ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 9 (2005) 1835–1880
Published: 25 September 2005

The Grushko decomposition of a finite graph of finite rank free groups: an algorithm

Guo-An Diao
Mark Feighn

School of Arts and Sciences, Holy Family University
Philadelphia, PA 19114, USA
and
Department of Mathematics and Computer Science
Rutgers University, Newark, NJ 07102, USA

Email: gdiao@holyfamily.edu and feighn@andromeda.rutgers.edu

Abstract

A finitely generated group admits a decomposition, called its *Grushko decomposition*, into a free product of freely indecomposable groups. There is an algorithm to construct the Grushko decomposition of a finite graph of finite rank free groups. In particular, it is possible to decide if such a group is free.

AMS Classification numbers Primary: 20F65

Secondary: 20E05

Keywords: Graph of groups, Grushko decomposition, algorithm, labeled graph

Proposed: Benson Farb
Seconded: Martin Bridson, Joan Birman

Received: 5 February 2005
Revised: 11 September 2005

1 Introduction

Theorem 1.1 (Grushko [15]) *A finitely generated group G is a free product of a finite rank free subgroup and finitely many freely indecomposable non-free subgroups.*

Up to reordering and conjugation, the non-free factors appearing in this *the Grushko decomposition of G* are unique. The rank of the free factor is also an invariant of G. The main result of this paper is:

Theorem 1.2 *There is an algorithm which produces the Grushko decomposition of a finite graph*[1] *of finite rank free groups*[2]*.*

A relative version is given in Section 2.11.

The class of finite graphs of finite rank free groups is fascinating and has received much attention. For example, mapping tori of free group automorphisms are in this class, see [3, 10, 12, 7]. Also, limit groups which appear in the recent work on the Tarski problem, see [19, 24], have a hierarchy in which those limit groups appearing on the first level are finite graphs of finite rank free groups.

The algorithm is given in Section 2. Theorem 2.8 is a refined version of Theorem 1.2 and is proved in Section 10. Work of Shenitzer or Swarup when combined with Whitehead's algorithm for deciding if a given element of a free group is primitive[3] gives the case of Theorem 1.2 where edge groups are cyclic.

Theorem 1.3 (Shenitzer–Swarup [26],[29], [31],[30], see also [28]) *There is an algorithm to decide whether a finite graph of finite rank free groups with cyclic edge groups is free.*

There are other notable situations where the Grushko decomposition may be found algorithmically. For example, given a presentation with one relation for a group G, the Grushko decomposition of G can be constructed (eg, [21]). Also, from a triangulation of a closed orientable 3–manifold M, the connected sum decomposition of M can be found (see [17]). The referee points out two other papers [13] and [20] that consider respectively hyperbolic groups and limit groups.

[1] in the sense of Bass–Serre [25]
[2] ie, vertex and edge groups are free of finite rank, see Section 2.7
[3] an element of some basis

Here is a sketch of the proof of Theorem 1.2. There are three steps. Suppose that S is a cocompact G–tree with finitely generated edge stabilizers. Suppose further that G is freely decomposable and that T is a G–tree with one orbit of edges and with trivial edge stabilizers. Give the product $S \times T$ the diagonal G–action. $S \times T$ is a union of squares (edge$\times$edge). There is a cocompact G–subcomplex $X_S(T)$ that is a G–deformation retract of $S \times T$ (see Section 5). As $X_S(T)$ is contained in $S \times T$, there is a natural map $X_S(T) \to T$. The preimage in $X_S(T)$ of the midpoint of an edge of T is a compact forest. A valence one vertex of this forest corresponds to a square in $X_S(T)$ that may be equivariantly collapsed. We may iteratively collapse until each component of this tree is a point. These points exhibit free decompositions of G that are compatible with the original splitting given by S.

An argument similar to the one in this first step was used by Bestvina–Feighn in [2] to among other things reprove Theorem 1.3. The use of products of trees is inspired by the Fujiwara–Papasoglu [11] approach to the Rips–Sela JSJ–theorem [23].

The second step is to translate these collapses of $X_S(T)$ into corresponding *simplifications* of the original graph of groups. This is straightforward and is done in Section 8. These first two steps do not use the hypothesis that edge and vertex groups are free.

In the third step (Section 9), we show how Gersten representatives [14] of conjugacy classes of subgroups of free groups can be used to detect simplifications. This is probably the heart of the paper.

A special case of Theorem 1.2 solves Problem F24b on the problem list at `http://www.grouptheory.org`.

For the convenience of readers interested primarily in using the algorithm, it is described in the next section (Section 2). Definitions are given, but proofs are, for the most part, deferred until later in the paper. The first two steps of the proof are more general and therefore somewhat cleaner, see Sections 3–8 which can be read independently of Section 2.

The first named author's thesis [9] included an algorithm to decide if a finite graph of finite rank free groups is free. The second named author warmly thanks Mladen Bestvina for helpful conversations and gratefully acknowledges the support of the National Science Foundation.

2 The algorithm

Mainly to establish notation, we first recall the definition of a graph of groups.

2.1 Graphs of groups

A reference for this section is [25]. A *graph* is a 1–dimensional CW–complex and is determined by the following combinatorial data: a 4–tuple $(V,\ \hat{E},\ \mathrm{op},\ \partial_0)$ where

- V and $\hat{E}$ are sets;
- $\mathrm{op}\colon\ \hat{E} \to \hat{E}$ satisfies
  (1) $\mathrm{op} \circ \mathrm{op} = \mathrm{Id}$, and
  (2) $\mathrm{op}(e) \neq e$, for all $e \in \hat{E}$; and
- $\partial_0\colon\ \hat{E} \to V$.

For $e \in \hat{E}$, we also write e^{-1} for $\mathrm{op}(e)$ and set $\partial_1 e = \partial_0 e^{-1}$. For $v \in V$, $\hat{E}(v) = \{e \in \hat{E} : \partial_0 e = v\}$. The *valence of* v is the cardinality $|\hat{E}(v)|$ of $\hat{E}(v)$. Such a 4–tuple is a *combinatorial graph.* The graph

$$\Gamma = \Gamma(V,\ \hat{E},\ \mathrm{op},\ \partial_0)$$

so determined has vertex set identified with V. The set $\hat{E}$ corresponds to the set of oriented edges of Γ; the set E of edges of Γ is identified with $\{\{e, e^{-1}\} : e \in \hat{E}\}$. The map op reverses edge orientations, and ∂_0 determines the characteristic maps of Γ. Up to isomorphism, a graph uniquely determines a combinatorial graph and *vice versa.* In particular, properties of one give properties of the other. The interior of an edge e is denoted $\mathring{e}$.

A *graph of groups* is a 4–tuple

$$\mathcal{G} = (\Gamma(V,\ \hat{E},\ \mathrm{op},\ \partial_0),\ \{G_v : v \in V\},\ \{G_e : e \in \hat{E}\},\ \{\varphi_e\colon\ e \in \hat{E}\})$$

where

- $\Gamma(V,\ \hat{E},\ \mathrm{op},\ \partial_0)$ is a connected combinatorial graph $\Gamma(\mathcal{G})$;
- for $e \in \hat{E}$ and $v \in V$, $G_e = G_{e^{-1}}$ and G_v are groups; and
- for $e \in \hat{E}$, $\varphi_e\colon\ G_e \to G_{\partial_0 e}$ is a monomorphism.

The groups G_e and G_v are respectively *edge* and *vertex groups.* The φ_e's are *bonding maps.* We say that $\mathcal{G}$ is *reduced* if

- $\varphi_e\colon\ G_e \to G_v$ is not an isomorphism for any valence one vertex v; and

- if v has valence two and if $\varphi_e\colon\ G_e \to G_v$ is an isomorphism, then $\hat{E}(v) = \{e, e^{-1}\}$ (in which case $\Gamma(\mathcal{G})$ is a *loop*).

Associated to a graph of groups $\mathcal{G}$ is an isomorphism type of group $\pi_1(\mathcal{G})$, see [25]. If $G \cong \pi_1(\mathcal{G})$ then we say that $\mathcal{G}$ is a *graph of groups decomposition for* G. Let $\mathcal{G}$ and $\mathcal{G}'$ be graphs of groups with the same underlying graphs, edge groups, and vertex groups. We say that $\mathcal{G}$ and $\mathcal{G}'$ are *conjugate*, written $\mathcal{G} \sim \mathcal{G}'$, if there is a sequence $\vec{h} = \{h_e \in G_{\partial_0 e} : e \in \hat{E}\}$ such that $\varphi'_e = i_{h_e} \circ \varphi_e$ where i_{h_e} denotes the inner automorphism induced by h_e, ie, $i_{h_e}(g) = h_e g h_e^{-1}$. If $\mathcal{G}$ and $\mathcal{G}'$ are conjugate, then $\pi_1(\mathcal{G}')$ and $\pi_1(\mathcal{G})$ are isomorphic.

A simplicial action of a group G on a tree T determines a graph of groups with underlying graph T/G and with vertex and edge groups given by vertex and edge stabilizers of in T. Conversely, a graph of groups $\mathcal{G}$ determines up to simplicial isomorphism a $G \cong \pi_1(\mathcal{G})$–tree $T(\mathcal{G})$. See [25]. $\mathcal{G}$ is a *trivial* graph of groups decomposition if $G \cong \pi_1(\mathcal{G})$ fixes a point of $T(\mathcal{G})$. If $\mathcal{G}$ is a non-trivial graph of groups decomposition for G then we also say that $\mathcal{G}$ is a *splitting* for G. If the edge groups of $\mathcal{G}$ are contained in some class of groups then we say that G *splits over this class.* For example, if $\mathcal{G}$ is a non-trivial graph of groups decomposition for G and all edge groups of $\mathcal{G}$ are trivial then we say that G splits over $\mathbf{1}$ where $\mathbf{1}$ denotes the trivial group.

If $\mathcal{G}$ is a graph of groups with $G \cong \pi_1(\mathcal{G})$ then the edges e of $\Gamma(\mathcal{G})$ with $G_e = \mathbf{1}$ determine a free product decomposition $F_m * G_1 * \cdots * G_n$ where m is the rank of the graph obtained from $\Gamma(\mathcal{G})$ by collapsing all edges f with $G_f \neq \mathbf{1}$ and the G_i's are the fundamental groups of graphs of groups given by the components of $\Gamma(\mathcal{G}) \setminus (\cup_e \{\mathring{e} \mid G_e = \mathbf{1}\})$. This decomposition is called the *decomposition of G determined by the edges of $\mathcal{G}$ with trivial stabilizer.*

Now we describe some operations on a graph of groups $\mathcal{G}$. These will be the simplifying moves of the algorithm. The moves will transform

$$\mathcal{G} = (\Gamma, \{G_v\}, \{G_e\}, \{\varphi_e\})$$

into $\mathcal{G}' = (\Gamma', \{G_{v'}\}, \{G_{e'}\}, \{\varphi_{e'}\})$. Much of the data will be the same for $\mathcal{G}$ and $\mathcal{G}'$, so in describing the moves we will usually only record the differences.

2.2 Reducing

If a bonding map at a valence one or two vertex is an isomorphism then there is an obvious simplification that we now describe.

Suppose $v \in V$ has valence one, ie, $\hat{E}(v) = \{e\}$. Suppose further that φ_e is an isomorphism. Then, define $\mathcal{G}'$ by setting $V' = V \setminus \{v\}$ and $\hat{E}' = \hat{E} \setminus \{e, e^{-1}\}$.

Next suppose that $v \in V$ has valence two and that Γ is not a loop, ie, $\hat{E}(v) = \{e, f \neq e^{-1}\}$. Suppose further that φ_e is an isomorphism. Then, define $\mathcal{G}'$ by setting $V' = V \setminus \{v\}$, setting $\hat{E}' = \hat{E} \setminus \{e, e^{-1}\}$, redefining $\partial_0 f$ to be $\partial_1 e$, and redefining φ_f to be $\varphi_{e^{-1}} \circ \varphi_e^{-1} \circ \varphi_f$.

If $\Gamma(\mathcal{G})$ is finite then, since each of these operations decreases the number of vertices, after finitely many operations we obtain a reduced graph of groups that has been obtained from $\mathcal{G}$ by *reducing*. See Figure 1.

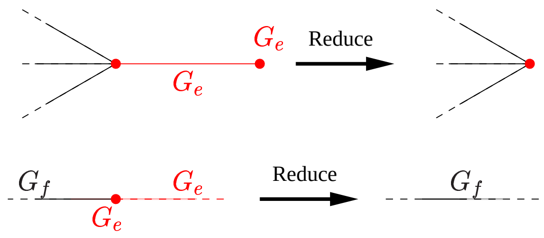

Figure 1: Reducing

In the remaining moves, some vertex group admits a graph of groups decomposition that is compatible with the bonding maps.

2.3 Blowing up

There are two types of blowing up. Suppose first that, for some $v \in V$, $G_v = G'_v * \langle t \rangle$ where t has infinite order, and $\varphi_e(G_e) \subset G'_v$ for $e \in \hat{E}(v)$. Then, define $\mathcal{G}'$ as follows:

- the vertex sets are the same, ie, $V' = V$;
- add a new oriented loop so that $\hat{E}' = \hat{E} \cup \{e_t, e_t^{-1}\}$, with $\partial_0 e_t = \partial_0 e_t^{-1} = v$, $G_{e_t} = \mathbf{1}$, and redefine G_v to be G'_v; and
- bonding maps are given by restricting the codomains of the bonding maps of $\mathcal{G}'$ if necessary.

Secondly, suppose that, for some $v \in V$, $G_v = G'_v * G''_v$ and, for each $e \in \hat{E}(v)$, $\varphi_e(G_e)$ is either contained in G'_v or G''_v. Then, define $\mathcal{G}'$ as follows:

- replace v by two vertices v' and v'', ie, $V' = V \cup \{v', v''\} \setminus \{v\}$;
- add a new oriented edge so that $\hat{E}' = \hat{E} \cup \{e_t, e_t^{-1}\}$ with $\partial_0 e_t = v'$, $\partial_1 e_t = v''$, and $G_{e_t} = \mathbf{1}$;

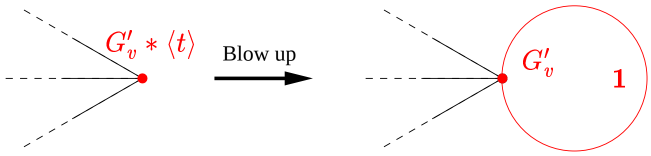

Figure 2: The first type of blowup

- if $e \in \hat{E}(v)$ then $\partial_0 e$ is v' or v'' depending on whether $G_e \subset G_{v'}$ or $G_e \subset G_{v''}$; and
- bonding maps are given by restricting the codomains of the bonding maps of $\mathcal{G}'$ if necessary.

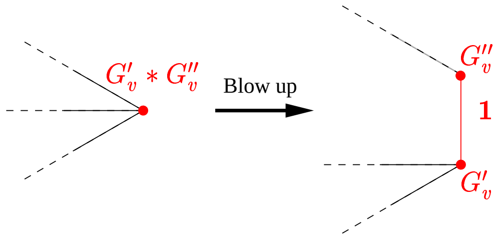

Figure 3: The second type of blowup

In each of these cases, we say that the new graph of groups is obtained from $\mathcal{G}$ by *blowing up*. See Figures 2 and 3.

2.4 Unpulling

Suppose that, for some $v \in V$ and $e \in \hat{E}(v)$, we have $G_v = G'_v * \mathbb{Z}$, $G_e = G'_e * \mathbb{Z}$, $\varphi_e(G'_e) \subset G'_v$, $\varphi_e(\mathbb{Z}) = \mathbb{Z}$, and $\varphi_f(G_f) \subset G'_v$ for $f \in \hat{E}(v) \setminus \{e\}$. Then, define $\mathcal{G}'$ as follows:

- $V' = V$;
- $\hat{E}' = \hat{E}$;
- redefine G_v to be G'_v, G_e to be G'_e; and
- bonding maps are given by restricting codomains of bonding maps of $\mathcal{G}'$ if necessary.

We say that the new graph of groups is obtained from $\mathcal{G}$ by *unpulling*.[4] See Figure 4.

[4]We use the term *unpulling* because the inverse operation *pulls an element of* G'_v *across the edge* e.

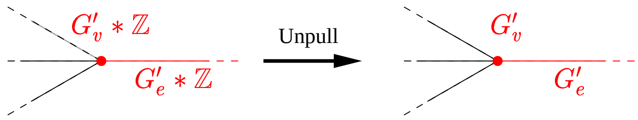

Figure 4: Unpulling

2.5 Unkilling

Suppose that, for some $v \in V$ and $e \in \hat{E}(v)$, $G_v = G'_v * \langle t \rangle$ where t has infinite order, $G_e = G'_e * G''_e$, $\varphi_e(G'_e) \subset G'_v$, $\varphi_e(G''_e) \subset tG'_v t^{-1}$, and $\varphi_f(G_f) \subset G'_v$ for $f \in \hat{E}(v) \setminus \{e\}$. Then, define $\mathcal{G}'$ as follows:

- $V' = V$;
- the oriented edge $\{e, e^{-1}\}$ is replaced with two oriented edges having the same endpoints as e:

$$\hat{E}' = \hat{E} \cup \{e', e'', (e')^{-1}, (e'')^{-1}\} \setminus \{e, e^{-1}\}$$

  with $\partial_0 e' = \partial_0 e'' = \partial_0 e$ and $\partial_1 e' = \partial_1 e'' = \partial_1 e$;
- $G_{e'} = G'_e$, $G_{e''} = G''_e$, G_v is redefined to be G'_v; and
- $\varphi_{e''} = i_{t^{-1}} \circ \varphi_e|_{G''_e}$ and other bonding maps are given by restricting domains and/or codomains of bonding maps of $\mathcal{G}'$ if necessary.

We say that the new graph of groups is obtained from $\mathcal{G}$ by *unkilling*.[5] See Figure 5.

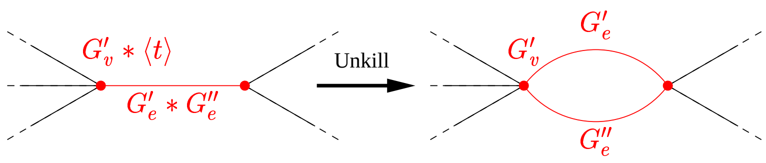

Figure 5: Unkilling

[5]We use the term *unkilling* because the inverse operation *kills a cycle*.

2.6 Cleaving

Suppose that, for some $v \in V$ and $e \in \hat{E}(v)$, $G_v = G'_v * G''_v$ non-trivially, $G_e = G'_e * G''_e$, $\varphi_e(G'_e) \subset G'_v$, $\varphi_e(G''_e) \subset G''_v$, and for $f \in \hat{E}(v) \setminus \{e\}$ either $\varphi_f(G_f) \subset G'_v$ or $\varphi_f(G_f) \subset G''_v$. Then, define $\mathcal{G}'$ as follows:

- v is replaced by two vertices: $V' = V \cup \{v', v''\} \setminus \{v\}$;
- the oriented edge $\{e, e^{-1}\}$ is replaced by two oriented edges:
$$\hat{E}' = \hat{E} \cup \{e', e'', (e')^{-1}, (e'')^{-1}\} \setminus \{e, e^{-1}\}$$
with $\partial_0 e' = v'$, $\partial_0 e'' = v''$, and $\partial_1 e' = \partial_1 e'' = \partial_1 e$;
- for $f \in \hat{E}(v) \setminus \{e\}$, $\partial_0 f = v'$ if $\varphi_e(G_f) \subset G'_v$ and $\partial_0 f = v''$ if $\varphi_e(G_f) \subset G''_v$;
- $G_{v'} = G'_v$, $G_{v''} = G''_v$, $G_{e'} = G'_e$, $G_{e''} = G''_e$; and
- bonding maps are given by restricting domains and/or codomains of bonding maps of $\mathcal{G}'$ if necessary.

We say that the new graph of groups is obtained from $\mathcal{G}$ by *cleaving*. See Figure 6. Each of the operations blowing up, unpulling, unkilling, and cleaving

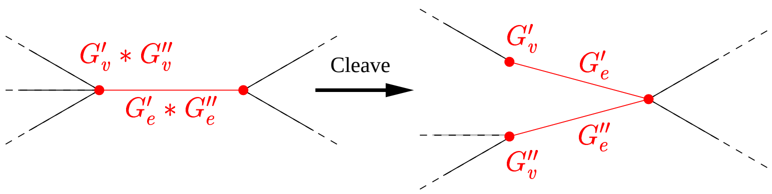

Figure 6: Cleaving

is a *simplification*.

Proposition 2.1 *If $\mathcal{G}'$ is obtained from $\mathcal{G}$ by reducing or simplifying, then $\pi_1(\mathcal{G}')$ and $\pi_1(\mathcal{G})$ are isomorphic.*

Proof In the first type of reducing move, $\pi_1(\mathcal{G}) \cong \pi_1(\mathcal{G}') *_{G_e} G_e$ where the map $G_e \to G_e$ is an isomorphism. By van Kampen's theorem, $\pi_1(\mathcal{G}) \cong \pi_1(\mathcal{G}')$. In all of the other cases, $\mathcal{G}$ is obtained from $\mathcal{G}'$ by a Stallings fold and so $\pi_1(\mathcal{G}) \cong \pi_1(\mathcal{G}')$, see [4, Section 2]. $\square$

Remark 2.2 If $\pi_1(\mathcal{G})$ is not infinite cyclic then the first type of blow up is a composition of a second type of blow up, an unkilling, and a reduction. Therefore, we will not have to consider the first type of blow up.

2.7 Our case

Given a graph of groups, we want to iteratively simplify until the resulting graph of groups can't be simplified. In order to do this algorithmically, we need be able to recognize when a simplification is possible. To this end, we restrict the graphs of groups that we will consider to the case where Γ is a finite graph, ie, where E is finite, and where, for $v \in V$ and $e \in E$, G_v and G_e are finite rank free groups. Such a $\mathcal{G}$ is a *finite graph of finite rank free groups*.

2.8 Labeled graphs

Graphs will have two uses in this paper. The first we have already seen–these are as the underlying graphs of graphs of groups and are denoted by Γ's. The other use will be to represent subgroups of free groups and these will be denoted by Σ's. We now explain this second usage. Let $F_{\mathcal{B}}$ denote the free group with basis $\mathcal{B}$. For $S \subset F_{\mathcal{B}}$, $S^{\pm 1}$ is defined to be $S \cup S^{-1}$ where S^{-1} is $\{s^{-1} : s \in S\}$. A *labeled graph* or a *$\mathcal{B}$–graph* is a connected graph $\Sigma = \Sigma(V, \hat{E}, \mathrm{op}, \partial_0)$ with a *labeling function* $\hat{E} \to \mathcal{B}^{\pm 1}$ such that the label assigned to $\mathrm{op}(c) = c^{-1}$ is the inverse of the label assigned to c. The *$\mathcal{B}$–rose* is a $\mathcal{B}$–graph $R_{\mathcal{B}}$ with one vertex and a bijective labeling function. We identify $\pi_1(R_{\mathcal{B}})$ with $F_{\mathcal{B}}$. (The homotopy class of the path formed by the edge labeled b is identified with b.) There is a natural map $\lambda_\Sigma \colon \Sigma \to R_{\mathcal{B}}$ sending 1–cells to 1–cells and preserving labels and orientations. Since λ_Σ determines the labeling function and *vice versa*, we will call λ_Σ the labeling function as well. The graph Σ is *based* if there is a distinguished vertex $*$. On the level of fundamental groups, the image of λ_Σ is a subgroup of $F_{\mathcal{B}}$ denoted $[(\Sigma, *)]$. If we forget the basepoint then the image is only defined up to conjugacy and Σ determines a conjugacy class $[[\Sigma]]$ of subgroups $F_{\mathcal{B}}$. We say that $(\Sigma, *)$ *represents* $[(\Sigma, *)]$ and that Σ *represents* $[[\Sigma]]$. More generally, if $\vec{\Sigma}$ is a sequence of labeled graphs then a sequence $[[\vec{\Sigma}]]$ of conjugacy classes of subgroups of $F_{\mathcal{B}}$ is determined. If there are basepoints $\vec{*}$ then a sequence $[(\vec{\Sigma}, \vec{*})]$ of subgroups of $F_{\mathcal{B}}$ is determined.

It is well-known, see eg [16, Section 1.A], that a generating set $[\Sigma, *]$ may be obtained as follows. Choose a maximal tree T for Σ and choose orientations for the edges not in T. The generating set is indexed by these oriented edges. Specifically, the generator corresponding to the oriented edge c is the word in $\mathcal{B}^{\pm 1}$ determined by reading the labels of the loop obtained by concatenating the path in T from $*$ to $\partial_0 c$, c, and the path in T from $\partial_1 c$ back to $*$.

2.9 Stallings and Gersten representatives

The *complexity* $\mathsf{c}(\Sigma)$ of the labeled graph Σ is $|E(\Sigma)|$ and the *complexity* $\mathsf{c}(\vec{\Sigma})$ of the sequence $\vec{\Sigma} = \{\Sigma_i\}_{i\in I}$ of labeled graphs is $\sum_{i\in I}\mathsf{c}(\Sigma_i)$. If $\vec{H}$ is a finite sequence of finitely generated subgroups of $F_{\mathcal{B}}$, the *Stallings representative for $\vec{H}$ with respect to $\mathcal{B}$* is the sequence $\vec{\Sigma}_{\mathsf{S}} = \Sigma_{\mathsf{S}}(\vec{H},\mathcal{B})$ of based $\mathcal{B}$–graphs of minimal complexity representing $\vec{H}$. We often omit the $\mathcal{B}$ from the notation. If $\vec{H}$ is represented by a finite sequence $\vec{H} = \{H_i = \langle S_i\rangle\}$ where each S_i is a finite set of words in $\mathcal{B}^{\pm 1}$, then there is an algorithm due to Stallings [27] to find $\vec{\Sigma}_{\mathsf{S}}$ from $\{S_i\}$, see also [10].

If $\vec{\mathcal{H}}$ is a sequence of conjugacy classes of non-trivial subgroups of $F_{\mathcal{B}}$, the *Stallings representative for $\vec{\mathcal{H}}$ with respect to $\mathcal{B}$* is the $\mathcal{B}$–graph $\Sigma_{\mathsf{S}}(\vec{\mathcal{H}})$ of minimal complexity representing $\vec{\mathcal{H}}$. In fact, if $\vec{\mathcal{H}}$ is represented by $\vec{H} = \{H_i = \langle S_i\rangle\}$ as above then $\Sigma_{\mathsf{S}}(\vec{\mathcal{H}})$ is the sequence $\mathsf{core}(\Sigma_{\mathsf{S}}(\vec{H}))$ of cores of elements of $\Sigma_{\mathsf{S}}(\vec{H})$. Recall that if Σ is graph then the *core* of Σ, denoted $\mathsf{core}(\Sigma)$, is the union of all immersed circuits in Σ, see also Section 9.2.1. In particular, there is an algorithm to find $\mathsf{core}(\Sigma_{\mathsf{S}}(\vec{H}))$ from $\{S_i\}$ as well as a sequence $\vec{h}$ of elements of $F_{\mathcal{B}}$ such that $\mathsf{core}(\Sigma_{\mathsf{S}}(\vec{H})) = \Sigma_{\mathsf{S}}(\vec{H}^{\vec{h}})$ where $\vec{H}^{\vec{h}}$ is the sequence of groups obtained by conjugating a component of $\vec{H}$ with the corresponding component of $\vec{h}$. In fact, if Σ is a component of $\Sigma_{\mathsf{S}}(\vec{H})$ then the corresponding component h of $\vec{h}$ can be taken to be the inverse of the word read along the shortest path from the basepoint $*$ of Σ to $\mathsf{core}(\Sigma)$. It is convenient to also allow conjugacy classes of trivial groups in $\vec{\mathcal{H}}$. Since the core of a tree is empty, we take the Stallings representative of the conjugacy class of the trivial group to be the empty set.

A *Gersten representative* $\Sigma_{\mathsf{G}}(\vec{\mathcal{H}})$ *for* $\vec{\mathcal{H}}$ is a sequence of $\mathcal{B}$–graphs of minimal complexity among sequences of $\mathcal{B}$–graphs representing $\alpha\vec{\mathcal{H}}$ as α varies over $\mathsf{Aut}(F_{\mathcal{B}})$. (If $\vec{H}$ represents $\vec{\mathcal{H}}$, then $\alpha\vec{H}$ represents $\alpha\vec{\mathcal{H}}$, where α is applied coordinate-wise.) If $\vec{\mathcal{H}}$ is represented by $\vec{H} = \{H_i = \langle S_i\rangle\}$ as above, then there is an algorithm that produces a $\Sigma_{\mathsf{G}}(\vec{\mathcal{H}})$ as well as an automorphism α such that $\mathsf{core}(\Sigma_{\mathsf{S}}(\alpha\vec{H})) = \Sigma_{\mathsf{G}}(\vec{\mathcal{H}})$, see [14], [18], and also Section 9.5.

Example 2.3 If $\mathcal{B} = \{a,b\}$ and if $H = \langle aaba^{-1}, ab^{-1}abba^{-1}\rangle$, then the graph $(\Sigma,*)$ pictured in Figure 7 represents H. (The open arrows denote 'a' and the closed 'b'.) The graph $(\Sigma_{\mathsf{S}},*)$ is the Stallings representative of H. For the automorphism $\alpha\colon F_{\mathcal{B}} \to F_{\mathcal{B}}$ given by $a \mapsto ab^{-1}$, $b \mapsto b$, a Gersten representative Σ_{G} of H is the core of the Stallings representative for αH.

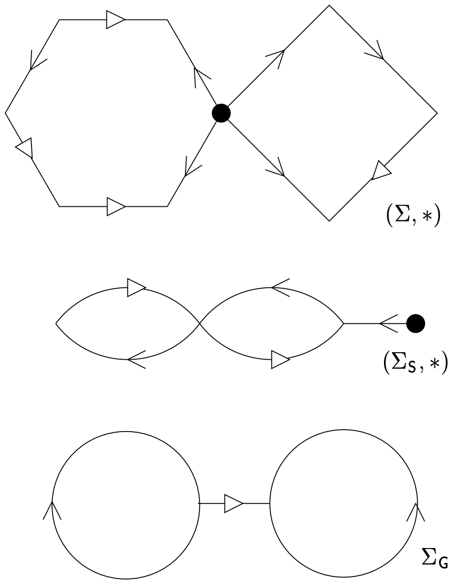

Figure 7

Notation 2.4 If $\vec{\Sigma}' = \{\Sigma_i'\}_{i\in I}$ and $\vec{\Sigma}'' = \{\Sigma_j''\}_{j\in J}$ are sequences of $\mathcal{B}$–graphs and if $\{0\} = I \cap J$ then $\vec{\Sigma}' \vee \vec{\Sigma}'' = \{\Sigma_k\}_{k\in I\cup J}$ is a sequence of graphs with labels in $\mathcal{B}$ of the following form:

$$\Sigma_k = \begin{cases} \Sigma_0' \vee \Sigma_0'', & \text{if } k = 0; \\ \Sigma_i', & \text{if } k \in I \setminus \{0\}; \text{ and} \\ \Sigma_j'', & \text{if } k \in J \setminus \{0\}. \end{cases}$$

Definition 2.5 Let $\vec{\Sigma} = \{\Sigma_i\}_{i\in I}$ be a sequence of $\mathcal{B}$–graphs.

(1) If there is a non-trivial partition $\mathcal{B} = \mathcal{B}' \sqcup \mathcal{B}''$ such that, for each $i \in I$, the labels of Σ_i are either all in $\mathcal{B}'$ or all in $\mathcal{B}''$, then we say that $\vec{\Sigma}$ can be *visibly blown up*.

(2) Suppose that $b \in \mathcal{B}$ appears as a label in only one element Σ of $\vec{\Sigma}$ and that only one oriented edge c_0 of Σ has label b. Suppose further that $c_0 \subset \mathsf{core}(\Sigma)$.

 (a) If c_0 does not separate Σ then we say that $\vec{\Sigma}$ can be *visibly unpulled*.

 (b) If c_0 does separate Σ then we say that Σ can be *visibly unkilled*.

(3) If there is a non-trivial partition $\mathcal{B} = \mathcal{B}' \sqcup \mathcal{B}''$ such that $\vec{\Sigma} = \vec{\Sigma}' \vee \vec{\Sigma}''$ where $\vec{\Sigma}'$ is a sequence of $\mathcal{B}'$–graphs and $\vec{\Sigma}''$ is a sequence of $\mathcal{B}''$–graphs then we say that $\vec{\Sigma}$ can be *visibly cleaved*.

If $\vec{\Sigma}$ can be visibly blown up, visibly unpulled, visibly unkilled, or visibly cleaved, then we say that $\vec{\Sigma}$ can be *visibly simplified.*

After the statement of the next proposition, we can describe the algorithm. The proof of this proposition is almost obvious, but requires some bookkeeping which is postponed until the appendix. This proposition is subsumed into Proposition A.3.

Notation 2.6 Let $\mathcal{G}$ be a graph of finite rank free groups with notation as in Section 2.1. For $v \in V$, let $\vec{H}(v)$ denote the sequence of subgroups of G_v represented by $\{\varphi_e(G_e) : e \in \hat{E}(v)\}$. Also let $\vec{\mathcal{H}}(v)$ denote the sequence of conjugacy classes of subgroups of G_v represented by $\vec{H}(v)$.

Proposition 2.7 *Suppose that for some $v \in V$, $\Sigma_{\mathsf{G}}(\vec{\mathcal{H}}(v))$ can be visibly simplified. Then, $\mathcal{G}^{out} \sim \mathcal{G}$ can be algorithmically found such that $\mathcal{G}^{out}$ can be simplified.*

2.10 The algorithm

Here is the algorithm. See Figure 8 for a flow chart. More details on the algorithm are given in Section 9, Section 10, and the appendix.

Step 0 Input $\mathcal{G}$, a finite graph of finite rank free groups.

Step 1 Reduce $\mathcal{G}$.

Step 2 If, for some $v \in V$, $\Sigma_{\mathsf{G}}(\vec{\mathcal{H}}(v))$ can be visibly simplified, then replace $\mathcal{G}$ by a simplified conjugate and return to Step 1. Else, done.

The main result of the paper is:

Theorem 2.8 *Suppose that a finite graph $\mathcal{G}$ of finite rank free groups is input into the above algorithm and that $\mathcal{G}^{out}$ is output. Then, the decomposition of $\pi_1(\mathcal{G})$ determined by the edges of $\mathcal{G}^{out}$ with trivial stabilizer is the Grushko decomposition of $\pi_1(\mathcal{G})$.*

The proof of Theorem 2.8 is found in Section 10.

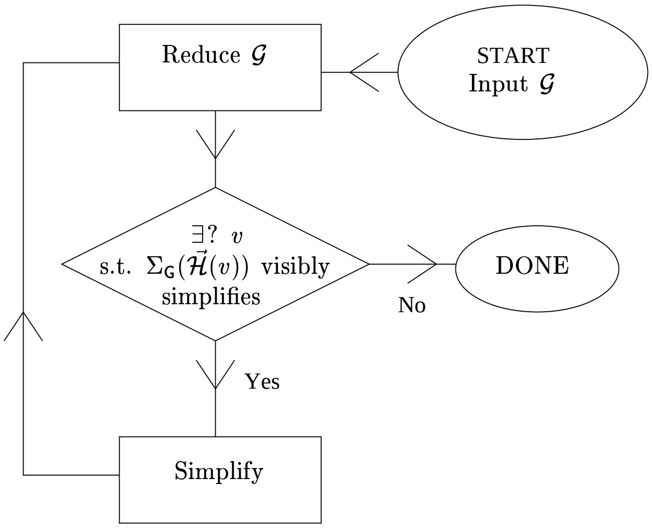

Figure 8: Flow chart

2.11 Relative version

In this section we describe a relative version of Theorem 2.8. If H is a subgroup of a group G, then we say that G is *freely decomposable rel* H if there is a free decomposition $G = G' * G''$ with $H \subset G'$ and G'' non-trivial. Otherwise, G is *freely indecomposable rel* H. The relative version of Grushko's theorem (Theorem 1.1) is:

Theorem 2.9 *Suppose H is a subgroup of the finitely generated group G. Then, G is a free product $G_H * G_1 * \cdots * G_n * F$ where $H \subset G_H$, G_H is freely indecomposable rel H, G_i for $1 \leq i \leq n$ is freely indecomposable and not free, and F is a finite rank free group.*

The subgroup G_H is unique in this *the Grushko decomposition of G rel H.* Up to reordering and conjugation, the G_i, $1 \leq i \leq n$, are unique. Also, the rank of F is an invariant of the pair $H \subset G$.

Suppose now that $\mathcal{G}$ is a finite graph of finite rank free groups, that $v_0 \in V$ has valence one with incident edge e_0, and that φ_{e_0} is an isomorphism. We are going to describe a slight modification of the algorithm of Section 2.10 that produces the Grushko decomposition of $\pi_1(\mathcal{G})$ rel G_{v_0}. Intuitively, in the modified algorithm we only reduce or visibly simplify only if the special edge group G_{e_0} is unchanged. Specifically, we modify the algorithm as follows.

Steps 0 and 1 are replaced by:

Step 0′ Input $\mathcal{G}$, a finite graph of finite rank free groups as above.

Step 1′ Reduce $\mathcal{G}$ rel v_0, ie, apply the reducing moves displayed in Figure 1 only if $e \notin \{e_0, e_0^{-1}\}$.

To describe the modification of Step 2, we need a definition. In Definition 2.5(2), the component Σ of $\vec{\Sigma}$ is *special*. In Definition 2.5(3), the component of $\vec{\Sigma}$ corresponding to $\Sigma_0 = \Sigma_0' \vee \Sigma_0''$ in Notation 2.4 is *special*. In Definition 2.5(1), none of the components of $\vec{\Sigma}$ are special. The components of $\Sigma_{\mathsf{G}}(\vec{\mathcal{H}}(v))$ are parametrized by the set $\hat{E}(v)$ of edges incident to v. If $\Sigma_{\mathsf{G}}(\vec{\mathcal{H}}(v))$ can be visibly simplified, then the edge corresponding to the special component is *special*. If the special edge e is not in $\{e_0, e_0^{-1}\}$, then the resulting simplification will not change G_{e_0}. Step 2 of the algorithm is replaced by:

Step 2′ If, for some $v \in V$, $\Sigma_{\mathsf{G}}(\vec{\mathcal{H}}(v))$ can be visibly simplified and the special edge is not in $\{e_0, e_0^{-1}\}$, then replace $\mathcal{G}$ by a simplified conjugate and return to Step 1′. Else, done.

Theorem 2.10 *Suppose that $\mathcal{G}$ as above is input into the modified algorithm and that $\mathcal{G}^{out}$ is output. Then, the decomposition of $\pi_1(\mathcal{G})$ determined by the edges of $\mathcal{G}^{out}$ with trivial stabilizer gives the Grushko decomposition of $\pi_1(\mathcal{G})$ rel G_{v_0}.*

The proof of the relative version requires only minor notational changes to the proof of Theorem 2.8 and is left to the reader.

3 Laminated square complexes and models

This section contains a discussion of certain laminated two complexes called models whose 2–cells are squares. For an interesting study of complexes built from squares see [6].

Let I denote the unit interval $[0,1]$. An *n–cube* is a metric space isometric to I^n. A metric space is a *cube* if it is an n–cube for some n. A *cube complex* is a union of cubes glued by isometries of faces. A finite dimensional cube complex X admits a maximal metric such that the inclusion $C \to X$ is a local isometry for each cube C of X [5]. A *square* is a 2–cube, and we only have need to consider *square complexes*, ie, cube complexes of dimension at most two. A *graph* is a 1–dimensional cube complex[6], a *tree* is a simply connected graph and a *forest* is a disjoint union of trees.

A decomposition of a square is *standard* if it is induced by projection to a codimension–1 face. A decomposition of a 1–cube is *standard* if all decomposition elements are points. It is *trivial* if the only decomposition element is the 1–cube itself. A *laminated square complex* is a simply connected square complex X with a decomposition $\mathcal{D}$ such that:

(M1) The link of every vertex of X is a flag complex.

(M2) For each square C of X, the induced decomposition of C is standard. In other words, the decomposition of C whose elements are the components of C intersected with elements of $\mathcal{D}$ is standard. For each 1–cube of X, the induced decomposition is either standard or trivial.

In this context, (M1) means that every link is a simplicial graph with no circuits of length three. A decomposition element is also called a *leaf*.

Proposition 3.1 *Let $(X,\mathcal{D})$ be a laminated square complex. Then,*

(1) *X is contractible and*

(2) *leaves are forests.*

Proof (M1) implies that the the metric on X is $CAT(0)$ [5]. Hence, (1). A vertex of a square with a standard decomposition is contained in exactly one edge that is contained in a leaf. So, links of vertices in X are bipartite and leaves are totally geodesic. In particular, leaves are 1–dimensional and contractible. □

A *model* is a laminated square complex $(X,\mathcal{D})$ such that:

(M3) Leaves are connected.

The next proposition is immediate from definitions.

Proposition 3.2 *Let $(X,\mathcal{D})$ be a laminated square complex and let $\hat{\mathcal{D}}$ be the decomposition of X whose elements are the connected components of elements of $\mathcal{D}$. Then, $(X,\hat{\mathcal{D}})$ is a model.* □

[6] 1–dimensional CW–complexes and 1–dimensional cube complexes are both called *graphs*. Since we will only be using combinatorial properties, the distinction is not important to us.

Definition 3.3 We say that $(X, \hat{\mathcal{D}})$ is *induced by* $(X, \mathcal{D})$.

Proposition 3.4 *If $(X, \mathcal{D})$ is a model, then the decomposition space $X/\mathcal{D}$ is a tree.*

Proof Since a leaf is totally geodesic, it intersects each square C in a connected set. In particular, the decomposition $C \cap \mathcal{D}$ of C obtained by intersecting C with elements of $\mathcal{D}$ is standard, $I = C/C \cap \mathcal{D} \to X/\mathcal{D}$ is injective, and $X/\mathcal{D}$ is naturally a graph. Leaves are connected, and so $X \to X/\mathcal{D}$ is π_1–surjective. Thus, $X/\mathcal{D}$ is a tree. □

If $(X, \mathcal{D})$ is a model, we say that X *is a model for the tree* $S = X/\mathcal{D}$ or that $X \to S$ *is a model.* The preimage in X of $s \in S$ is X_s. If c is an edge of S and if $s, s' \in \mathring{c}$, then X_s and $X_{s'}$ have the same isomorphism type X_c. We will sometimes abuse notation and identify X_c with X_s for $s \in \mathring{c}$.

Notation 3.5 If the group G acts on the set X and if S is a subset of X, then $G_{X,S}$ is the stabilizer of S, ie, the subgroup elements $g \in G$ such that $g(S) = S$. If $S = \{s\}$ then we also write $G_{X,s}$ for $G_{X,S}$. We will suppress the X if the space is understood.

An action of a group G on $(X, \mathcal{D})$ is an (isometric) action of G on X permuting cubes and decomposition elements. In this case, $(X, \mathcal{D})$ is a G*–model.* There is an induced action of G on $S = X/\mathcal{D}$. The quotients X/G and S/G are denoted $\bar{X}$ and $\bar{S}$ respectively. We say that G *acts without inversions* if, for all cubes $C \subset X$, G_C fixes C pointwise. By subdividing X, we may arrange that G acts without inversions. Hence, *we always assume that our actions are without inversions.* Note that the space X_s is G_s–invariant and X_c is G_c–invariant.

4 Trees

We review some tree basics. A G–tree S is *minimal* if it has no proper G–invariant subtrees. It is *trivial* if has a fixed point. In this case we also say that the G–action is *elliptic.* If S has a unique minimal invariant G–subtree, then it is denoted S_G. This occurs, for example, if G contains a *hyperbolic* element, that is an element that fixes no point of S [1]. A *morphism* $S \to T$ of G–trees is a simplicial G–map. It is *strict* if no edge is mapped to a point. If there is a morphism $S \to T$, then S *resolves* T.

If e is an edge of S, then we say that S' *is obtained from* S *by collapsing* e if S' is the result of equivariantly collapsing e. A morphism $S \to S'$ is a *collapse* if S' is obtained from S by iteratively collapsing edges. If edges e_1 and e_2 of S share the vertex s, then we say that S' *is obtained from folding* e_1 *and* e_2 if S' is the result of equivariantly identifying e_1 and e_2 with an isometry fixing s. The resulting morphism $S \to S'$ is a *fold.*

5 Examples of Models

Example 5.1 First a non-example. Glue two squares along three sides and laminate so that restrictions to squares are standard and so that the unglued sides form a leaf. The result is not a laminated square complex even though it is contractible (there are vertices whose links consist of distinct edges with the same endpoints–such a link is not a simplicial graph). Notice that not all leaves are trees.

Example 5.2 The quotient $\bar{X} \to \bar{S}$ of a G–model $X \to S$ with G free of rank 2 is depicted in Figure 9. The preimage in $\bar{X}$ of a point in the interior of the edge of $\bar{S}$ is isomorphic to a circle. The preimage of the vertex is a 'pair of eyeglasses'. The stabilizer of a vertex of S is free of rank two; the stabilizer of an edge of S is infinite cyclic.

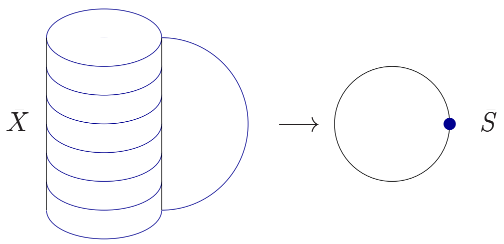

Figure 9

Example 5.3 If S and T are G–trees, then

$$S \times T = \cup\{e \times f : e \text{ is an edge of } S,\ f \text{ is an edge of } T\}$$

is a union of squares with projection maps $q_S\colon\ S \times T \to S$ and $q_T : S \times T \to T$. The induced decomposition with quotient S (respectively T) gives $S \times T$ the structure of a model for S (respectively T).

Example 5.4 Let $(X, \mathcal{D})$ be a G–model and let Y be a simply connected G–subcomplex of X. Then, Y with the decomposition $\mathcal{D}(Y) = \{D \cap Y \mid D \in \mathcal{D}\}$ is a laminated square complex and $Y/\mathcal{D}(Y) \to X/\mathcal{D}$ is an inclusion. If we let $\hat{\mathcal{D}}(Y)$ be the decomposition of Y induced by $\mathcal{D}(Y)$ (see Definition 3.3), then $(Y, \hat{\mathcal{D}}(Y)))$ is a model and $Y/\hat{\mathcal{D}}(Y) \to X/\mathcal{D}$ is a morphism. We call $\hat{\mathcal{D}}(Y)$ the *restricted decomposition of* Y.

Example 5.5 *Main Example* Suppose that S and T are G–trees and that, for each $s \in S$, there is a unique minimal G_s–invariant subtree T_s of T. The union $X_S(T) = \bigsqcup_{s \in S}(\{s\} \times T_s)$ is a subcomplex of $S \times T$ and is simply-connected (being a union of simply-connected spaces along simply-connected spaces). $X_S(T) \to S$ is a model as is $X_S(T) \to \hat{T}$ where $\hat{T}$ is the quotient of the decomposition induced from $X_S(T) \subset S \times T \to T$ by restriction (see Example 5.4). If $\overline{S} = S/G$ and all $\overline{T}_s = T_s/G_s$ are compact, then $\overline{X_S(T)} = X_S(T)/G$ is also compact.

Proposition 5.6 *$\hat{T}$ as in Example 5.5 is minimal.*

Proof We may identify T_s with the image of the injection $\{s\} \times T_s \to \hat{T}$. By construction, $\hat{T} = \cup_{s \in S} T_s$. If $\hat{t} \in \hat{T}$ is not contained in an invariant G–subtree R of $\hat{T}$ and if $\hat{t} \in T_s$, then $R \cap T_s$ is a proper G_s–invariant subtree of T_s, contradiction. □

6 Operations on models

In this section, we assume that G is a group, S is a G–tree, and $X \to S$ is a model. We will describe operations on X. In each case, the result X' is a model for S' where S' resolves S. The operations are geometric generalizations of the simplifications of Section 2.

6.1 0–Simplifying

Let $C = I \subset X$ be a 1–cube and set $C_0 = \{0\}$. Suppose that

- C meets cubes of X other than faces of C only in $\{1\}$; and
- the restriction of the decomposition to C is $\{C\}$.

Let X' be the result of equivariantly replacing C by $\{1\}$, ie,

$$X' = X \setminus \cup_{g \in G} g \cdot [0, 1).$$

We say that X' with the restricted decomposition is obtained from X by *0–simplifying C from C_0*. Here $S' = S$. See Figure 10.

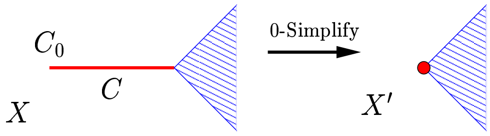

Figure 10

6.2 I–Simplifying

Let $C = I \times I^n \subset X$ $(n = 0$ or $1)$ be a cube and set $C_0 = \{0\} \times I^n$. Suppose that

- C meets cubes of X other than faces of C only in $\{1\} \times I^n$; and
- C_0 is a decomposition element.

Let X' be the result of equivariantly replacing C by $\{1\} \times I^n$. We say that X' with the restricted decomposition is obtained from X by *I–simplifying* C *from* C_0. $S \to S'$ is a collapse. See Figure 11.

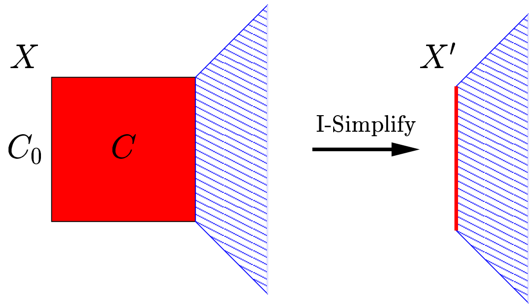

Figure 11

6.3 II–Simplifying

Let $C = I^2 \subset X$ be a square and set $C_0 = \{0\} \times I$. Suppose that

- C meets cubes of X other than faces of C only in
$$(\{1\} \times I) \cup (I \times \{1\});$$
- it is not possible to I–simplify from C_0, ie, $[0,1) \times \{1\}$ meets a cube other than a face of C; and
- C_0 is an element of the decomposition restricted to C.

The model X' with restricted decomposition elements obtained by equivariantly replacing C by

$$(\{1\} \times I) \cup (I \times \{1\})$$

is the result of *II–simplifying* C *from* C_0. Note that $S' = S$. See Figure 12.

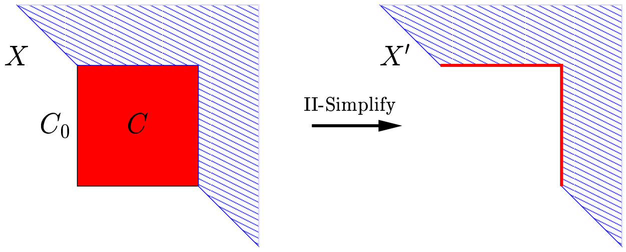

Figure 12

6.4 III–Simplifying

Let $C = I^2 \subset X$ be a square and set $C_0 = \{0\} \times I$. Suppose that

- C meets cubes of X other than faces of C only in $(\partial I \times I) \cup (I \times \{1\})$;
- it is not possible to I–, or II–simplify from C_0; and
- C_0 is an element of the decomposition restricted to C.

The model X' with restricted decomposition elements obtained by equivariantly replacing C by $(\partial I \times I) \cup (I \times \{1\})$ is the result of *III–simplifying* C *from* C_0. Note that $S' \to S$ is a non-trivial fold. See Figure 13.

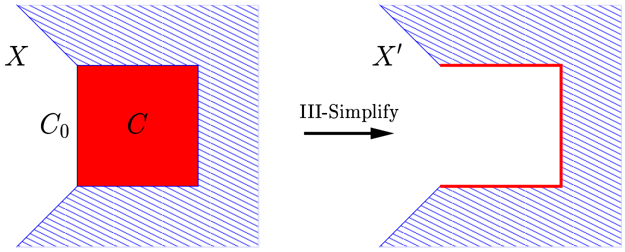

Figure 13

6.5 Blowing up

If there is a cube $C = I \subset X$ such that $\mathring{I}$ meets no cube other than faces of C and such that the decomposition restricted to C is $\{C\}$ then we may refine the decomposition by G–equivariantly replacing the decomposition element X_C containing C with

$$\{\{g \cdot t\} \mid t \in \mathring{I},\ g \in G_{X_C}\} \sqcup \{\text{components of } X_C \setminus G_{X_C} \cdot \mathring{I}\}.$$

We say that $X' \to S'$ is obtained from $X \to S$ by *blowing up* C. The induced map $S' \to S$ collapses to points the edges of S' corresponding to the orbit of C, explaining the term "blowing up". See Figure 14.

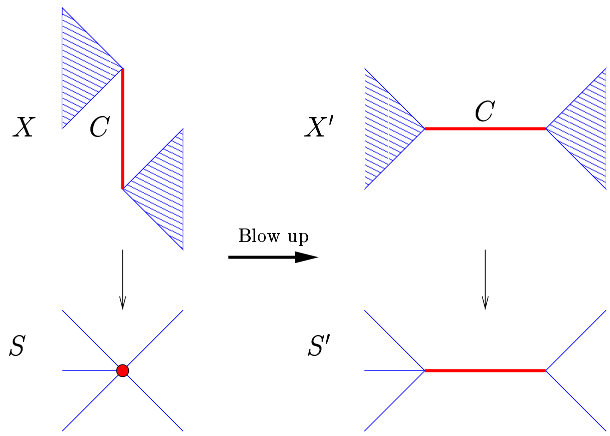

Figure 14

Definition 6.1 Let $S' \to S$ a morphism. If there is a model $Y \to S$ and a cube C of Y with face C_0 such that 0–simplifying C from C_0 yields $Y' \to S'$, then we say that S' is obtained from S by *0–G_0–simplifying* or equivalently by *0–simplifying over* G_0 where G_0 is the stabilizer of C_0 in Y. The definitions for I–G_0–, II–G_0–, and III–G_0–simplifying are analogous. If there is a model $Y \to S$ with a cube C such that $Y' \to S'$ is the result of blowing up C, then we say that S' is obtained from S by *G_0–blowing up* or equivalently by *blowing up over* G_0 where G_0 is the stabilizer of C in Y. In each of these cases, we say that S' is obtained from S by *simplifying over* G_0 or just by *simplifying* if G_0 is understood.

Remark 6.2 Since the identity map $S \to S$ is an example of a model, blowing up as defined in Section 2.3 is an example of **1**–blowing up.

7 Generalized Shenitzer–Swarup

In the next lemma, we use the notation of Example 5.5.

Lemma 7.1 *Suppose that G is a group and that S and T are G–trees such that*

(1) *$\overline{S}$ is compact;*

(2) *for each $s \in S$ there is a unique minimal cocompact G_s–subtree T_s of T; and*

(3) *for each edge $f \subset T$, the action of G_f on S is elliptic.*

Then, there is a sequence

$$\{X_s(T) = X_0 \to S = S_0, X_1 \to S_1, \cdots, X_N \to S_N\}$$

of I–, II–, and III–simplifications such that $(X_N)_{\hat{f}}$ is a point for each edge $\hat{f}$ of $\hat{T}$. Further, all the simplifications are over subgroups of edge stabilizers of T.

Proof Recall from Example 5.5 that $X_0 \to \hat{T}$ is obtained by restricting $S \times T \to T$. It is not possible to 0–simplify $X_0 \to S_0$. Indeed, in order to 0–simplify X_0 there would have to be a 1–cube as in the definition of 0–simplifying. By the construction of X_0, the restriction of the decomposition giving $X_0 \to \hat{T}$ to this 1–cube is standard. This is impossible since, by Proposition 5.6, $\hat{T}$ is minimal. Further, if X_i is obtained from X_0 by a sequence of I–, II–, and III–simplifications, then the restriction to X_i of the decomposition giving $X_0 \to \hat{T}$ still has decomposition space $\hat{T}$. (It is the decomposition space S_i of the restriction of $S \times T \to S$ that can change.) In particular, it is also not possible to 0–simplify X_i.

Suppose we have constructed the sequence

$$\{X_S(T) = X_0 \to S = S_0,\ X_1 \to S_1,\ \cdots,\ X_i \to S_i\}.$$

We will describe how to proceed. Let $\hat{t}$ be a point in the interior of an edge of $\hat{T}$. The preimage $(X_i)_{\hat{t}}$ of $\hat{t}$ under $X_i \to \hat{T}$ is a $G_{\hat{t}}$–subtree of $S \times \{\hat{t}\}$. By (1) and (2), X_0, and so also X_i, is cocompact. Therefore, $(X_i)_{\hat{t}}/G_{\hat{t}}$ is compact. Since T' resolves T, by (3) the action of $G_{\hat{t}}$ on $S \times \{\hat{t}\}$, and hence also on $X_{\hat{t}}$ is elliptic. We see that $X_{\hat{t}}/G_{\hat{t}}$ is a finite tree. If this finite tree is not a single vertex then $X_{\hat{t}}$ contains a valence one vertex whose stabilizer equals the stabilizer of the incident edge.

If there is such a valence one vertex, then this vertex is contained in a cube $C_0^i = I \subset X_i$ that projects to a point in S_i. In this case, simplify from C_0^i to obtain $X_{i+1} \to S_{i+1}$. Stop if, for each edge $\hat{f}$ of $\hat{T}$, $\hat{T}_{\hat{f}}$ is a point.

The process must eventually stop since there are only finitely many G–orbits of cubes in $X_S(T)$.

The final claim of the lemma follows from the observation that C_0^i projects to an edge of $\hat{T}$ and so the stabilizer G_0^i of C_0^i fixes an edge of $\hat{T}$. Since $\hat{T}$ resolves T, G_0^i fixes an edge of T as well. □

Theorem 7.2 *Let S be a cocompact G–tree with finitely generated edge stabilizers and with G finitely generated. Suppose that G splits over a finite group. Then, S may be iteratively I–, II–, and III–simplified and then blown up to a G–tree S' such that the decomposition of G given by edges of S' with finite stabilizer is non-trivial. Further, the simplifications and blow ups are all over finite groups. In particular, all point stabilizers of S' are finitely generated.*

Proof Choose T to be a minimal G–tree with one orbit of edges and with finite edge stabilizers. If an edge stabilizer of S is finite then we may set $S' = S$ and we are done. We may assume then that the edges stabilizers of S are infinite.

Since edge stabilizers of S are finitely generated and since G is finitely generated, for each $s \in S$, G_s is finitely generated, see for example [8, Lemma 32]. The edge stabilizers of T are finite and by assumption G_s is infinite and so either G_s is contained in a unique vertex stabilizer of T or some element of G_s acts hyperbolically on T. In particular, there is a unique minimal cocompact G_s–subtree T_s of T. Therefore we may apply Lemma 7.1 to simplify $X_S(T)$ to obtain $X_N \to S_N$.

Blow up $X_N \to S_N$ to obtain $X' \to S'$. Since $(X_N)_{\hat{f}}$ is a point for each edge $\hat{f}$ of $\hat{T}$, S' resolves $\hat{T}$. By Proposition 5.6, $\hat{T}$ is minimal and so S' is non-trivial. □

Corollary 7.3 (Generalized Shenitzer–Swarup) *Let S be a minimal G–tree with finitely generated edge stabilizers. Suppose that G splits over $\mathbf{1}$. Then, S may be iteratively $\mathbf{1}$–simplified to a tree S' such that the decomposition of G determined by the edges of S' with trivial stabilizer is non-trivial.* □

The focus of this paper is on splittings over $\mathbf{1}$, ie, on free decompositions. In a future paper, we plan to explore splittings over small groups. Here is a sample analogue of Lemma 7.1 in that setting. Again, we use the notation of Example 5.5.

Theorem 7.4 *Suppose that G is a freely indecomposable group. Suppose that S and T are G–trees such that*

1. *$\overline{S}$ is compact;*
2. *for each $s \in S$, there is a unique minimal cocompact G_s–subtree T_s of T; and*
3. *edge stabilizers of T are infinite cyclic and T has one orbit of edges.*

Then, $X_S(T) \to S$ may be iteratively simplified to $X' \to S'$ where, for each edge $\hat{f}$ of $\hat{T}$, $X'_{\hat{f}}$ is either a point with infinite cyclic stabilizer or a line with infinite cyclic stabilizer with generator acting by a non-trivial translation. Further, these simplifications are over $\mathbf{1}$ or $\mathbb{Z}$.

Proof Since $\hat{T}$ resolves T, $\hat{T}$ is minimal, and G is freely indecomposable, it follows that the edge stabilizers of $\hat{T}$ are infinite cyclic. Thus, for $\hat{f}$ an edge of $\hat{T}$, $(X_S(T))_{\hat{f}}$ is a $\mathbb{Z}$–tree. If this tree is not a point or a line, then it has a valence one vertex and a simplification is possible. Iterate. □

8 Algebraic consequences

This section will be needed for algorithmic questions. We use the notation of Section 6. The goal is to describe the effect of simplifying on edge and vertex stabilizers.

Definition 8.1 Let S be a G–tree and let $X \to S$ be a model with a cube C with face C_0 such that $X' \to S'$ is the result of III–simplifying C from C_0. Further, let e be the image in S of C, let s_0 be the image of C_0, and let $s_e \in \mathring{e}$. Set $C_{s_e} = X_{s_e} \cap C$. Denote by $\overline{C}_{s_e}$ the image of C_{s_e} in $\overline{X}_{s_e}$ and by $\overline{C}_0$ the image of C_0 in $\overline{X}_{s_0}$. There are three cases.

1. $\overline{C}_{s_e}$ separates $\overline{X}_{s_e}$ and $\overline{C}_0$ separates $\overline{X}_{s_0}$.
2. $\overline{C}_{s_e}$ separates $\overline{X}_{s_e}$, but $\overline{C}_0$ does not separate $\overline{X}_{s_0}$.
3. $\overline{C}_{s_e}$ does not separate $\overline{X}_{s_e}$ and $\overline{C}_0$ does not separate $\overline{X}_{s_0}$.

Let G_0 be the stabilizer in X of C_0. In Case (1), we say the simplification is a G_0*–cleaving*, in Case (2) a G_0*–unkilling*, and in Case (3) a G_0*–unpulling*.

For the moment, we forget models and make some purely algebraic definitions. Here G_0 is a subgroup of the group G and $\vec{\mathcal{H}}$ is a sequence of conjugacy classes of subgroups of G.

Definition 8.2 If there are subgroups $G' \subset G$ and $G'' \subset G$ containing G_0 such that

- $G = G' *_{G_0} G''$, ie, the natural map $G' *_{G_0} G'' \to G$ is an isomorphism;
- some $\mathcal{H}_0 \in \vec{\mathcal{H}}$ has a representative $H_0 \in \mathcal{H}_0$ with subgroups $H_0' \subset H_0$ and $H_0'' \subset H_0$ satisfying
  - $G_0 \subset H_0' \subset G'$;
  - $G_0 \subset H_0'' \subset G''$; and
  - $H_0 = H_0' *_{G_0} H_0''$; and
- for all $\mathcal{H} \neq \mathcal{H}_0$ in $\vec{\mathcal{H}}$, there is $H \in \mathcal{H}$ such that either $H \subset G'$ or $H \subset G''$

then we say that $\vec{\mathcal{H}}$ *can be* G_0*–cleaved in* G.

Definition 8.3 If there is a subgroup $G' \subset G$ containing G_0, a monomorphism $h\colon G_0 \to G'$, and $t \in G$ such that

- $G = G' *_h = \langle G', t \mid tgt^{-1} = h(g),\ g \in G_0 \rangle$;
- some $\mathcal{H}_0 \in \vec{\mathcal{H}}$ has a representative $H_0 \in \mathcal{H}_0$ with subgroups H_0' and H_0'' satisfying
  - $G_0 \subset H_0' \subset G'$;
  - $h(G_0) \subset H_0'' \subset G'$; and
  - $H_0 = H_0' *_{G_0} t^{-1} H_0'' t$; and
- for all $\mathcal{H} \neq \mathcal{H}_0$ in $\vec{\mathcal{H}}$ there is $H \in \mathcal{H}$ with $H \subset G'$

then we say that $\vec{\mathcal{H}}$ *can be* G_0*–unkilled in* G.

Definition 8.4 If there is a subgroup $G' \subset G$ containing G_0, a monomorphism $h\colon G_0 \to G'$, and $t \in G$ such that

- $G = G' *_h$;
- some $\mathcal{H}_0 \in \vec{\mathcal{H}}$ has a representative $H_0 \in \mathcal{H}_0$ with a subgroup H_0' satisfying
  - $t \in H_0$;
  - $H_0' \subset G'$;
  - $G_0 \subset H_0'$;

  - $h(G_0) \subset H_0'$; and
  - $H_0 = H_0' *_h$; and
- for all $\mathcal{H} \neq \mathcal{H}_0$, there is $H \in \mathcal{H}$ such that $H \subset G'$

then we say that $\vec{\mathcal{H}}$ *can be* G_0*–unpulled in* G.

If $\vec{\mathcal{H}}$ can be G_0–cleaved, G_0–unkilled, or G_0–unpulled then we say it can be *III–*G_0*–simplified.*

Recall that if s is a vertex of the G–tree S then $\vec{\mathcal{H}}(s)$ denotes the sequence of conjugacy classes of subgroups of the stabilizer of s represented by the stabilizers of edges incident to s. The sequence is indexed by the oriented edges $\overline{S} = S/G$ that are incident to the image in $\overline{S}$ of s.

Lemma 8.5 *Let S be a G–tree. If S can be G_0–cleaved, G_0–unkilled, or G_0–unpulled then there is a vertex s of S such that $\vec{\mathcal{H}}(s)$ can be G_0–cleaved, G_0–unkilled, or G_0–unpulled in G_s.*

Proof Assume that S can be III–G_0–simplified. Let $X \to S$ be a model with a cube $C = I \times I$ and face $C_0 = \{0\} \times I$ such that III–simplifying C from C_0 produces $X' \to S'$. In particular, G_0 is the stabilizer of C_0. Let s_0 be the image of C_0 in S, let e be the image of C in S, and let H_0 be the stabilizer of e. Choose $s_e \in \mathring{e}$, and set $C_{s_e} = X_{s_e} \cap C$.

The desired splitting of G_{s_0} is obtained by collapsing all edges of the G_{s_0}–tree X_{s_0} that are not in the orbit of C_0. The desired H_0–tree is obtained by collapsing all edges of the G_0–tree X_{s_e} that are not in the orbit of C_{s_e}. Thus, $\vec{\mathcal{H}}(s)$ can be III–G_0–simplified. □

Definition 8.6 Suppose that G is a group and that $\vec{\mathcal{H}}$ is a sequence of conjugacy classes of subgroups of G. Suppose that $G = G' *_{G_0} G''$ or $G = G' *_{G_0}$ and that, for all $\mathcal{H} \in \vec{\mathcal{H}}$, $\mathcal{H}$ is conjugate into G' or G''. Then we say that $\vec{\mathcal{H}}$ can be G_0*–blown up in* G.

The proofs of the Lemmas 8.7 and 8.8 are very similar to that of Lemma 8.5 and are not provided.

Lemma 8.7 *Let G be a group and let S be a G–tree. If S can be G_0–blown up then there is a vertex s of S such that $\vec{\mathcal{H}}(s)$ can be G_0–blown up in G_s.* □

Lemma 8.8 *Let G be a group and let S be a G–tree. If S can be G_0–I–simplified then S has a valence one vertex with stabilizer G_0 whose incident edge has isomorphic stabilizer.* □

Remark 8.9 Recall that a II–simplification has no effect on S.

9 Algorithmic Results

9.1 More labeled graphs

A map of labeled graphs $g\colon \Sigma_1 \to \Sigma_2$ is a *morphism* if

- the induced map between universal covers is a morphism; and
- g is label-preserving, ie, the following diagram commutes.

$$\begin{tikzcd} \Sigma_1 \arrow[rr, "g"] \arrow[dr, "\lambda_{\Sigma_1}"'] & & \Sigma_2 \arrow[dl, "\lambda_{\Sigma_2}"] \\ & R_{\mathcal{B}} & \end{tikzcd}$$

It is *strict* if this induced map is also strict. Stallings introduced labeled graphs into the study of free groups. The next lemma is key.

Lemma 9.1 (Stallings [27]) *An immersion of labeled graphs induces an injection of fundamental groups.* □

An *edge path* in a labeled graph Σ is a strict morphism $I \to \Sigma$ where I is an oriented compact interval. If I is a point then the edge path is *trivial.* A nontrivial edge path may be identified with a sequence of oriented edges $e_0 \cdots e_m$ where, for $1 \leq i \leq m$, $\partial_1 e_{i-1} = \partial_0 e_i$. The product of edge paths σ_1 and σ_2 is denoted $\sigma_1\sigma_2$. An edge path is *closed* if the initial and terminal vertices of I have the same image.

A *loop* in Σ is a strict morphism $S \to \Sigma$ where S is an oriented circle. A loop may be represented by a cyclic sequence of edges of Σ. An oriented edge of Σ is *crossed* by a path or a loop if appears in the edge sequence representing the path or loop. The graph Σ is *tight* if its labeling function $\lambda_\Sigma\colon \Sigma \to R_{\mathcal{B}}$ is an immersion. We record a simple corollary of Lemma 9.1.

Corollary 9.2 *If $g\colon I \to \Sigma$ is a tight non-trivial edge path then the element of $\pi_1(R_{\mathcal{B}})$ represented by $\lambda_\Sigma \circ g$ is non-trivial. Equivalently, if $e_0 \cdots e_m$ represents a tight non-trivial edge path and if the label of e_i is $b_i^{\delta_i}$ ($\delta_i = \pm 1$) then $b_0^{\delta_0} \cdots b_m^{\delta_m}$ is non-trivial in $F_{\mathcal{B}}$.* □

9.2 Operations on graphs

9.2.1 Coring

A *core graph* is a graph such that every edge is crossed by an immersed loop. By Zorn's lemma, every graph Σ has a unique maximal core subgraph, its *core*, denoted $\mathsf{core}(\Sigma)$. A core graph contains no valence 0 or 1 vertices. If Σ is connected and has finite fundamental group, then $\mathsf{core}(\Sigma)$ is finite. The core of a tree is empty. If Σ is labeled, then so is $\mathsf{core}(\Sigma)$. In fact, core is a functor from the category of labeled graphs and immersions to the category of labeled core graphs and immersions. The map $\mathsf{core}(\Sigma) \subset \Sigma$ is natural with respect to this functor. The conjugacy class $[[H]]$ of a subgroup of $F_{\mathcal{B}}$ is uniquely represented by the core $\Sigma(H)$ of the cover $R_{\mathcal{B},H}$ of $R_{\mathcal{B}}$ corresponding to H. The simple proof of the next lemma is left to the reader.

Lemma 9.3 *Let e be an edge of the labeled graph Σ.*

- *Suppose that e does not separate Σ. Then, $e \subset \mathsf{core}(\Sigma)$ if and only if there is an immersed loop crossing e exactly once.*
- *Suppose that e separates Σ. Then, $e \subset \mathsf{core}(\Sigma)$ if and only if there is an immersed loop crossing each of e and e^{-1} exactly once.* □

9.2.2 Folds and tightening

A morphism $g\colon \Sigma_1 \to \Sigma_2$ of graphs is a *fold* if the induced map between universal covers is a fold. A fold induces a surjection on the level of fundamental groups. It is a homotopy equivalence unless the edges that are identified share both initial and terminal vertices [10].

A finite graph Σ may be iteratively folded until it is tight. If Σ is not finite, then the direct limit of the system of finite sequences of folds is well-defined. The result is the *tightening* of Σ and is denoted $\mathsf{tight}(\Sigma)$. Fix a base vertex for Σ (if Σ is non-empty) and let H denote the image $(\lambda_\Sigma)_\#(\pi_1(\Sigma)) \subset \pi_1(R_{\mathcal{B}})$. Then, λ_Σ lifts to $\Sigma \to R_{\mathcal{B},H}$. The graph $\mathsf{tight}(\Sigma)$ may be identified with the image of this lift. In fact, tight is a functor from the category of labeled graphs and strict morphisms to the category of tight labeled graphs and immersions. The quotient map $\Sigma \to \mathsf{tight}(\Sigma)$ is natural with respect to this functor. More generally, if $b \in \mathcal{B}$ and if Σ is a labeled graph, then we define $\mathsf{tight}_b(\Sigma)$ as above except that only edges labeled b or b^{-1} are folded.

9.2.3 Applying an automorphism

If Σ is a labeled graph and $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$, then $\alpha\Sigma$ is the labeled graph obtained by replacing each labeled oriented edge e of Σ by the sequence of labeled oriented edges αe. More precisely, if the oriented edge e has the label b and if $\alpha(b) = w$ where w is a reduced word of length k in $\mathcal{B}$, then αe is obtained from e by subdividing e into k subedges. The i^{th} letter of w has the form c^δ where $c \in \mathcal{B}$, and $\delta = \pm 1$. The i^{th} subedge of e is given the label c and an orientation agreeing with that of e if δ is positive and the opposite orientation otherwise. The operation of applying the automorphism α is a functor from the category of labeled graphs and morphisms to itself. The construction gives a cellular map $\alpha\colon \Sigma \to \alpha\Sigma$ that is well defined up to a homotopy rel vertices and that is natural (but not a morphism).

Lemma 9.4 *If σ is an immersed edge path in $\alpha\Sigma$, then there are an immersed edge path $\hat\sigma$ in Σ represented by $e_0 \cdots e_m$, an initial edge subpath σ_0 of $\alpha(e_0)$, and a terminal edge subpath σ_m of $\alpha(e_m)$ such that*

(1) $\sigma_0 \neq \alpha(e_0)$;

(2) $\sigma_m \neq \alpha(e_m)$; *and*

(3) $\alpha(\hat\sigma)$ *is the immersed edge path* $\sigma_0\sigma\sigma_m$.

Proof We may view $\alpha\Sigma$ as being obtained from Σ by subdividing and relabeling. With this in mind, any immersed edge path σ in $\alpha\Sigma$ gives an immersed path σ' in Σ that may not have endpoints vertices. This path extends uniquely to an immersed edge path $\hat\sigma$ that is minimal with respect to containing σ'. $\square$

9.2.4 Collapsing edges

If e is an edge of the labeled graph Σ, then $g\colon \Sigma \to \Sigma'$ is a collapse of e if the induced map between universal covers is the morphism collapsing a lift of e. In this case, we denote Σ' by $\mathsf{collapse}_e(\Sigma)$. More generally, if $\mathcal{E}$ is a set of edges in Σ then we may collapse each edge in $\mathcal{E}$ to a point and obtain $\mathsf{collapse}_{\mathcal{E}}(\Sigma)$. If $g\colon \Sigma \to \Sigma'$ is a morphism, and if $\mathcal{E}'$ is a set of edges in Σ', then there is an induced morphism $\mathsf{collapse}_{\mathcal{E}'}(g)\colon \mathsf{collapse}_{g^{-1}(\mathcal{E}')}(\Sigma) \to \mathsf{collapse}_{\mathcal{E}'}(\Sigma')$. To each edge e' in $\mathsf{collapse}_{\mathcal{E}}(\Sigma)$, we may associate the unique edge e of Σ such that $\mathsf{collapse}_{\mathcal{E}}(e) = e'$. The proof of the next lemma is left to the reader.

Lemma 9.5 (See [10]) (1) *The quotient map $\Sigma \to \mathsf{collapse}_{\mathcal{E}}(\Sigma)$ induces a surjection of fundamental groups.*

(2) *If* $e' \subset \mathsf{core}(\mathsf{collapse}_{\mathcal{E}}(\Sigma))$ *then* $e \subset \mathsf{core}(\Sigma)$.

Remark 9.6 If $\mathcal{E}$ is the set of edges labeled b, then the operations tight_b and $\mathsf{collapse}_{\mathcal{E}}$ commute.

9.2.5 Operations on tight labeled core graphs

If Σ is a tight labeled core graph and if $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$, then $\alpha_{\#}\Sigma$ is the labeled core graph $\mathsf{core}(\mathsf{tight}(\alpha\Sigma))$ obtained by coring the tightening of $\alpha\Sigma$. If Σ represents $[[H]]$, then $\alpha_{\#}\Sigma$ represents $[[\alpha H]]$.

9.2.6 Sequences

All the above notions extend to sequences of labeled graphs. For example, if $\vec{\Sigma} = \{\Sigma_k\}$ is a sequence of labeled graphs, then a path in $\vec{\Sigma}$ is a path $I \to G_{k_0}$ for some choice k_0 of k, $\alpha_{\#}\vec{\Sigma}$ denotes $\{\alpha_{\#}\Sigma_k\}$, etc.

A sequence $\vec{\mathcal{H}}$ of conjugacy classes of subgroups of $F_{\mathcal{B}}$ is uniquely represented by $\Sigma(\vec{\mathcal{H}})$. In the following definitions, $\vec{\Sigma}$ is a sequence of labeled graphs. For a labeled graph Σ, $|\Sigma|_b$ is the number of oriented edges of Σ with label b. If $\vec{\Sigma} = \{\Sigma_i\}$ is a sequence of labeled graphs then $|\vec{\Sigma}|_b$ is the sum of the $|\Sigma_i|_b$.

9.3 Elementary Whitehead automorphisms

A reference for this section is [18]. An *extended permutation of* $F_{\mathcal{B}}$ is an automorphism of $F_{\mathcal{B}}$ induced by a permutation of $\mathcal{B}^{\pm 1}$. An *elementary Whitehead automorphism* is an automorphism α of $F_{\mathcal{B}}$ that is either an extended permutation or has the following form. There is an element $b \in \mathcal{B}^{\pm 1}$ and a subset A of $\mathcal{B}^{\pm 1} \setminus \{b^{\pm 1}\}$ such that

- if $c \in A \setminus A^{-1}$ then $\alpha(c) = bc$;
- if $c \in A \cap A^{-1}$ then $\alpha(c) = bcb^{-1}$; and
- if $c \notin A \cup A^{-1}$ then $\alpha(c) = c$.

We call b *the distinguished label of* α.

Remark 9.7 Let $\vec{\Sigma}$ be a sequence of labeled graphs and let α be an elementary Whitehead automorphism with distinguished label b. There is a 1–1 correspondence between the set of edges of $\vec{\Sigma}$ not labeled b and the edges of

$\alpha\vec{\Sigma}$ not labeled b. In $\alpha\vec{\Sigma}$, there are *old* and *new* edges labeled b. The terminal vertex of each new edge has valence 2 and the other incident edge is not labeled b. Such a valence 2 vertex is *new*; other vertices are *old*. The subgraph of $\alpha\vec{\Sigma}$ consisting of new edges is a forest each component of which is a cone over a set of new vertices with base an old vertex. All edges of the cone have initial vertex the base. See Figure 15.

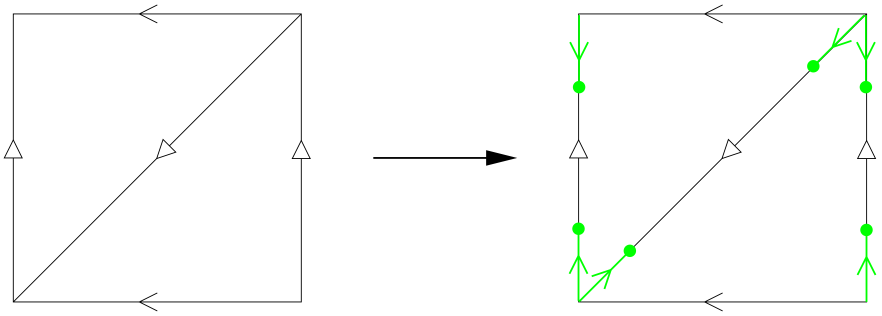

Figure 15: $\alpha(a) = a,\ \alpha(b) = aba^{-1}$

Remark 9.8 For $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$, the sequence of folds needed to tighten $\alpha R_{\mathcal{B}}$ algorithmically gives a factorization of α as a product of elementary Whitehead automorphisms.

The next lemma is a consequence of Step 1 of the proof of the proposition on page 455 of [4].

Lemma 9.9 *Let $g\colon \Sigma_0 \to \Sigma_1$ be a strict morphism of labeled graphs that is surjective on the level of fundamental groups. Then, there is a fold $g'\colon \Sigma_0 \to \Sigma'$ such that g factors as*

$$\Sigma_0 \xrightarrow{g'} \Sigma' \to \Sigma_1. \quad \square$$

Lemma 9.10 *Let $g\colon \Sigma_0 \to \Sigma_1$ be a strict morphism of labeled graphs that is surjective on the level of fundamental groups. Suppose that, for some $b \in \mathcal{B}$, $|\Sigma_1|_b < |\Sigma_0|_b$. Then, there are strict morphisms making the following diagram commute*

$$\begin{tikzcd} I \arrow[r, "\sigma"] \arrow[d, "g'"'] & \Sigma_0 \arrow[d, "g"] \\ T \arrow[r] & \Sigma_1 \end{tikzcd}$$

where

- *where* σ *is an immersed edge path represented by* $e_0 e_1 \cdots e_m$*;*
- T *is a labeled tree;*
- e_0 *and* e_m^{-1} *are labeled by* $b' = b$ *or* b^{-1}*;*
- *for* $0 < i < m$*,* e_i *is not labeled by* b *or* b^{-1}*; and*
- $g'(e_0) = g'(e_m^{-1})$*.*

Proof Since $|\Sigma_1|_b < |\Sigma_0|_b$ there are distinct edges e and e' in Σ_0 each labeled b that are identified under g. Consider the lift $\tilde{g}\colon \tilde{\Sigma}_0 \to \tilde{\Sigma}_1$ to universal covers. Because g induces a surjection on the level of fundamental groups, there are lifts $\tilde{e}$ and $\tilde{e}'$ of e and e' to $\tilde{\Sigma}_0$ that are identified under $\tilde{g}$. Choose $\tilde{e}$ and $\tilde{e}'$ with this property so that the subtree I they span has minimal diameter (with respect to the edge metric). The edge path $\sigma\colon I \to \Sigma_0$ is the restriction of the first covering projection. The edge path σ factors as $I \to T = \tilde{g}(I) \to \Sigma_1$ where the first factor is induced by the restriction of $\tilde{g}$ to I and the second factor is the restriction of the second covering projection. □

Lemma 9.11 *Let* $\vec{\Sigma}$ *be a sequence of tight labeled graphs and let* $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ *be an elementary Whitehead automorphism with distinguished label* b*. Let* $\mathcal{E}$ *(respectively* $\mathcal{E}'$*) be the set of edges of* Σ *(respectively* $\mathsf{tight}(\alpha\Sigma)$*) that are labeled* b*. Then, the following diagram commutes and the lower horizontal arrow is an isomorphism.*

$$\begin{tikzcd}
\vec{\Sigma} \arrow[r] \arrow[d] & \mathsf{tight}(\alpha\vec{\Sigma}) \arrow[d] \\
\mathsf{collapse}_{\mathcal{E}}(\vec{\Sigma}) \arrow[r] & \mathsf{collapse}_{\mathcal{E}'}(\mathsf{tight}(\alpha\vec{\Sigma}))
\end{tikzcd}$$

In particular, there is a natural 1–1 correspondence between edges of $\vec{\Sigma}$ *not labeled* $b^{\pm 1}$ *and edges of* $\mathsf{tight}(\alpha\vec{\Sigma})$ *not labeled* $b^{\pm 1}$*.*

Proof Let $\mathcal{E}''$ be the set of edges of $\alpha\Sigma$ that are labeled b. We have a commuting diagram

$$\begin{tikzcd}
\vec{\Sigma} \arrow[r] \arrow[d] & \alpha\vec{\Sigma} \arrow[r] \arrow[d] & \mathsf{tight}(\alpha\vec{\Sigma}) \arrow[d] \\
\mathsf{collapse}_{\mathcal{E}}(\vec{\Sigma}) \arrow[r] & \mathsf{collapse}_{\mathcal{E}''}(\alpha\vec{\Sigma}) \arrow[r] & \mathsf{collapse}_{\mathcal{E}'}(\mathsf{tight}(\alpha\vec{\Sigma}))
\end{tikzcd}$$

It is clear that the lower left horizontal arrow is an isomorphism and that the lower right arrow is strict and surjective.

In order to obtain a contradiction, assume

$$\mathsf{collapse}_{\mathcal{E}''}(\alpha\vec{\Sigma}) \to \mathsf{collapse}_{\mathcal{E}'}(\mathsf{tight}(\alpha\vec{\Sigma}))$$

is not injective. Since this map is π_1–surjective, by Lemma 9.9 there are two edges not labeled b or b^{-1} with the same image. It follows that there are two edges not labeled b or b^{-1} with the same image under $\alpha\vec{\Sigma} \to \mathsf{tight}(\alpha\vec{\Sigma})$. Using Lemma 9.10 and taking a subpath if necessary, there is an immersed edge path $\sigma\colon I \to \alpha\vec{\Sigma}$ represented by $e_0 e_1 \cdots e_m$ such that

- the label of e_0 and e_m^{-1} is $b' \neq b^{\pm 1}$;
- the label of e_i is $b^{\pm 1}$ for all $0 < i < m$; and
- $I \to \alpha\vec{\Sigma} \to \mathsf{tight}(\alpha\vec{\Sigma})$ factors through a tree.

If $\hat{\sigma}$ is the immersed edge path in $\vec{\Sigma}$ determined by σ as in Lemma 9.4 then

- the label of $\hat{e}_0$ and $\hat{e}_{\hat{m}}^{-1}$ is b'; and
- $\hat{e}_i$ is labeled $b^{\pm 1}$ for $1 < i < \hat{m}$.

Since $\hat{\sigma}$ is an immersion, all the $\hat{e}_i$, $1 < i < \hat{m}$, are consistently oriented. It is easy to see then that $I \to \alpha\vec{\Sigma} \to \mathsf{tight}(\alpha\vec{\Sigma})$ cannot factor through a tree, contradiction. □

Lemma 9.12 *Let $\vec{\Sigma}$ be a sequence of tight labeled core graphs and let α be an elementary Whitehead automorphism with distinguished label b, let e be an edge of $\vec{\Sigma}$ not labeled $b^{\pm 1}$, and let e' be the corresponding edge in $\mathsf{tight}(\alpha\vec{\Sigma})$. Then, e' is in $\alpha_{\#}\vec{\Sigma} = \mathsf{core}(\mathsf{tight}(\alpha\vec{\Sigma}))$. In particular, there is a natural 1–1 correspondence between edges of $\vec{\Sigma}$ not labeled $b^{\pm 1}$ and edges of $\alpha_{\#}\vec{\Sigma}$ not labeled $b^{\pm 1}$.*

Proof Suppose that e separates (respectively does not separate) its component. By Lemma 9.3, there is an immersed loop $g\colon S \to \vec{\Sigma}$ crossing e (respectively crossing e' and e'^{-1} each) exactly once. It follows from Lemma 9.11 that the immersed loop $\mathsf{tight}(\alpha g)$ crosses e' (respectively e' and e'^{-1} each) exactly once. By Lemma 9.3, e' is contained in $\alpha_{\#}\vec{\Sigma}$. □

9.4 Complexity

If $\vec{\mathcal{H}}$ is a finite sequence of conjugacy classes of finitely generated subgroups of $F_{\mathcal{B}}$ and if $b \in \mathcal{B}$, then $|\vec{\mathcal{H}}|_b$ is the number of edges in $\Sigma(\vec{\mathcal{H}})$ that are labeled with b. The *complexity of* $\vec{\mathcal{H}}$, denoted $\mathsf{c}(\vec{\mathcal{H}})$, is the number of edges in $\Sigma(\vec{\mathcal{H}})$ or equivalently $\sum_{b\in\mathcal{B}} |\vec{\mathcal{H}}|_b$.

We will also need a finer measure of complexity of $\vec{\mathcal{H}}$. Define the *lexity* of $\vec{\mathcal{H}}$, denoted $\mathsf{lex}(\vec{\mathcal{H}})$, to be the sequence of non-negative integers $\{|\Sigma(\vec{\mathcal{H}})|_b\}_{b\in\mathcal{B}}$ arranged in non-decreasing order. The set $\mathcal{L}$ of non-decreasing sequences of non-negative integers is well-ordered lexicographically. Let $\mathsf{minlex}(\vec{\mathcal{H}})$ denote $\min_{b\in\mathcal{B}}\{|\Sigma(\vec{\mathcal{H}})|_b\}$.

Lemma 9.13 *Let $\vec{\mathcal{H}}$ be a finite sequence of conjugacy classes of finitely generated subgroups of $F_{\mathcal{B}}$ and let α be an elementary Whitehead automorphism. Then $\mathsf{c}(\alpha\vec{\mathcal{H}}) < \mathsf{c}(\vec{\mathcal{H}})$ if and only if $\mathsf{lex}(\alpha\vec{\mathcal{H}}) < \mathsf{lex}(\vec{\mathcal{H}})$. If further α has distinguished label b, then $\mathsf{lex}(\alpha\vec{\mathcal{H}}) \le \mathsf{lex}(\vec{\mathcal{H}})$ if and only if $|\alpha\vec{\mathcal{H}}|_b \le |\vec{\mathcal{H}}|_b$ with equality if and only if $|\alpha\vec{\mathcal{H}}|_b = |\vec{\mathcal{H}}|_b$.*

Proof Since extended permutations preserve both c and lex, we may suppose that α has distinguished label $b \in \mathcal{B}$. It follows from Lemma 9.12 that if $b' \neq b^{\pm 1}$, then the number of times that b' appears in $\Sigma(\alpha\vec{\mathcal{H}})$ is the same as the number of times that b' appears in $\Sigma(\vec{\mathcal{H}})$. □

9.5 Gersten's Theorem

Let $\vec{\mathcal{H}}$ be a finite sequence of conjugacy classes of finitely generated subgroups of $F_{\mathcal{B}}$. If $\mathsf{c}(\vec{\mathcal{H}}) = \min\{\mathsf{c}(\alpha\vec{\mathcal{H}}) \mid \alpha \in \mathsf{Aut}(F_{\mathcal{B}})\}$ then $\vec{\mathcal{H}}$ is a *Gersten representative for the orbit* $\mathsf{Aut}(F_{\mathcal{B}})\vec{\mathcal{H}}$. We also write that $\vec{\mathcal{H}}$ is a Gersten representative for any element of the orbit. Since $\vec{\mathcal{H}}$ is an element of this orbit, we often simply write that $\vec{\mathcal{H}}$ is a Gersten representative. A finite set of generators for a representative $H \in \mathcal{H}$ for each $\mathcal{H} \in \vec{\mathcal{H}}$ is a *finite generating system* for $\vec{\mathcal{H}}$. SM Gersten [14][18] gave an algorithm that when input a finite generating system for $\vec{\mathcal{H}}$ outputs the finite set of Gersten representatives for $\vec{\mathcal{H}}$.

Theorem 9.14 [14],[18]

(1) *If $\vec{\mathcal{H}}$ is a finite sequence of conjugacy classes of finitely generated subgroups of $F_{\mathcal{B}}$ that is not a Gersten representative, then there is an elementary Whitehead automorphism α such that $\mathsf{c}(\alpha\vec{\mathcal{H}}) < \mathsf{c}(\vec{\mathcal{H}})$.*

(2) *If $\vec{\mathcal{H}}$ and $\vec{\mathcal{H}}'$ are Gersten representatives for $\vec{\mathcal{H}}$, then there is a finite sequence $\{\alpha_k\}_{k=1}^m$ of elementary Whitehead automorphisms and a sequence*

$$\{\vec{\mathcal{H}} = \vec{\mathcal{H}}_0, \vec{\mathcal{H}}_1, \cdots, \vec{\mathcal{H}}_m = \vec{\mathcal{H}}'\}$$

of Gersten representatives such that $\alpha_k\vec{\mathcal{H}}_{k-1} = \vec{\mathcal{H}}_k$ for $1 \le k \le m$.

Corollary 9.15 *If $\vec{\mathcal{H}}$ is a finite sequence of conjugacy classes of finitely generated subgroups of $F_{\mathcal{B}}$, then there is an algorithm that when input a finite generating system for $\vec{\mathcal{H}}$ outputs a Gersten representative for $\vec{\mathcal{H}}$.*

Proof Iteratively apply elementary Whitehead automorphisms to $\vec{\mathcal{H}}$ until complexity cannot be decreased. The resulting sequence is a Gersten representative by Theorem 9.14(1). □

Corollary 9.16 *Let $\vec{\mathcal{H}}$ be a finite sequence of conjugacy classes of finitely generated subgroups of $F_{\mathcal{B}}$. Then, there is an algorithm that when input a finite generating system for $\vec{\mathcal{H}}$ outputs the finitely many Gersten representatives of $\vec{\mathcal{H}}$.*

Proof By Corollary 9.15, we may assume that $\vec{\mathcal{H}}$ is a Gersten representative. Consider the graph whose vertices are finite sequences of conjugacy classes of finitely generated subgroups of $F_{\mathcal{B}}$ of complexity equal to $\mathsf{c}(\vec{\mathcal{H}})$, and where two vertices $\vec{\mathcal{H}}_1$ and $\vec{\mathcal{H}}_2$ are connected by an edge if there is an elementary Whitehead automorphism α such that $\alpha\vec{\mathcal{H}}_1 = \vec{\mathcal{H}}_2$. By Theorem 9.14(2), the component of this graph containing $\vec{\mathcal{H}}$ has vertices that are precisely the Gersten representatives of $\vec{\mathcal{H}}$. □

9.6 Consequences of Lemma 9.13

In this section we show that simplifications can be detected using Gersten representatives. Throughout this section, $\vec{\mathcal{H}}$ is a finite sequence of conjugacy classes of finitely generated subgroups of $F_{\mathcal{B}}$.

Lemma 9.17 *The following are equivalent.*

(1) *There is an $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ such that $\mathsf{minlex}(\alpha\vec{\mathcal{H}}) = 0$.*

(2) *For some (any) Gersten representative $\vec{\mathcal{H}}'$ of $\vec{\mathcal{H}}$, $\mathsf{minlex}(\vec{\mathcal{H}}') = 0$.*

Proof By Lemma 9.13, for any $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ and any Gersten representative $\vec{\mathcal{H}}'$ of $\vec{\mathcal{H}}$, $\mathsf{lex}(\alpha\vec{\mathcal{H}}) \geq \mathsf{lex}(\vec{\mathcal{H}}')$. Thus, (1) $\iff$ (2). □

Lemma 9.18 *Let $\mathcal{B}' \subset \mathcal{B}$ and suppose that $\vec{\mathcal{H}}'$ is a finite sequence of conjugacy classes of finitely generated subgroups of $F_{\mathcal{B}'}$. The following are equivalent.*

(1) *$\vec{\mathcal{H}}'$ is a Gersten representative with respect to $\mathcal{B}'$.*

(2) *$\vec{\mathcal{H}}'$ is a Gersten representative with respect to $\mathcal{B}$.*

Proof That (2) $\implies$ (1) is clear. Suppose (1), but not (2). By Lemma 9.13 there is then an elementary Whitehead automorphism $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ with distinguished label $b \notin (\mathcal{B}')^{\pm 1}$ such that $|\alpha\vec{\mathcal{H}}'|_b < |\vec{\mathcal{H}}'|_b = 0$, contradiction. $\square$

Recall that the terms *visibly blown up*, *visibly unpulled*, *visibly unkilled*, *visibly uncleaved*, and *visibly simplified* were defined in Definition 2.5.

Lemma 9.19 *The following are equivalent.*

(1) *There is an $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ such that $\Sigma(\alpha\vec{\mathcal{H}})$ can be visibly blown up.*

(2) *For some (every) Gersten representative $\vec{\mathcal{H}}'$ of $\vec{\mathcal{H}}$, $\Sigma(\vec{\mathcal{H}}')$ can be visibly blown up.*

Proof The lemma will follow from:

Claim If $\Sigma(\vec{\mathcal{H}})$ can be visibly blown up and if α is an elementary Whitehead automorphism with $\mathsf{lex}(\alpha\vec{\mathcal{H}}) \leq \mathsf{lex}(\vec{\mathcal{H}})$ then either $\Sigma(\alpha\vec{\mathcal{H}})$ can be visibly blown up or there is an automorphism α' with $\mathsf{lex}(\alpha'\vec{\mathcal{H}}) < \mathsf{lex}(\vec{\mathcal{H}})$ such that $\Sigma(\alpha'\vec{\mathcal{H}})$ can be visibly blown up.

We now prove the claim. Suppose $\mathcal{B} = \mathcal{B}' \sqcup \mathcal{B}''$ is a non-trivial partition such that each element of $\vec{\mathcal{H}}$ has a representative in either $F_{\mathcal{B}'}$ or $F_{\mathcal{B}''}$. If α is an extended permutation, then the claim is clear.

Alternatively, let $b \in \mathcal{B}^{\pm 1}$ be the distinguished label of α and suppose without loss that $b \in (\mathcal{B}')^{\pm 1}$. Let α' be the automorphism that agrees with α on $\mathcal{B}'$ and that is the identity on $\mathcal{B}''$. It is clear that $\Sigma(\alpha'\vec{\mathcal{H}})$ can be visibly blown up. Now, $\Sigma(\alpha\vec{\mathcal{H}})$ can be visibly blown up if and only if, for each $\mathcal{H} \in \vec{\mathcal{H}}$ with a representative in $F_{\mathcal{B}''}$, $\alpha\mathcal{H}$ has a representative in $F_{\mathcal{B}''}$. By Lemma 9.13, this occurs if and only if $|\alpha\mathcal{H}|_b > 0$ for such $\mathcal{H}$ and this occurs if and only if $|\alpha'\vec{\mathcal{H}}|_b < |\alpha\vec{\mathcal{H}}|_b$. $\square$

Lemma 9.20 *Suppose that for some (any) Gersten representative $\vec{\mathcal{H}}'$ of $\vec{\mathcal{H}}$ we have $\mathsf{minlex}(\vec{\mathcal{H}}') \neq 0$. Then, the following are equivalent.*

(1) *There is an $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ such that $\mathsf{minlex}(\alpha\vec{\mathcal{H}}) = 1$.*

(2) *For some (any) Gersten representative $\vec{\mathcal{H}}'$ of $\vec{\mathcal{H}}$, $\mathsf{minlex}(\vec{\mathcal{H}}') = 1$.*

(3) *$\vec{\mathcal{H}}$ can be either visibly unkilled or visibly unpulled.*

Proof In the presence of $\mathsf{minlex}(\vec{\mathcal{H}}') \neq 0$, (1) $\iff$ (2) by Lemma 9.13. Suppose that (2) holds. There are two cases: there is a label $b \in \mathcal{B}$ that appears exactly once in $\Sigma(\vec{\mathcal{H}}')$ and (a) the edge labeled b separates its component and (b) edge labeled b does not separate its component. It is an easy exercise to show that in case (a) $\Sigma(\vec{\mathcal{H}})$ can be visibly unkilled and in case (b) $\Sigma(\vec{\mathcal{H}})$ can be visibly unpulled.

Suppose that (3) holds. Then there is free factorization $F_{\mathcal{B}} = F^1 *_{\langle 1\rangle} = F^1 * \langle t\rangle$ as in Definition 8.3 or Definition 8.4. Let $\mathcal{B}^1$ be a basis for F^1. Choose an $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ so that $\alpha(\mathcal{B}^1 \sqcup \{t\}) = \mathcal{B}$. Then $\mathsf{minlex}(\alpha\vec{\mathcal{H}}) = 1$. □

The next lemma will be used to prove Lemma 9.23.

Lemma 9.21 *Suppose that*

(1) $\mathsf{minlex}(\vec{\mathcal{H}}) > 1$;

(2) $\mathcal{B} = \mathcal{B}' \sqcup \mathcal{B}''$ *is a non-trivial partition;*

(3) $\Sigma(\vec{\mathcal{H}}) = \Sigma(\vec{\mathcal{H}}') \vee \Sigma(\vec{\mathcal{H}}'')$ *where* $\vec{\mathcal{H}}'$ *is a Gersten representative in* $F_{\mathcal{B}'}$ *and* $\vec{\mathcal{H}}''$ *is a Gersten representative in* $F_{\mathcal{B}''}$*; and*

(4) $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ *is an elementary Whitehead automorphism with distinguished label* $b \in (\mathcal{B}')^{\pm 1}$ *such that* $\mathsf{lex}(\alpha\vec{\mathcal{H}}) \leq \mathsf{lex}(\vec{\mathcal{H}})$.

Then,

- $\alpha\vec{\mathcal{H}}'' = \vec{\mathcal{H}}''$;
- $\alpha\vec{\mathcal{H}}'$ *is a Gersten representative for* $\vec{\mathcal{H}}'$*; and*
- $\Sigma(\alpha\vec{\mathcal{H}}) = \Sigma(\alpha\vec{\mathcal{H}}') \vee \Sigma(\vec{\mathcal{H}}'')$.

In particular, $\vec{\mathcal{H}}$ *satisfying (1), (2), and (3) is a Gersten representative.*

Proof Note:

- $\Sigma(\alpha\vec{\mathcal{H}}) = \mathsf{tight}\left(\Sigma(\alpha\vec{\mathcal{H}}') \vee I \vee \Sigma(\alpha\vec{\mathcal{H}}'')\right)$ for some labeled graph I homeomorphic to a compact interval. This follows because in tightening $\alpha\Sigma(\vec{\mathcal{H}})$, we can tighten $\alpha\Sigma(\vec{\mathcal{H}}')$ and $\alpha\Sigma(\vec{\mathcal{H}}'')$ first.
- The subgraph of $\Sigma(\alpha\vec{\mathcal{H}}'')$ consisting of edges labeled b is a tree whose components are single (non-loop) edges. This follows from Remark 9.7 and the following commutative diagram where $\mathcal{E}$ is the set of edges of $\Sigma(\alpha\vec{\mathcal{H}}'')$ that are labeled b.

$$\begin{tikzcd}
\Sigma(\vec{\mathcal{H}}'') \arrow[rr] \arrow[dr, "="'] & & \Sigma(\alpha\vec{\mathcal{H}}'') \arrow[dl, "\mathsf{collapse}_{\mathcal{E}}"] \\
& \Sigma(\vec{\mathcal{H}}'') &
\end{tikzcd}$$

It follows that in tightening $\Sigma(\alpha\vec{\mathcal{H}}') \vee I \vee \Sigma(\alpha\vec{\mathcal{H}}'')$ at most one edge of $\Sigma(\alpha\vec{\mathcal{H}}')$ folds with an edge of $\Sigma(\alpha\vec{\mathcal{H}}'')$. Hence,

$$\begin{aligned}\mathsf{c}(\vec{\mathcal{H}}') + \mathsf{c}(\vec{\mathcal{H}}'') = \mathsf{c}(\vec{\mathcal{H}}) &\geq \mathsf{c}(\alpha\vec{\mathcal{H}}) \\ &\geq \mathsf{c}(\alpha\vec{\mathcal{H}}') + \mathsf{c}(\alpha\vec{\mathcal{H}}'') - 1 \geq \mathsf{c}(\vec{\mathcal{H}}') + \mathsf{c}(\vec{\mathcal{H}}'') - 1.\end{aligned}$$

Thus, either $\mathsf{c}(\alpha\vec{\mathcal{H}}'') = \mathsf{c}(\vec{\mathcal{H}}'')$ or $\mathsf{c}(\alpha\vec{\mathcal{H}}'') = \mathsf{c}(\vec{\mathcal{H}}'') + 1$. In the former case, we are done. The latter case cannot occur. Indeed, otherwise $|\alpha\vec{\mathcal{H}}''|_b = 1$. But then, by Lemmas 9.18 and 9.20 , $\mathsf{minlex}(\vec{\mathcal{H}}'') \leq 1$ and hence $\mathsf{minlex}(\vec{\mathcal{H}}) \leq 1$, contradiction. □

Remark 9.22 Without (1), Lemma 9.21 is false. Consider $\mathcal{B} = \{a, b\} \sqcup \{c\}$, $H = \langle ab^{-1}, c\rangle$, and $\alpha(a) = a$, $\alpha(b) = b$, $\alpha(c) = bc$.

Lemma 9.23 *Suppose that $\mathsf{minlex}(\vec{\mathcal{H}}') > 1$ for some (any) Gersten representative $\vec{\mathcal{H}}'$ of $\vec{\mathcal{H}}$. Then, the following are equivalent.*

(1) *There is $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ such that $\Sigma(\alpha\vec{\mathcal{H}})$ can be visibly cleaved.*

(2) *For some (every) Gersten representative $\vec{\mathcal{H}}'$ of $\vec{\mathcal{H}}$, $\Sigma(\vec{\mathcal{H}}')$ can be visibly cleaved.*

Proof (2) $\Longrightarrow$ (1) is clear. We now show (1) $\Longrightarrow$ (2). Suppose that $\Sigma(\alpha\vec{\mathcal{H}}) = \Sigma(\vec{\mathcal{H}}') \vee \Sigma(\vec{\mathcal{H}}'')$ where all labels of $\Sigma(\vec{\mathcal{H}}')$ are in $\mathcal{B}'^{\pm 1}$ and all labels of $\Sigma(\vec{\mathcal{H}}'')$ are in $\mathcal{B}''^{\pm 1}$ for some non-trivial partition $\mathcal{B} = \mathcal{B}' \sqcup \mathcal{B}''$. Choose $\alpha' \in \mathsf{Aut}(F_{\mathcal{B}'})$ and $\alpha'' \in \mathsf{Aut}(F_{\mathcal{B}''})$ such that $\alpha'\vec{\mathcal{H}}'$ and $\alpha''\vec{\mathcal{H}}''$ are Gersten representatives. Let $\alpha \in \mathsf{Aut}(F_{\mathcal{B}})$ agree with α' on $F_{\mathcal{B}'}$ and α'' on $F_{\mathcal{B}''}$. Then, $\Sigma(\alpha\vec{\mathcal{H}}) = \Sigma(\alpha'\vec{\mathcal{H}}') \vee I' \vee I'' \vee \Sigma(\alpha''\vec{\mathcal{H}}'')$ where I' and I'' are labeled graphs homeomorphic to compact intervals, all labels of I' are in $(\mathcal{B}')^{\pm 1}$, and all labels of I'' are in $(\mathcal{B}'')^{\pm 1}$. We will now show that $\Sigma(\alpha'\vec{\mathcal{H}}') \vee \Sigma(\alpha''\vec{\mathcal{H}}'')$ is also a representative for $\vec{\mathcal{H}}$. Suppose that I'' is not trivial and that the edge e of I'' with initial vertex in $\Sigma(\alpha'\vec{\mathcal{H}}') \vee I'$ is labeled b. Let α_b be the elementary Whitehead automorphism that is conjugation by b on $\mathcal{B}'$ and the identity on $\mathcal{B}''$. Then, $\Sigma(\alpha_b\alpha\vec{\mathcal{H}})$ is obtained from $\Sigma(\alpha'\vec{\mathcal{H}}') \vee I' \vee I'' \vee \Sigma(\alpha''\vec{\mathcal{H}}'')$ by collapsing e. We may continue until I'' and symmetrically I' are trivial. It follows from Lemma 9.21 that the result is a Gersten representative and that all Gersten representatives have this form. □

10 Proof of the Main Theorem

Proposition 10.1 *The algorithm of Section 2.10 is in fact an algorithm.*

Proof To detect a reduction, it is only necessary to be able to decide algorithmically if a homomorphism $\alpha\colon F_{\mathcal{B}} \to F_{\mathcal{B}'}$ between free groups is an isomorphism. According to Stallings [27], α is injective if and only if the rank of the Stallings representative of $\alpha F(\mathcal{B})$ with respect to $\mathcal{B}'$ is $|\mathcal{B}|$. It is surjective if and only if this Stallings representative is $R_{\mathcal{B}'}$. Thus Step 1 is algorithmic.

Step 2 only depends on being able to find a Gersten representative and this is algorithmic by Corollary 9.15.

After reducing and $\mathbf{1}$–simplifying, complexity has been reduced where complexity is the sequence of ranks of conjugacy classes of edge stabilizers viewed as an element of $\mathcal{L}$. Therefore this process stops. □

We are finally in a position to prove Theorem 2.8.

Proof of Theorem 2.8 By Proposition 10.1, we may assume that $\mathcal{G}$ is reduced. Let S be the corresponding G–tree. We may also assume that no edge stabilizer is trivial of S. (Otherwise, G has an obvious free decomposition and we may work with the factors instead of G.) If G is freely decomposable, then by Lemma 8.5, there is a vertex $v \in V$ such that $\vec{\mathcal{H}}(v)$ can be $\mathbf{1}$–simplified. In particular, by definition there is a basis $\mathcal{B}$ for G_v with respect to which $\Sigma(\vec{\mathcal{H}}(v))$ may be visibly simplified. Since $\mathsf{Aut}(F_{\mathcal{B}})$ acts transitively on bases, $\Sigma_{\mathsf{G}}(\vec{\mathcal{H}}(v))$ may be visibly simplified by Lemmas 9.19, 9.20, and 9.23. □

We end with a few questions.

Question 1 Is there an algorithm to decide if the fundamental group of a finite graph of finite rank free groups is a surface group?

Question 2 Is there an algorithm to decide if the fundamental group of a finite graph of finite rank free groups splits over $\mathbb{Z}$?

Question 3 Is there an algorithm to find the JSJ–decomposition of the fundamental group of a finite graph of finite rank free groups?

One can't hope to go too far in this direction since according to C. Miller [22] the isomorphism problem for finite graphs of finite rank free groups is unsolvable.

Appendix

A Bookkeeping

In this section, details are provided as to how we record a finite graph of finite rank free groups $\mathcal{G}$ and how that data changes under a simplification. See Section 2.1 for notation.

A finite graph of finite rank free groups $\mathcal{G}$ is given by the following data:

- for each $e \in \hat{E}$, a basis $\mathcal{B}_e$ for G_e such that $\mathcal{B}_e = \mathcal{B}_{e^{-1}}$;
- for $v \in V$, a basis $\mathcal{B}_v$ for G_v; and
- $\vec{w} = \{\vec{w}_e\}_{e\in\hat{E}}$ where $\vec{w}_e = \{w_{e,b}\}_{b\in\mathcal{B}_e}$ is a sequence of reduced words in $\mathcal{B}_v^{\pm 1}$ representing $\{\varphi_e(b)\}_{b\in\mathcal{B}_e}$.

The Stallings algorithm referred to in Section 2.9 can be used to decide if a sequence $\vec{w}_e$ determines a monomorphism. Indeed, check if the rank of Σ_{S} obtained from $\vec{w}_e$ is equal to the rank of G_e.

Definition A.1 If $\{\vec{w}_e\}_{e\in\hat{E}}$ is a sequence as above, then we say another sequence $\{\vec{w}_e^{out}\}_{e\in\hat{E}}$ is *conjugate* to $\{\vec{w}_e\}_{e\in\hat{E}}$, written $\{\vec{w}_e^{out}\}_{e\in\hat{E}} \sim \{\vec{w}_e\}_{e\in\hat{E}}$, if there are $\{\psi_e = \psi_{e^{-1}} \in \mathsf{Aut}(G_e)\}_{e\in\hat{E}}$, $\{\psi_v \in \mathsf{Aut}(G_v)\}_{v\in V}$, and $\vec{h} = \{h_e \in G_v\}_{e\in\hat{E}}$ so that $\vec{w}_e^{out}$ represents $\{\psi_v \circ i_{h_e} \circ \varphi_e \circ \psi_e(b)\}_{b\in\mathcal{B}_e}$. If $\{\psi_e\}$ and $\{\psi_v\}$ are viewed as changing bases, then we see that $\{\vec{w}_e\}_{e\in\hat{E}}$ and $\{\vec{w}_e^{out}\}_{e\in\hat{E}}$ determine conjugate graphs of groups.

Definition A.2 For $e \in \hat{E}$ and $v = \partial_0 e$, we say that given bases $\mathcal{B}_e$ and $\mathcal{B}_v$ are *good* if they determine decompositions of G_e and G_v that give a visual simplification. Specifically, we say that $\mathcal{B}_v$ and $\mathcal{B}_e$ are *good* in any of the following four cases.

blowing up There is a distinguished element $b_v \in \mathcal{B}_v$ so that $\varphi_f(\mathcal{B}_f) \subset \langle\hat{b}_v\rangle$, $f \in \hat{E}(v)$. There is no condition on $\mathcal{B}_e$ in this case. (We use the notation $\hat{b}_v = \mathcal{B}_v \setminus \{b_v\}$.)

unpulling There are distinguished elements $b_v \in \mathcal{B}_v$ and $b_e \in \mathcal{B}_e$ so that $\varphi_e(\hat{b}_e) \subset \langle\hat{b}_v\rangle$, $\varphi_e(b_e) = b_v$, and $\varphi_f(G_f) \subset \langle\hat{b}_v\rangle$, $f \in \hat{E}(v) \setminus \{e\}$.

unkilling There is a distinguished element $b_v \in \mathcal{B}_v$ and a partition $\mathcal{B}_e = \mathcal{B}_e' \sqcup \mathcal{B}_e''$ such that $\varphi_e(\mathcal{B}_e') \subset \langle\hat{b}_v\rangle$, $\varphi(\mathcal{B}_e'') \subset b_v\langle\hat{b}_v\rangle b_v^{-1}$, and $\varphi_f(G_f) \subset \langle\hat{b}_v\rangle$, $f \in \hat{E}(v) \setminus \{e\}$.

cleaving There are partitions $\mathcal{B}_v = \mathcal{B}'_v \sqcup \mathcal{B}''_v$ and $\mathcal{B}_e = \mathcal{B}'_e \sqcup \mathcal{B}''_e$ such that $\varphi_e(\mathcal{B}'_e) \subset \langle \mathcal{B}'_v \rangle$, $\varphi(\mathcal{B}''_e) \subset \langle \mathcal{B}''_v \rangle$, and for $f \in \hat{E}(v) \setminus \{e\}$ either $\varphi_f(G_f) \subset \langle \mathcal{B}'_v \rangle$ or $\varphi_f(G_f) \subset \langle \mathcal{B}''_v \rangle$.

If $\mathcal{B}_e$ and $\mathcal{B}_v$ are good bases, then the corresponding simplification can be performed. The bases associated to edges and vertices of the simplified graphs are as follows. Unless explicitly mentioned, the words $w_{e,b}$ representing $\varphi_e(b)$ do not change. We use the notation as above and in Sections 2.3–2.6.

blowing up After blowing up, $\mathcal{B}_v = \hat{b}_v$.

unpulling After unpulling, $\mathcal{B}_e = \hat{b}_e$ and $\mathcal{B}_v = \hat{b}_v$.

unkilling After unkilling, $\mathcal{B}_{e'} = \mathcal{B}'_e$, $\mathcal{B}_{e''} = \mathcal{B}''_e$, $\mathcal{B}_v = \hat{b}_v$, and $w_{e'',b}$ represents $b_v^{-1}\varphi_e(b)b_v$ for $b \in \mathcal{B}''_v$.

cleaving After cleaving, $\mathcal{B}_{v'} = \mathcal{B}'_v$, $\mathcal{B}_{v''} = \mathcal{B}''_v$, $\mathcal{B}_{e'} = \mathcal{B}'_e$, and $\mathcal{B}_{e''} = \mathcal{B}''_e$.

Proposition A.3 *Let $\mathcal{G}$ be a finite graph of finite rank free groups given as in the beginning of this section. Suppose that, for some $v \in V$, $\Sigma_{\mathsf{G}}(\vec{\mathcal{H}}(v))$ can be visibly simplified. Then,*

(1) *there is $\mathcal{G}^{out} \sim \mathcal{G}$ such that $\mathcal{G}^{out}$ can be simplified; and*

(2) *If $\vec{w}$ specifies the bonding maps of $\mathcal{G}$, then a conjugate sequence $\vec{w}^{out}$ may be found algorithmically so that in $\mathcal{G}^{out}$ bases are good.*

Proof Since (2) implies (1), it is enough to prove (2). The conjugate sequence $\vec{w}^{out}$ will be specified by supplying change of basis automorphisms $\{\psi_f = \psi_{f^{-1}} \in \mathsf{Aut}(G_f)\}_{f\in\hat{E}}$ and $\{\psi_u \in \mathsf{Aut}(G_u)\}_{u\in V}$ as well as the conjugating elements $\vec{h} = \{h_f \in G_f\}_{f\in\hat{E}}$ as in Definition A.1. The change of basis automorphisms can be given by specifying new bases.

Gersten's algorithm supplies $\alpha \in \mathsf{Aut}(G_v)$ so that $\mathsf{core}(\Sigma_{\mathsf{S}}(\alpha\vec{H}(v))) = \Sigma_{\mathsf{G}}(\vec{\mathcal{H}}(v))$. By taking $\psi_v = \alpha$, we obtain a conjugate sequence (still denoted $\vec{w}$) such that $\mathsf{core}(\Sigma)$ can be visibly simplified where $\Sigma = \Sigma_{\mathsf{S}}(\vec{H}(v))$. We are using the convention that all unmentioned change of basis and conjugating automorphisms are identities.

Recall that each component of Σ has a basepoint. Choose shortest paths in each component of Σ from the basepoint to the core and take as conjugating elements the words read off along the inverses of these paths. The resulting conjugate sequence has basepoints in $\mathsf{core}(\Sigma)$. Hence, we may further assume that $\Sigma = \mathsf{core}(\Sigma)$.

If Σ can be visibly blown up, then bases are good and if $b \in \mathcal{B}_v$ is an element that does not appear as a label on Σ we take b for the distinguished element b_v.

In each of the remaining cases, there is a distinguished $e \in \hat{E}(v)$. Let Σ_e be the component of Σ indexed by e, ie, $\Sigma_e = \Sigma_{\mathsf{S}}(\varphi_e(G_e))$. There is the natural factorization $G_e \to \pi_1(\Sigma_e, *) \to G_v$ of φ_e where the first map is the isomorphism coming from Stallings algorithm (see Section 2.9) and the second map is induced by natural map of Section 2.8. We will use this first map to identify G_e with $\pi_1(\Sigma_e, *)$. If a basis for $\pi_1(\Sigma_e, *)$ is given via a maximal tree T for Σ_e, it is easy to write this isomorphism in terms of the given bases. To make this identification explicit, it is necessary to invert the automorphism and this can be done algorithmically, see Remark 9.8. Via our identification, T determines a new basis for G_e and hence a conjugate sequence.

If Σ can be visibly unpulled, then choose an edge c of Σ_e that does not separate its component and whose label $b \in \mathcal{B}_v$ appears exactly once in Σ. Choose a maximal tree T for Σ_e so that $c \not\subset T$ and let $\mathcal{B}_e^{out}$ be the basis determined by T. There is an element $b_e \in \mathcal{B}_e^{out}$ corresponding to c. Let b_v be φ_e–image in G_v of b_e. Set $\mathcal{B}_v^{out} = \{b_v\} \sqcup \hat{b}$ which is a basis for G_v since b appears exactly once in b_v when expressed as a $\mathcal{B}_v$–word. The new bases are good and determine change of basis automorphisms giving rise to the desired $\vec{w}^{out}$.

If Σ can be visibly unkilled, then choose an edge c of Σ_e that separates its component and whose label $b \in \mathcal{B}_v$ appears exactly once in Σ. By changing the orientation of c and inverting b if necessary, we may assume that $\partial_0 c$ and the basepoint are in the same component of the graph obtained by removing c from Σ_e. Let T be a maximal tree for Σ_e. This gives the desired new basis $\mathcal{B}_e^{out}$. Choosing a shortest path in T from the $\partial_0 c$ to the basepoint gives rise to a conjugating element h_e which has the effect of changing the basepoint of Σ_e to $\partial_0 c$. The new bases are good. Indeed, the partition of $\mathcal{B}_e^{out}$ is induced by the separating edge. More precisely, let $(\mathcal{B}_e^{out})''$ be the elements of $\mathcal{B}_e^{out}$ that contain the letter b and let $(\mathcal{B}_e^{out})'$ be the complement. Set $b_v = b$.

If Σ can be visibly cleaved, then write $\Sigma = \Sigma' \vee \Sigma''$ respecting the non-trivial partition $\mathcal{B}_v' \sqcup \mathcal{B}_v''$ of $\mathcal{B}_v$. A choice of maximal tree T for Σ_e gives the desired new basis $\mathcal{B}_e^{out}$ and a shortest path from the wedge point to the basepoint determines a conjugating element h_e. With these choices, the new bases are good. Indeed, the partition of $\mathcal{B}_e^{out}$ is induced by the wedge. More precisely, $(\mathcal{B}_e^{out})'$ corresponds to the set of edges of $\Sigma' \cap \Sigma_e$ not in T and $(\mathcal{B}_e^{out})''$ corresponds to the remaining edges not in T. $\square$

Example A.4 Let $\mathcal{G}$ have underlying graph Γ as in Figure 16. Suppose that $\mathcal{B}_e = \{a_1, a_2\}$, $\mathcal{B}_v = \{b_1, b_2\}$, and $\mathcal{B}_f = \{z\}$. We will not specify G_u since it will not change. Suppose that $\varphi_e(a_1) = b_1^2 b_2^2$, $\varphi_e(a_2) = b_1^2 b_2^2 b_1^2$, $\varphi_f(z) = b_1$, $\varphi_{f^{-1}}(z) = b_2$. Then, $\Sigma(v) = \Sigma_{\mathsf{S}}(\vec{H}(v))$ is a Gersten representative, ie, $\Sigma_{\mathsf{S}}(\vec{H}(v)) = \Sigma_{\mathsf{G}}(\vec{\mathcal{H}}(v))$. $\Sigma(v)$ is displayed in Figure 16, and it can be visibly cleaved. The given bases are not good. A good basis for G_e corresponds to the one determined by the wedge point in $\Sigma(v)$. The change of basis automorphism $\psi_e \in \mathsf{Aut}(G_e)$ is given by $\psi_e(a_1) = a_1^{-1} a_2$ and $\psi_e(a_2) = a_2^{-1} a_1 a_1$, ie, $\varphi_e \circ \psi_e(a_1) = b_1^2$ and $\varphi_e \circ \psi_e(a_2) = b_2^2$.

If $\mathcal{G}'$ is the result of cleaving $\mathcal{G}$, then $\mathcal{B}_{e'} = \{a_1\}$, $\mathcal{B}_{e''} = \{a_2\}$, $\mathcal{B}_{v'} = \{b_1\}$, $\mathcal{B}_{v''} = \{b_2\}$, $\varphi_{e'}(a_1) = b_1^2$, $\varphi_{e''}(a_2) = b_2^2$, $\varphi_{(e')^{-1}}(a_1) = \varphi_{e^{-1}} \circ \psi_e(a_1)$, and $\varphi_{(e'')^{-1}}(a_2) = \varphi_{e^{-1}} \circ \psi_e(a_2)$.

The next step in the algorithm would be to reduce $\mathcal{G}'$. The example could have been complicated by post-composing φ_e, φ_f, and $\varphi_{f^{-1}}$ by some $\psi_v \in \mathsf{Aut}(G_v)$. In that case, we would use Gersten's algorithm first (and discover ψ_v).

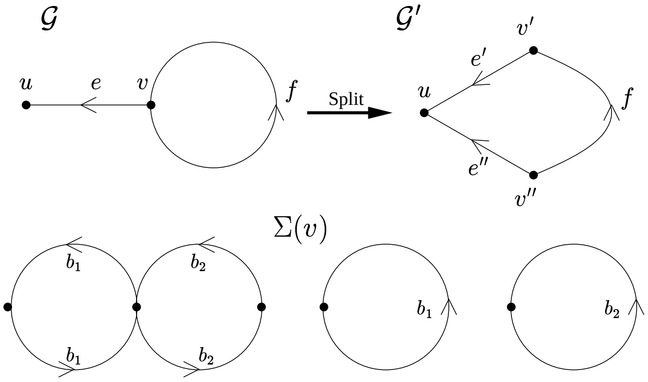

Figure 16

References

[1] **H Bass**, *Some remarks on group actions on trees*, Comm. Algebra 4 (1976) 1091–1126

[2] **M Bestvina**, **M Feighn**, *Outer limits*, preprint (1994)

[3] **M Bestvina**, **M Feighn**, *A combination theorem for negatively curved groups*, J. Differential Geom. 35 (1992) 85–101

[4] **M Bestvina**, **M Feighn**, *Bounding the complexity of simplicial group actions on trees*, Invent. Math. 103 (1991) 449–469

[5] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundlehren series 319, Springer-Verlag, Berlin (1999)

[6] **M R Bridson**, **D T Wise**, *$\mathcal{VH}$ complexes, towers and subgroups of $F \times F$*, Math. Proc. Cambridge Philos. Soc. 126 (1999) 481–497

[7] **P Brinkmann**, *Splittings of mapping tori of free group automorphisms*, Geom. Dedicata 93 (2002) 191–203

[8] **D E Cohen**, *Combinatorial group theory: a topological approach*, London Mathematical Society Student Texts 14, Cambridge University Press, Cambridge (1989)

[9] **G-A Diao**, *Is a graph of finitely generated free groups free? An algorithm*, PhD thesis, Rutgers University, Newark (2003)

[10] **M Feighn**, **M Handel**, *Mapping tori of free group automorphisms are coherent*, Ann. of Math. (2) 149 (1999) 1061–1077

[11] **K Fujiwara**, **P Papasoglu**, *JSJ-decompositions of finitely presented groups and complexes of groups*, to appear in GAFA

[12] **R Geoghegan**, **M L Mihalik**, **M Sapir**, **D T Wise**, *Ascending HNN extensions of finitely generated free groups are Hopfian*, Bull. London Math. Soc. 33 (2001) 292–298

[13] **V Gerasimov**, *Detecting connectedness of the boundary of a hyperbolic group*, preprint (1999)

[14] **S M Gersten**, *On Whitehead's algorithm*, Bull. Amer. Math. Soc. (N.S.) 10 (1984) 281–284

[15] **IA Grushko**, *On generators of a free product of groups*, Matem. Sbornik N. S. 8 (1940) 169–182

[16] **A Hatcher**, *Algebraic topology*, Cambridge University Press, Cambridge (2002)

[17] **W Jaco**, **D Letscher**, **J H Rubinstein**, *Algorithms for essential surfaces in 3-manifolds*, from: "Topology and geometry: commemorating SISTAG", Contemp. Math. 314, Amer. Math. Soc., Providence, RI (2002) 107–124

[18] **S Kalajdžievski**, *Automorphism group of a free group: centralizers and stabilizers*, J. Algebra 150 (1992) 435–502

[19] **O Kharlampovich**, **A Myasnikov**, *Irreducible affine varieties over a free group. I. Irreducibility of quadratic equations and Nullstellensatz*, J. Algebra 200 (1998) 472–516

[20] **O Kharlampovich**, **A Myasnikov**, *Effective JSJ decompositions*, from: "Groups, languages, algorithms", (Borovik, editor), Contemp. Math. 378, Amer. Math. Soc., Providence, RI (2005) 87–212

[21] **R C Lyndon**, **P E Schupp**, *Combinatorial group theory*, Classics in Mathematics, Springer-Verlag, Berlin (2001)

[22] **C F Miller, III**, *On group-theoretic decision problems and their classification*, Princeton University Press, Princeton, N.J. (1971)

[23] **E Rips**, **Z Sela**, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition*, Ann. of Math. (2) 146 (1997) 53–109

[24] **Z Sela**, *Diophantine geometry over groups. I. Makanin-Razborov diagrams*, Publ. Math. Inst. Hautes Études Sci. (2001) 31–105

[25] **J-P Serre**, *Trees*, Springer Monographs in Mathematics, Springer-Verlag, Berlin (2003)

[26] **A Shenitzer**, *Decomposition of a group with a single defining relation into a free product*, Proc. Amer. Math. Soc. 6 (1955) 273–279

[27] **J R Stallings**, *Topology of finite graphs*, Invent. Math. 71 (1983) 551–565

[28] **J R Stallings**, *Foldings of G-trees*, from: "Arboreal group theory (Berkeley, CA, 1988)", Math. Sci. Res. Inst. Publ. 19, Springer, New York (1991) 355–368

[29] **G A Swarup**, *Decompositions of free groups*, J. Pure Appl. Algebra 40 (1986) 99–102

[30] **J H C Whitehead**, *On certain sets of elements in a free group*, Proc. London Math. Soc. 41 (1936) 48–56

[31] **J H C Whitehead**, *On equivalent sets of elements in a free group*, Ann. of Math. (2) 37 (1936) 782–800