{k-1}(e)$ and $f^{k-1}(e')$ are distinct single edges that are identified by f .

The resulting map $f': \mathcal{G}' \rightarrow \mathcal{G}'$ may be a topological representative, in which case the Perron–Frobenius eigenvalue λ' satisfies $\lambda' = \lambda$. In this case $\tilde{f}'^{k-2}(\tilde{e})$ and $\tilde{f}'^{k-2}(\tilde{e}')$ are single edges that are identified by f . In the contrary case, nontrivial tightening occurs. After collapsing a maximal pretrivial forest and a maximal invariant forest, the resulting irreducible topological representative $f'': \mathcal{G}'' \rightarrow \mathcal{G}''$ has Perron–Frobenius eigenvalue λ'' satisfying $\lambda'' < \lambda$.

Repeating this dichotomy k times if necessary, we have either decreased λ , or we have folded e and e' (which are distinct edges in the graph of groups) so that p is now an inessential valence-one vertex.

We remove inessential valence-one and non-problematic valence-two vertices by the appropriate homotopies. Note that if \mathbb{G} and φ satisfy the first assumption, all valence-two vertices present were created by subdivision, not same-orbit folding, and thus are not prob-

lematic. Since valence-one homotopy always decreases the Perron–Frobenius eigenvalue, the resulting irreducible topological representative $f_1: \mathcal{G}_1 \rightarrow \mathcal{G}_1$ has Perron–Frobenius eigenvalue λ_1 satisfying $\lambda_1 < \lambda$. ■

Remark 3.10. As in the original, the proof of Theorem 3.2 provides in outline an algorithm that takes as input a topological representative of an irreducible outer automorphism and returns a train track map. To make it a true algorithm in general, one needs an “oracle” that can compute images of the various homomorphisms f_e and f_v , compute products of elements in the vertex groups \mathcal{G}_v and tell when two vertex group elements are equal.

A *reduction* for an outer automorphism $\varphi \in \text{Out}(\pi_1(\mathcal{G}))$ is a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ which has no inessential valence-one vertices and no invariant forests but has a nontrivial invariant subgraph. If φ has a reduction, then it is *reducible*, i.e. not irreducible. Let $F = A_1 * \cdots * A_n * F_k$ be a free product, represented as the fundamental group of a graph of groups \mathbb{G} with trivial edge groups, vertex groups the A_i and ordinary fundamental group free of rank k . Define the *complexity* of F relative to \mathbb{G} to be the quantity $n + 2k - 1$. If F' is a free factor of F relative to \mathbb{G} , we may define the complexity of F' relative to \mathbb{G} analogously. The final result of this section is the following characterization of reducibility for outer automorphisms φ represented on \mathbb{G} -marked graphs of groups with trivial edge groups.

Proposition 3.11. *Let F be a free product. An outer automorphism φ represented on a \mathbb{G} -marked graph of groups with trivial edge groups is reducible relative to \mathbb{G} if and only if there are free factors F^1, \dots, F^m of F with positive complexity such that $F^1 * \cdots * F^m$ is a free factor of F and φ cyclically permutes the conjugacy classes of the F^i .*

Proof. Suppose first that φ is reducible relative to \mathcal{A} ; let $f: \mathcal{G} \rightarrow \mathcal{G}$ be a reduction and let $G_i = f^i(G_1), 0 \leq i \leq m - 1$ denote distinct noncontractible components of an f -invariant subgraph. Then each $\pi_1(\mathcal{G}|_{G_i})$ determines a free factor F^i with positive complexity such that $F^1 * \cdots * F^m$ is a free factor of F and such that φ cyclically permutes the conjugacy classes of the F^i .

Conversely, suppose F^1, \dots, F^m are free factors with positive complexity as in the statement of the proposition. Take F^{m+1} a free factor so that $F = F^1 * \cdots * F^m * F^{m+1}$. Suppose that n_i and k_i are the data determining the complexity of F^i for $1 \leq i \leq m + 1$. Let \mathbb{G}_i be the thistle with n_i prickles and k_i petals (if $n_{m+1} = k_{m+1} = 0$, then \mathbb{G}_{m+1} is a vertex) and distinguished vertex \star_i . For each i satisfying $1 \leq i \leq m$ choose automorphisms $\Phi_i: F \rightarrow F$ representing φ such that $\Phi_i(F^i) = F^{i+1}$, with indices taken mod m , and let $f_i: \mathbb{G}_i \rightarrow \mathbb{G}_{i+1}$ be the corresponding topological representatives taking \star_i to \star_{i+1} . Define \mathcal{G} to be the union of the \mathbb{G}_i for $1 \leq i \leq m + 1$ together with, for $1 \leq i \leq m$, an oriented edge E_i connecting \star_i to \star_{m+1} .

Collapsing the E_i to a point yields a homotopy equivalence $\mathcal{G} \rightarrow \mathbb{G}$, where \mathbb{G} is the thistle with n prickles and k petals. Identifying the image of $\pi_1(\mathbb{G}_i, \star_i)$ with F^i will serve as (the inverse of) a marking. We will use Φ_1 to create a topological representative $f: \mathcal{G} \rightarrow$

\mathcal{G} for φ . Define $f(\mathbb{G}_i) = f_i(\mathbb{G}_i)$ for $1 \leq i \leq m$. By assumption there exist $c_i \in F$ such that $\Phi_1(x) = c_i \Phi_i(x) c_i^{-1}$. Choose γ_i a closed tight edge path based at \star_{m+1} representing c_i (so γ_1 is the trivial path) and define $f(E_i) = \gamma_i E_{i+1}$ with indices taken mod m . Finally, define $f(\mathbb{G}_{m+1})$ by Φ_1 and the marking on \mathbb{G}_{m+1} .

The topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ is a reduction for φ unless \mathcal{G} has an invariant contractible forest. Since thistles have contractible subgraphs, there are a few possibilities. If there is a family of non-loop edges e_1, \dots, e_m with $e_i \in \mathbb{G}_i$ and $f(e_i) = e_{i+1}$ with indices mod m , we may collapse each of these edges. (Note that up to equivalence, if $f(e_i) = e_{i+1}$ as a map of graphs, then $f(e_i) = e_{i+1}$ as a map of graphs of groups.) Likewise if some non-loop edge of \mathbb{G}_{m+1} is sent to itself, we may collapse it. If each $c_i = 1 \in F$, then the E_i also form an invariant forest that is contractible if the subgraph they span contains at most one vertex with vertex group some A_i . After all these forest collapsings, the only worry is that F^{m+1} has nonpositive complexity and the E_i would be collapsed, leaving G as the only f -invariant subgraph. In this case, choose A an edge of \mathbb{G}_1 sharing an initial vertex with E_1 , and change f via a homotopy with support in E_1 so that $f(E_1) = f(A) f(\bar{A}) E_2$, then fold the initial segment of E_1 mapping to $f(A)$ with all of A . The resulting graph is combinatorially identical to \mathcal{G} but the markings differ. Now $f(E_1) = f(\bar{A}) E_2$ and $f(E_k) = \bar{A} E_1$, so the E_i no longer form an invariant forest. ■

4. Relative train track maps

The purpose of this section is to prove the general case of Theorem A. The strategy is to adapt arguments in [2, Section 5] and [6, Section 2].

Filtrations. A *filtration* on a marked graph of groups \mathcal{G} with respect to a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ is an increasing sequence $\emptyset = G_0 \subset G_1 \subset \dots \subset G_m = G$ of f -invariant subgraphs. The subgraphs are not required to be connected.

Strata. The r th *stratum* of \mathcal{G} is the subgraph H_r containing those edges of G_r not contained in G_{r-1} . An edge path has *height* r if it is contained in G_r and meets the interior of H_r . If both edges of a turn T are contained in a stratum H_r , then T is a *turn in* H_r . If a path has height r and contains no illegal turns in H_r , then it is *r-legal*.

When we think of a stratum H_r or a filtration element G_r as a graph of groups in its own right, the vertex and edge groups of H_r and G_r are equal to what they are in \mathcal{G} . In the language of Bass [1], we work with subgraphs of groups, not subgraphs of subgroups.

Transition submatrices. Relabeling the edges of G and thus permuting the rows and columns of the transition matrix M so that the edges of H_i precede those of H_{i+1} , M becomes block upper-triangular, with the i th block M_i equal to the square submatrix of M containing those rows and columns corresponding to edges in H_i .

A filtration is *maximal* when each M_i is either irreducible or the zero matrix. If M_i is irreducible, call H_i an *irreducible stratum* and a *zero stratum* otherwise. If H_i is irre-

ducible, M_i has an associated Perron–Frobenius eigenvalue $\lambda_i \geq 1$. If $\lambda_i > 1$, then H_i is an *exponentially growing stratum*. Otherwise $\lambda_i = 1$, we say H_i is *non-exponentially growing* and M_i is a transitive permutation matrix.

Associated to a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ there is a maximal filtration $\emptyset = G_0 \subset \dots \subset G_m = G$ defined as follows [2, p. 33]. Order the edges of G , and let $M = (m_{ij})$ be the resulting transition matrix for f . Construct a graph E with a vertex v_i for each edge e_i of G , and m_{ij} oriented edges from v_j to v_i . Two edges e_1 and e_2 belong to the same irreducible stratum if there exist an oriented path from v_1 to v_2 and an oriented path from v_2 to v_1 . An edge e_1 does not belong to an irreducible stratum if there is no oriented edge path from v_1 to itself. A collection of such edges may determine a zero stratum if for each pair of edges e_1 and e_2 in the collection, there is no oriented edge path from v_1 to v_2 nor from v_2 to v_1 . (Perhaps it is easiest to therefore just let each zero stratum be a single edge.) Let H and H' be two resulting strata; we define a partial order on strata. Put H before H' if there are edges $e_1 \in H$ and $e_2 \in H'$ such that there is an oriented path from v_2 to v_1 . Complete this partial order to a total order arbitrarily. Thus a maximal filtration associated to f is not unique, although the irreducible strata are. We will think of a maximal filtration as part of the data of a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$.

The following lemma is an observation we made in the proof of Lemma 3.8.

Lemma 4.1. *Let $f: \mathcal{G} \rightarrow \mathcal{G}$ be a topological representative of an outer automorphism $\varphi \in \text{Out}(\pi_1(\mathcal{G}))$ with irreducible stratum H_r and the property that no iterate of Φ maps a generalized edge group of \mathcal{G} properly into a conjugate of itself for some and hence any automorphism Φ representing φ . All edge groups in H_r are isomorphic, and in fact if e_i is an edge of H_r in the f -image of the edge e of H_r , then the map f_{e,e_i} is an isomorphism.*

Proof. Let e_1 and e_2 be edges of H_r . By irreducibility, there are some $k_1 \geq 1$ and $k_2 \geq 1$ such that the edge path $f^{k_1}(e_1)$ contains e_2 and similarly $f^{k_2}(e_2)$ contains e_1 . This implies that there are injective homomorphisms $\mathcal{G}_{e_1} \rightarrow \mathcal{G}_{e_2}$ and $\mathcal{G}_{e_2} \rightarrow \mathcal{G}_{e_1}$. The double composition $\mathcal{G}_{e_1} \rightarrow \mathcal{G}_{e_2} \rightarrow \mathcal{G}_{e_1}$ must be an isomorphism, so we conclude that each composing homomorphism is an isomorphism. In fact, by irreducibility, we can arrange so that f_{e_1,e_i} is a composing homomorphism of the map $\mathcal{G}_{e_1} \rightarrow \mathcal{G}_{e_2}$ for e_i any edge of H_r contained in the f -image of the edge e_1 . ■

Eigenvalues. Let H_{r_1}, \dots, H_{r_k} be the exponentially growing strata for $f: \mathcal{G} \rightarrow \mathcal{G}$. We define $\text{PF}(f)$ to be the sequence of associated Perron–Frobenius eigenvalues $\lambda_{r_1}, \dots, \lambda_{r_k}$ in nonincreasing order. We order the set

$$\{\text{PF}(f) \mid f: \mathcal{G} \rightarrow \mathcal{G} \text{ is a topological representative of } \varphi\}$$

lexicographically; thus if $\text{PF}(f) = \lambda_1, \dots, \lambda_k$ and $\text{PF}(f') = \lambda'_1, \dots, \lambda'_\ell$, then $\text{PF}(f) < \text{PF}(f')$ if there is some j with $\lambda_j < \lambda'_j$ and $\lambda_i = \lambda'_i$ for i satisfying $1 \leq i < j$, or if $k < \ell$ and $\lambda_i = \lambda'_i$ for i satisfying $1 \leq i \leq k$.

Relative train track maps. Throughout the paper, we will assume our filtrations are maximal unless otherwise specified. Given σ a path in \mathcal{G} , let $f_{\#}(\sigma)$ denote a tight path homotopic rel endpoints to $f(\sigma)$. (If one wants $f_{\#}(\sigma)$ to be unique, one could insist that $f_{\#}(\sigma)$ be in normal form.) We will denote a maximal filtration preserved by $f: \mathcal{G} \rightarrow \mathcal{G}$ as $\emptyset = G_0 \subset \dots \subset G_m = G$. The map f is a *relative train track map* if for every exponentially growing stratum H_r , we have:

- (EG-i) Directions in H_r are mapped to directions in H_r by Df ; it follows that every turn with one edge in H_r and the other in G_{r-1} is legal.
- (EG-ii) If $\sigma \subset G_{r-1}$ is a homotopically nontrivial path with endpoints in $H_r \cap G_{r-1}$, then some (and hence every) $f_{\#}(\sigma)$ is nontrivial as well.
- (EG-iii) If $\sigma \subset G_r$ is a tight r -legal path, then $f(\sigma)$ is an r -legal path.

The following theorem is the main result of this section.

Theorem 4.2. *Assuming an oracle that can compute products of elements in vertex groups, can compute images of injective homomorphisms between edge groups and vertex groups of \mathcal{G} and can tell when two vertex group elements are equal, and that one of the following conditions holds, there is an algorithm that takes as input a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ of $\varphi \in \text{Out}(\pi_1(\mathcal{G}))$ and improves it to a relative train track map $f': \mathcal{G}' \rightarrow \mathcal{G}'$. The conditions are as follows.*

- (1) *Edge groups of \mathcal{G} are finitely generated and for some and hence every Φ representing φ , no generalized edge group of \mathcal{G} is mapped properly into a conjugate of itself by some iterate of Φ .*
- (2) *Vertex groups of \mathcal{G} are finitely generated and edge groups have finite index in the incident vertex groups. Additionally there are finitely many isomorphism types of graphs of groups \mathcal{G}' homotopy equivalent to \mathcal{G} with the property that every edge of \mathcal{G}' is surviving.*

We sketch the outline of the proof: we begin with a topological representative that is *bounded*, a term which will be defined below. We use two new operations, described in Lemma 4.5 and Lemma 4.6 so that the resulting topological representative satisfies (EG-i) and (EG-ii). If (EG-iii) is not satisfied, as in [2] and [6], we modify the algorithm in the proof of Theorem 3.2 to reduce $\text{PF}(f)$, the set of Perron–Frobenius eigenvalues for the exponentially growing strata of $f: \mathcal{G} \rightarrow \mathcal{G}$, while remaining bounded. The boundedness assumption ensures that we will hit a minimum value after a finite number of moves, at which point (EG-iii) will be satisfied.

Let us say a few words about the input of the algorithm: a finite, connected graph of groups is a graph together with groups and homomorphisms between them (for which we assume we have an oracle). The data of a topological representative is the finite connected graph of groups \mathcal{G} together with the filtration, a finite list of subgraphs of G , a list of finite edge paths $f(e) \in \mathcal{G}$ for each edge e of G , and a finite list of injective homomorphisms between vertex and edge groups of \mathcal{G} ; one for each f_v and each f_{e_i} . The oracle guarantees

that we can, for instance, tell when two edge paths $f(E)$ and $f(E')$ share a common initial segment (perhaps after passing to a homotopic topological representative or changing the marking).

Bounded representatives. As we observed in Lemma 3.9, there exists L such that if \mathcal{G} is a marked graph of groups without inessential valence-one vertices and either

- (1) without inessential valence-two vertices, or
- (2) which satisfies our second assumption and for which every inessential valence-two vertex is problematic,

then \mathcal{G} has at most L edges. Our first assumption coupled with the assumption that φ was irreducible allowed us to remove all inessential valence-two vertices that appeared, but in the general case certain inessential valence-two vertices are useful: one needs to introduce them so that (EG-i) is satisfied, for instance. As it happens, our method for showing that (EG-ii) is satisfied may in general even introduce problematic valence-two vertices. Instead, call a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ *bounded* if there are at most L exponentially growing strata, and if, for each exponentially growing stratum H_r , the associated Perron–Frobenius eigenvalue λ_r is also the Perron–Frobenius eigenvalue of a matrix with at most L rows and columns. As in the proof of Theorem 3.2, if $f: \mathcal{G} \rightarrow \mathcal{G}$ is bounded, the set of $\text{PF}(f')$ for $f': \mathcal{G}' \rightarrow \mathcal{G}'$ a bounded representative of φ satisfying $\text{PF}(f') \leq \text{PF}(f)$ is finite, so operations decreasing $\text{PF}(f)$ will eventually reach a minimum among bounded representatives, which we will denote PF_{\min} . Notice as well that the property of being bounded is a property of the sequence of numbers $\text{PF}(f)$.

Elementary moves revisited. In [2, Lemmas 5.1–5.4], Bestvina and Handel revisit the four elementary moves *subdivision*, *valence-one homotopy*, *valence-two homotopy* and *folding* to analyze their impact on $\text{PF}(f)$. All of these moves except valence-two homotopy produce a topological representative $f': \mathcal{G}' \rightarrow \mathcal{G}'$ such that the associated Perron–Frobenius eigenvalues satisfy $\text{PF}(f') \leq \text{PF}(f)$.

Let us discuss valence-two homotopy. Following [2, p. 35], suppose $e_i \in H_i$ and $e_j \in H_j$ are the edges incident to a valence-two vertex v . We assume $i \leq j$. If $i = j$ and H_i is exponentially growing, choose i and j so that the eigenvector coefficient of e_i is greater than or equal to that of e_j . Here is the key point: in all cases we perform the valence-two homotopy across e_i . Call such a valence-two homotopy *performable* if after making these choices, we have that the inclusion $\iota_{e_i}: \mathcal{G}_{e_i} \rightarrow \mathcal{G}_v$ is an isomorphism.

Lemma 4.3. *Suppose we are in the situation of the first assumption. All valence-two homotopies are performable, perhaps after rearranging strata.*

Proof. We continue to use the notation above. Suppose at first that H_i is a zero stratum. Then the restriction of f to G_i is a homotopy equivalence of G_i with G_{i-1} . In particular, since v is a valence-one vertex of G_i , we must have that $\iota_{e_i}: \mathcal{G}_{e_i} \rightarrow \mathcal{G}_v$ is an isomorphism. Therefore a valence-two homotopy of e_j across e_i is performable.

So assume that H_i is irreducible. Recall the partial order on strata, where $H_i \leq H_k$ if some edge in H_k is eventually mapped over some edge in H_i (and hence any edge in H_i , since H_i is irreducible). If H_j is a zero stratum, we may after dividing it into two zero strata, assume that $H_j = \{e_j\}$. This done, if we have $H_i \leq H_j$ in this *partial* order, and we have that v is an inessential valence-two vertex but $\iota_{e_i}: \mathcal{G}_{e_i} \rightarrow \mathcal{G}_v$ is not an isomorphism, we have a contradiction: by assumption the edge e_j is eventually mapped over the edge e_i , so there is an injective homomorphism $\mathcal{G}_v \cong \mathcal{G}_{e_j} \rightarrow \mathcal{G}_{e_i}$, which therefore must map $\iota_{e_i}(\mathcal{G}_{e_i})$ properly into itself, contradicting the first assumption.

Finally, if we do *not* have $H_i \leq H_j$ in this partial order, then we may freely move H_j below H_i when we complete the partial order to a total order, and thus may swap the roles of i and j if need be. ■

The proof of [2, Lemma 5.4] shows that if $i < j$ and H_i is exponentially growing, then $\text{PF}(f') < \text{PF}(f)$. In the case where $i = j$ and H_i is exponentially growing, it may happen that λ_i is replaced by some number of eigenvalues λ' that all satisfy $\lambda' \leq \lambda_i$, so it is possible that $\text{PF}(f') > \text{PF}(f)$. Nonetheless, we have the following result. Call an elementary move *safe* if performing it on a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ yields a new topological representative $f': \mathcal{G}' \rightarrow \mathcal{G}'$ with $\text{PF}(f') \leq \text{PF}(f)$. Thus all elementary moves with the exception of valence-two homotopy are always safe.

Lemma 4.4 ([2, Lemma 5.5]). *If $f: \mathcal{G} \rightarrow \mathcal{G}$ is a bounded topological representative and $f': \mathcal{G}' \rightarrow \mathcal{G}'$ is obtained from f by a sequence of safe moves with $\text{PF}(f') < \text{PF}(f)$, then there is a bounded topological representative $f'': \mathcal{G}'' \rightarrow \mathcal{G}''$ with $\text{PF}(f'') < \text{PF}(f)$.*

Proof. The proof is essentially identical to [2, Lemma 5.5]. Suppose first that our topological representatives satisfy the first assumption. By performing valence-one and safe valence-two homotopies, we may assume that $f': \mathcal{G}' \rightarrow \mathcal{G}'$ has the property that \mathcal{G}' has no inessential valence-one vertices and that each inessential valence-two vertex v has the property that the two edges incident to v belong to the same exponentially growing stratum. Thus $f': \mathcal{G}' \rightarrow \mathcal{G}'$ has at most L strata, exponentially growing or no, and $\text{PF}(f')$ is obtained from $\text{PF}(f)$ by replacing some of the Perron–Frobenius eigenvalues with strictly smaller eigenvalues λ'_i . For the eigenvalues that are not replaced, the fact that f was bounded implies that these eigenvalues are the Perron–Frobenius eigenvalues for matrices with at most L rows and columns. Thus we only need to show that the λ'_i are also the Perron–Frobenius eigenvalues for matrices with at most L rows and columns. We do this by performing dangerous valence-two homotopies, replacing each λ'_i with some collection of λ''_{ij} , satisfying $\lambda''_{ij} \leq \lambda'_i$ until each resulting stratum has at most L edges. We still have that $f: \mathcal{G}'' \rightarrow \mathcal{G}''$ has at most L strata, so this topological representative $f'': \mathcal{G}'' \rightarrow \mathcal{G}''$ is bounded.

In the situation of the second assumption, the argument is essentially the same. By Lemma 3.9, there is a uniform bound to the number of problematic valence-two vertices, so we need only focus on the inessential valence-two vertices which are *not* problematic. We then proceed exactly as above. ■

Invariant core subdivision. We recall the construction of the *invariant core subdivision* of an exponentially growing stratum H_r . Assume that a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ linearly expands edges over edge paths with respect to some metric on G . If $f(H_r)$ is not entirely contained in H_r , then the set

$$I_r := \{x \in H_r \mid f^k(x) \in H_r \text{ for all } k > 0\}$$

is an f -invariant Cantor set. The *invariant core* of an edge e in H_r is the smallest closed subinterval of e containing the intersection of I_r with the interior of e . The endpoints of invariant cores of edges in H_r form a finite set which f sends into itself. Declaring elements of this finite set to be vertices is called *invariant core subdivision*. The stratum H_r determines a new exponentially growing stratum H'_r whose edges are the invariant cores of edges in H_r .

The following lemma says that invariant core subdivision can be used to create topological representatives whose exponentially growing strata satisfy (EG-i).

Lemma 4.5 ([2, Lemma 5.13]). *If $f': \mathcal{G}' \rightarrow \mathcal{G}'$ is obtained from $f: \mathcal{G} \rightarrow \mathcal{G}$ by an invariant core subdivision of an exponentially growing stratum H_r , then $\text{PF}(f') = \text{PF}(f)$, and the map Df' maps directions in the resulting exponentially growing stratum H'_r to itself, so H'_r satisfies (EG-i). If H_j is another exponentially growing stratum for $f: \mathcal{G} \rightarrow \mathcal{G}$ that satisfies (EG-i) or (EG-ii), then the resulting exponentially growing stratum H'_j for $f': \mathcal{G}' \rightarrow \mathcal{G}'$ still satisfies those properties.*

In fact, invariant core subdivision affects only edges in H_r . If new vertices are created, then one or more non-exponentially growing strata are added to the filtration below H_r .

Collapsing inessential connecting paths. The following lemma says that an application of operations already defined may be used to construct topological representatives whose exponentially growing strata satisfy (EG-ii).

Lemma 4.6 ([2, Lemma 5.14]). *Let $f: \mathcal{G} \rightarrow \mathcal{G}$ be a bounded topological representative with exponentially growing stratum H_r . If $\alpha \subset G_{r-1}$ is a path with endpoints in $H_r \cap G_{r-1}$ such that $f_{\#}(\alpha)$ is trivial, we construct a new bounded topological representative $f': \mathcal{G}' \rightarrow \mathcal{G}'$ such that if H'_r is the stratum of \mathcal{G}' determined by H_r , then either (if the endpoints of α are distinct) $H'_r \cap G'_{r-1}$ has fewer points than $H_r \cap G_{r-1}$ or (if the endpoints of α are equal) a vertex group of $H'_r \cap G'_{r-1}$ has increased.*

If $k > r$ and H_k satisfies (EG-ii), then H'_k , the stratum determined by H_k , satisfies (EG-ii). If $k \geq r$ and H_k satisfies (EG-i), then H'_k satisfies (EG-i).

Proof. We follow the outline of the proof of [2, Lemma 5.14]. Let V be the vertex set of G . Subdivide at each point of $\alpha \cap f^{-1}(V)$, obtaining a topological representative $f_1: \mathcal{G}(1) \rightarrow \mathcal{G}(1)$ and an identifying homotopy equivalence $p_1: \mathcal{G} \rightarrow \mathcal{G}(1)$ whose map of underlying graphs is a homeomorphism but not a cellular map. Define $\alpha_1 = p_1(\alpha)$, and write $\alpha_1 = g_0 e_1 \dots e_k g_k$. There is a map of graphs of groups $h_1: \mathcal{G}(1) \rightarrow \mathcal{G}$ such that $p_1 h_1 = f_1$.

We may write

$$W_1 = h_1(\alpha_1) = (h_1)_{v_0}(g_0)g_{\bar{e}_1}h_1(e_1)g_{e_1} \dots g_{\bar{e}_k}h_1(e_k)g_{e_k}^{-1}(h_1)_{v_k}(g_k).$$

(Here each $h_1(e_i)$ should be understood as the edge determined by the map of underlying graphs.) Since $[W_1]$ is trivial, there is some backtracking, i.e. there exists ℓ such that the edges $h_1(e_\ell)$ and $h_1(\bar{e}_{\ell+1})$ are equal and $g_{e_\ell}(h_1)_{v_\ell}(g_\ell)g_{\bar{e}_{\ell+1}}$ belongs to $\iota_{e_\ell}(\mathcal{G}_{e_\ell})$. The same statement is true of f_1 , so we may (possibly after twisting the marking or changing the fundamental domain as in Example 2.3) fold e_ℓ and $\bar{e}_{\ell+1}$. Note that it is possible that $e_\ell = \bar{e}_{\ell+1}$, in which case the fold increases the edge group \mathcal{G}_{e_ℓ} . We get a resulting homotopy equivalence $f_2: \mathcal{G}(2) \rightarrow \mathcal{G}(2)$ and the resulting quotient map (which may be a homeomorphism of underlying graphs) $p_2: \mathcal{G}(1) \rightarrow \mathcal{G}(2)$. As before, there is a map of graphs of groups $h_2: \mathcal{G}(2) \rightarrow \mathcal{G}$ such that now $p_2 p_1 h_2 = f_2$. If the edges e and e' were folded to create an edge e'' , then (thinking of these edges as segments of edges of \mathcal{G}) we have $f(e) = f(e')$ as length-one edge paths – this is why we twisted the marking – and we define $h_2(e'') = f(e) = f(e')$. Define $\alpha_2 = (p_2 p_1)_\#(\alpha)$, and define $W_2 = h_2(\alpha_2)$ as above. We have that W_2 is obtained from W_1 by canceling some backtracking, so α_2 has fewer edges than α_1 . We have that $[W_2]$ is trivial, so we may repeat the above argument at most k times to produce $f_k: \mathcal{G}(k) \rightarrow \mathcal{G}(k)$ such that $\alpha_k = (p_k p_{k-1} \dots p_2 p_1)_\#(\alpha)$ is the trivial path. Finally, let $f'': \mathcal{G}'' \rightarrow \mathcal{G}''$ be the topological representative obtained from f_k by tightening and collapsing the maximal pretrivial forest.

Since folding decreases $\text{PF}(f)$ or leaves it the same, we have $\text{PF}(f'') \leq \text{PF}(f)$. If $\text{PF}(f'') = \text{PF}(f)$, then f'' is bounded since f was, so we let $f' = f''$. If not, then we apply Lemma 4.4 to produce a bounded topological representative $f': \mathcal{G}' \rightarrow \mathcal{G}'$ such that $\text{PF}(f'') \leq \text{PF}(f') < \text{PF}(f)$.

The argument now finishes as in [2, Lemma 5.14]. If the endpoints of α were distinct, then the exponentially growing stratum H'_r determined by H_r satisfies

$$|H'_r \cap G'_{r-1}| < |H_r \cap G_{r-1}|.$$

If the endpoints were *not* distinct, the vertex group of H'_r determined by the endpoint of α is now larger than it was in H_r , in the sense that there is a natural injective but not surjective identifying homomorphism.

As in [2, Lemma 5.14], if $k > r$ and H_k satisfies (EG-ii), then the corresponding stratum H'_k of $f': \mathcal{G}' \rightarrow \mathcal{G}'$ still satisfies (EG-ii). Likewise, if $k \geq r$ and H_k satisfies (EG-i), then the corresponding stratum H'_k satisfies (EG-i). ■

If \mathcal{G} is a finite graph, has finitely generated edge groups, and has no inessential valence-one vertices, then the map $f: \mathcal{G} \rightarrow \mathcal{G}$ satisfies the assumptions of [5, Theorem 2.1], and thus can be written as a (finite) product of folds and what Dunwoody calls “vertex morphisms.” In fact, the vertex morphisms are unnecessary, because $f_\#$ is an isomorphism. Since each of the folds performed in Lemma 4.6 is a fold factor of f , after performing finitely many such folds, we must have that the exponentially growing stratum of interest H_r satisfies (EG-ii).

Lemma 4.7 ([6, Lemma 2.4]). *Assuming an oracle that can compute products of elements in vertex groups, can compute images of injective homomorphisms between edge groups and vertex groups of \mathcal{G} and can tell when two vertex group elements are equal, there is an algorithm that checks whether a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ is a relative train track map.*

Proof. Since (EG-i) is a finite property (whether the image of a direction belongs to H_r is a property of the underlying edge, and H_r has finitely many edges), we may assume that each exponentially growing stratum satisfies (EG-i).

Suppose H_r is an exponentially growing stratum. A *connecting path* for H_r is a tight path α in G_{r-1} with endpoints in $H_r \cap G_{r-1}$. Since (EG-i) holds, vertices in $H_r \cap G_{r-1}$ are sent to vertices in $H_r \cap G_{r-1}$. For paths with distinct endpoints, we claim that for each component C of G_{r-1} , (EG-ii) for paths with distinct endpoints is equivalent to the condition that distinct vertices of $H_r \cap C$ are sent to distinct vertices of $H_r \cap G_{r-1}$. Indeed, if this holds, then tight paths with distinct endpoints are sent to tight paths with distinct endpoints which are thus homotopically nontrivial. If not, then there is a pair of distinct vertices v and w in $H_r \cap C$ identified by f . In this case there is a connecting path α with endpoints v and w whose $f_{\#}$ -image is trivial (consider what a homotopy inverse does to $f(v') = f(w')$).

Finally, we consider connecting paths with the same endpoint. Let v be a vertex in $H_r \cap C$. If the map $f_v: \mathcal{G}_v \rightarrow \mathcal{G}_{f(v)}$ is an isomorphism, there is nothing to check. The map f induces an isomorphism $f_{\#}: \pi_1(\mathcal{G}, v) \rightarrow \pi_1(\mathcal{G}, f(v))$, so we may consider the subgroup $f_{\#}^{-1}(\mathcal{G}_{f(v)})$. It is elliptic, and in fact fixes a vertex of the Bass–Serre tree Γ (consider again what a homotopy inverse to f does to the vertex $f(v)$). There is a tight path σ such that each element of $f_{\#}^{-1}(\mathcal{G}_{f(v)})$ may be represented by a path of the form $\sigma g \bar{\sigma}$. (This path may not be tight, but may be tightened by a homotopy.) Each of these paths is inessential, in the sense that their $f_{\#}$ -image is trivial, and they are connecting paths for H_r if they are contained in G_{r-1} . Thus a necessary condition for (EG-ii) is that for each such $g \in \mathcal{G}_{f(v)} \setminus f_v(\mathcal{G}_v)$, some and hence any tight path homotopic to $\sigma g \bar{\sigma}$ is *not* contained in G_{r-1} . In fact, this condition is sufficient. This is a finite property, since for each vertex v we need only consider the path σ . Therefore we may assume (EG-ii) holds.

Finally, (EG-iii) for H_r is equivalent to checking that $f(e)$ is r -legal for each edge $e \in H_r$. Since we assume H_r satisfies (EG-i), in the situation of the first assumption, Lemma 4.1 implies that any nondegenerate turn in H_r whose directions determine the same underlying oriented edge of G is legal. Thus if a turn is illegal, it becomes degenerate as soon as the underlying oriented edges of G are identified. This implies that checking (EG-iii) is a finite property: for each of the finitely many turns in H_r crossed by $f(e)$, we need only check that either the underlying edges of the turn are periodic, so the turn never degenerates, or the underlying edges are eventually identified, in which case we only need check whether the actual turn degenerates at that stage. In the case of the second assumption, there are only finitely many directions at a given vertex, so it is clear that (EG-iii) is a finite property. ■

Proof of Theorem 4.2. As in the proof of Theorem 3.2, we begin with a topological representative $f: \mathcal{G} \rightarrow \mathcal{G}$ on a graph of groups satisfying one of our standing assumptions. Assume further that \mathcal{G} is reduced. By assumption, f is bounded. Consider the highest exponentially growing stratum H_r of \mathcal{G} . We check whether H_r satisfies (EG-i) and (EG-ii) using Lemma 4.7. If not, apply Lemma 4.5 and Lemma 4.6 to create a new topological representative, still called $f: \mathcal{G} \rightarrow \mathcal{G}$ such that the resulting exponentially growing stratum H_r satisfies (EG-i) and (EG-ii). Repeat with the next highest exponentially growing stratum until all exponentially growing strata satisfy these properties. Check whether the resulting topological representative, which we still call $f: \mathcal{G} \rightarrow \mathcal{G}$, satisfies (EG-iii). If it does, we are done.

If not, then there is some edge e in an exponentially growing stratum H_r such that $f(e)$ is not r -legal. We apply the algorithm in the proof of Theorem 3.2: there is a point P in H_r where f^k is not injective at P for some $k > 1$. We subdivide and then repeatedly fold. As in the proof of Theorem 3.2, no edge-group-increasing folds are necessary in this step in the case of the first assumption. In the contrary case, we have a bound on the number of edge-group-increasing folds. Either we have reduced the eigenvalue for H_r or produced a valence-one vertex. We remove all valence-one vertices via homotopies and perform all possible valence-two homotopies which do not increase $\text{PF}(f)$. At this point we have created a new topological representative $f': \mathcal{G}' \rightarrow \mathcal{G}'$ with $\text{PF}(f') < \text{PF}(f)$, but f' may not be bounded. Apply Lemma 4.4 to produce a new bounded topological representative $f'': \mathcal{G}'' \rightarrow \mathcal{G}''$ with $\text{PF}(f'') < \text{PF}(f)$. If (EG-i) and (EG-ii) are not satisfied by f'' , we may restore these properties by applying Lemma 4.5 and Lemma 4.6. We saw that these lemmas preserve boundedness and do not increase $\text{PF}(f'')$. Because $\text{PF}(f)$ can only be decreased finitely many times before reaching PF_{\min} , eventually this process terminates, yielding a relative train track map. ■

Corollary 4.8. *If $f: \mathcal{G} \rightarrow \mathcal{G}$ is a topological representative satisfying (EG-i) and with $\text{PF}(f) = \text{PF}_{\min}$, then f is bounded and the exponentially growing strata of f satisfy (EG-iii).*

Acknowledgments. Part of this work originally appeared in the author's thesis [11]. The author would like to thank Lee Mosher for many helpful conversations and comments on early drafts of this work, Chloé Papin for conversations which led to realizing that further assumptions were necessary, Mark Feighn and Mladen Bestvina for suggestions on alternate assumptions, and the anonymous referee for numerous comments and a very careful reading which helped improve the exposition of this article.

References

- [1] H. Bass, Covering theory for graphs of groups. *J. Pure Appl. Algebra* **89** (1993), no. 1-2, 3–47
Zbl [0805.57001](#) MR [1239551](#)
- [2] M. Bestvina and M. Handel, Train tracks and automorphisms of free groups. *Ann. of Math. (2)* **135** (1992), no. 1, 1–51 Zbl [0757.57004](#) MR [1147956](#)
- [3] M. Clay, Deformation spaces of G -trees and automorphisms of Baumslag-Solitar groups. *Groups Geom. Dyn.* **3** (2009), no. 1, 39–69 Zbl [1226.20022](#) MR [2466020](#)
- [4] D. J. Collins and E. C. Turner, Efficient representatives for automorphisms of free products. *Michigan Math. J.* **41** (1994), no. 3, 443–464 Zbl [0820.20035](#) MR [1297701](#)
- [5] M. J. Dunwoody, Folding sequences. In *The Epstein birthday schrift*, pp. 139–158, Geom. Topol. Monogr. 1, Geom. Topol. Publ., Coventry, 1998 Zbl [0927.20013](#) MR [1668347](#)
- [6] M. Feighn and M. Handel, Algorithmic constructions of relative train track maps and CTs. *Groups Geom. Dyn.* **12** (2018), no. 3, 1159–1238 Zbl [1456.20044](#) MR [3845002](#)
- [7] M. Forester, Deformation and rigidity of simplicial group actions on trees. *Geom. Topol.* **6** (2002), 219–267 Zbl [1118.20028](#) MR [1914569](#)
- [8] M. Forester, Splittings of generalized Baumslag-Solitar groups. *Geom. Dedicata* **121** (2006), 43–59 Zbl [1117.20023](#) MR [2276234](#)
- [9] S. Francaviglia and A. Martino, Stretching factors, metrics and train tracks for free products. *Illinois J. Math.* **59** (2015), no. 4, 859–899 Zbl [1382.20031](#) MR [3628293](#)
- [10] V. Guirardel and G. Levitt, Deformation spaces of trees. *Groups Geom. Dyn.* **1** (2007), no. 2, 135–181 Zbl [1134.20026](#) MR [2319455](#)
- [11] R. A. Lyman, *Train tracks on graphs of groups and outer automorphisms of hyperbolic groups*. ProQuest LLC, Ann Arbor, MI, 2020 MR [4119876](#)
- [12] S. Meinert, The Lipschitz metric on deformation spaces of G -trees. *Algebr. Geom. Topol.* **15** (2015), no. 2, 987–1029 Zbl [1364.20029](#) MR [3342683](#)
- [13] P. Scott and T. Wall, Topological methods in group theory. In *Homological group theory* (Proc. Sympos., Durham, 1977), pp. 137–203, London Math. Soc. Lecture Note Ser. 36, Cambridge University Press, Cambridge, 1979 Zbl [0423.20023](#) MR [564422](#)
- [14] J.-P. Serre, *Trees*. Springer Monogr. Math., Springer, Berlin, 2003 Zbl [1013.20001](#) MR [1954121](#)
- [15] M. Sykietis, Stable representatives for symmetric automorphisms of groups and the general form of the Scott conjecture. *Trans. Amer. Math. Soc.* **356** (2004), no. 6, 2405–2441 Zbl [1041.20019](#) MR [2048523](#)

Received 18 February 2021.

Rylee Alanza Lyman

Department of Mathematics, Rutgers University–Newark, Smith Hall, 101 Warren Street,
Newark, NJ 07102, USA; rylee.lyman@rutgers.edu