Zeitschrift: Commentarii Mathematici Helvetici
Herausgeber: Schweizerische Mathematische Gesellschaft
Band: 68 (1993)

Artikel: Equivariant outer space and automorphisms of free-by-finite groups.
Autor: Krstic, Sava / Vogtmann, Karen
DOI: https://doi.org/10.5169/seals-51767

Comment. Math. Helvetici **68** (1993) 216–262

0010–2571/93/020216–47$1.50 + 0.20/0

Equivariant outer space and automorphisms of free-by-finite groups

SAVA KRSTIĆ[1] AND KAREN VOGTMANN[2]

§1. Introduction

A well-known elementary application of algebraic topology to group theory is the theorem that any group which acts freely on a (simplicial) tree is a free group, and its corollary that any subgroup of a free group is free. A deep generalization of this is the theorem of Karrass–Pietrowski–Solitar, Cohen and Scott which classifies groups with a free subgroup of finite index (free-by-finite groups) as those groups which act on trees with finite stabilizers of bounded order (see, e.g. [8], chapter 4).

In [7], Culler and Vogtmann study the group Out (F_n) of outer automorphisms of a finitely generated free group F_n by constructing an "outer space" X_n of free actions of F_n on simplicial trees. In this paper we generalize their construction to study automorphism groups of finitely generated free-by-finite groups. By results of McCool [13], information about these automorphism groups can be derived by studying centralizers of finite subgroups of Out (F_n). This has been used by Krstić to prove that these groups are finitely generated [11], and then independently by Kalajdžievski and Krstić to prove that they are finitely presented [9, 12].

Given a finite subgroup G of Out (F_n), we construct a simplicial complex L_G on which the centralizer $C(G)$ acts with finite stabilizers and finite quotient. This complex L_G is an equivariant deformation retract of the fixed point subcomplex of outer space X_n. The main theorem of the paper is that L_G is contractible. In addition, we compute the dimension of L_G, thereby giving an upper bound on the virtual cohomological dimension (vcd) of $C(G)$. Under mild hypotheses on G, we show that the dimension of L_G is in fact the vcd. Since the quotient of L_G is finite, $C(G)$ has finitely generated homology in all dimensions. These homological finiteness properties translate directly into similar properties for automorphism groups of

---

[1] Partially supported by Science Fund of Serbia grant #0401A, through Matematički Institut
[2] Partially supported by NSF grant #DMS-8702070

free-by-finite groups. In particular, we show that the vcd of the outer automorphism group of a free product of n finite groups is equal to $n-2$, as conjectured by Collins in [3, 4].

The paper is organized as follows. In Section 2 we review the connection between automorphism groups of free-by-finite groups and centralizers of finite subgroups of Out (F_n). In Section 3 we briefly review [7] and define the complex L_G as the geometric realization of a partially ordered set whose elements are "marked G-graphs" and whose relation is given by collapsing an invariant forest. In Section 4 we develop some theory of invariant forests in G-graphs, and show how maximal invariant forests are related. In Section 5 we give a combinatorial description of the link of a minimal vertex in L_G, which will be used in the proof that L_G is contractible and in computing the dimension of L_G. In Section 6 we define a norm on minimal vertices of L_G and show how the norm changes under a move to a "nearby" minimal vertex. The proof that L_G is contractible is given in Sections 7 and 8. In Section 9 we compute the dimension of L_G, thereby giving an upper bound on the vcd of $C(G)$, and show that this is often equal to the vcd. Other corollaries are collected in Section 10.

The first author would like to thank Cornell University for its hospitality during the time most of this paper was written.

§2. Automorphisms of free-by-finite groups and centralizers

A. McCool's results

In this section, we recall some results of McCool relating automorphisms of free-by-finite groups and centralizers of finite subgroups of Out (F_n). For further details and proofs, see [13] or Section 2 of [12].

Let E be an extension of a finitely generated free group F_n by a finite group K:

$$1 \to F_n \to E \to K \to 1.$$

Let $\mathrm{Aut}_0(E)$ be the subgroup of Aut (E) consisting of automorphisms which send F_n to itself and induce the identity on K. The group $\mathrm{Aut}_0(E)$ has finite index in Aut (E).

Conjugation in E induces a map $\theta : K \to \mathrm{Out}(F_n)$. Let G be the image $\theta(K)$, $C(G)$ the centralizer of G in Out (F_n), and $D(G)$ the preimage of $C(G)$ under the projection Aut $(F_n) \to$ Out (F_n).

PROPOSITION 2.1. *The map* $\mathrm{Aut}_0(E) \to D(G)$ *given by restriction to* F_n *is an isomorphism.* □

Since virtual cohomological finiteness properties are preserved by passing to subgroups of finite index, we have

COROLLARY 2.2. *The virtual cohomological dimension* (vcd) *of* $D(G)$ *is equal to the* vcd *of* $\mathrm{Aut}(E)$. $\mathrm{Aut}(E)$ *is VFL if and only if* $D(G)$ *is VFL.* □

The following commutative diagram allows us to pass to $\mathrm{Out}(E)$:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & F_n & \longrightarrow & \mathrm{Aut}_0(E) & \longrightarrow & \mathrm{Out}_0(E) & \longrightarrow & 1 \\ & & \downarrow{\scriptstyle =} & & \downarrow{\scriptstyle \cong} & & \downarrow & & \\ 1 & \longrightarrow & F_n & \longrightarrow & D(G) & \longrightarrow & C(G) & \longrightarrow & 1 \end{array}$$

where $\mathrm{Out}_0(E)$ is the image of $\mathrm{Aut}_0(E)$ in $\mathrm{Out}(E)$.

COROLLARY 2.3. *The* vcd *of* $C(G)$ *is equal to the* vcd *of* $\mathrm{Out}(E)$. $\mathrm{Out}(E)$ *is VFL if and only if* $C(G)$ *is VFL.* □

B. Graphs of groups and free-by-finite groups

By the theorem of Karrass, Pietrowski and Solitar [10] and the Bass–Serre theory [17], any finitely generated free-by-finite group E as above acts on a tree T with finite stabilizers and finite quotient $E\backslash T$. In particular, there is a graph of groups presentation $\mathcal{G}$ for E based on the graph $E\backslash T$, with finite vertex and edge groups. If E fits in the exact sequence

$$1 \to F_n \to E \to K \to 1$$

then the quotient map $T \to E\backslash T$ factors through $\Gamma = F_n\backslash T$. The graph Γ has fundamental group F_n, and K acts on Γ by graph automorphisms with quotient $E\backslash T$, generating the same graph of groups presentation $\mathcal{G}$ for E.

EXAMPLE 2.4. Let $E = \langle x, y, z \mid x^2, y^2, z^2 \rangle$ be the free product of three copies of $\mathbf{Z}_2$. Representing E as the fundamental group of a graph of groups as in Figure 1a, we get an action of E on the tree T depicted in Figure 1b. The commutator subgroup $[E, E]$ is free of rank 5 (free generators are $xyxy$, $yzyz$, $zxzx$, $xyzyzx$, and $yzxzxy$). We have $K = E/[E, E] = \langle x \mid x^2 \rangle \times \langle y \mid y^2 \rangle \times \langle y \mid y^2 \rangle \cong \mathbf{Z}_2{}^3$. Figure 1c represents the quotient graph $\Gamma = [E, E]\backslash T$. We can think of Γ as the 1-skeleton of

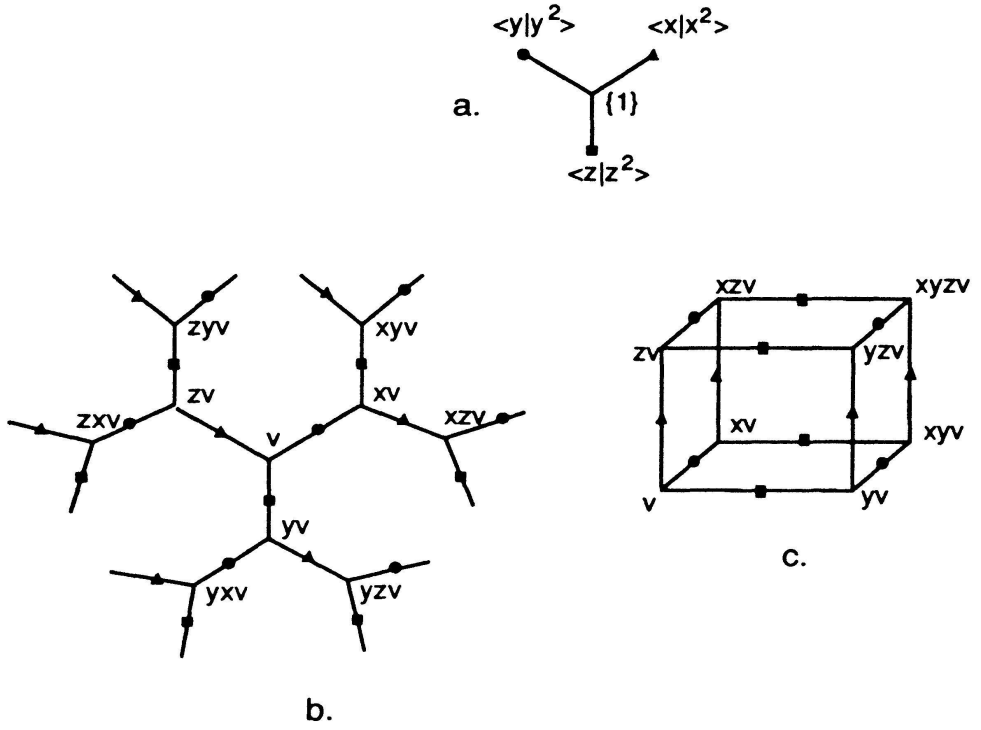

Figure 1

the cube centered at the origin, with edges parallel to the coordinate axes. Then the generators x, y, z of K act on Γ by reflections in the coordinate planes.

§3. The complexes K_G and L_G

A. Review of marked graphs and forest collapse

We briefly recall the definition of the "outer space" X_n. In fact, we will need only the simplicial complex K_n described in [7], which is an equivariant deformation retract of X_n. Our description differs slightly from that given in [7]. A vertex of K_n is a minimal free action of F_n on a simplicial tree. It is useful to describe such an action by considering the quotient graph of the action, as follows. By a *graph* we mean a Serre graph Γ, with vertices $V(\Gamma)$, oriented edges $E(\Gamma)$, involution $e \mapsto e^{-1}$ on $E(\Gamma)$, and initial and terminal vertex maps $\iota, \tau : E(\Gamma) \to V(\Gamma)$. We reserve the right to think of a graph as a topological space as well as a combinatorial object, where a map of graphs always sends vertices to vertices and is locally injective on edges which are not collapsed. Fix a graph R_n with one vertex v and $2n$ oriented edges, and identify F_n with $\pi_1(R_n, v)$. A vertex of K_n is an equivalence class of pairs (s, Γ), where Γ is a connected graph with vertices of valence at least three, and

$s : R_n \to \Gamma$ is a homotopy equivalence. Two pairs (s, Γ) and (s', Γ') are equivalent if there is a graph isomorphism $h : \Gamma \to \Gamma'$ such that $h \circ s \simeq s'$. An equivalence class of pairs is called a *marked graph*; the equivalence class of (s, Γ) will be denoted $\sigma = [s, \Gamma]$.

A vertex σ is said to be obtained from σ' by a *forest collapse* if representatives (s, Γ) and (s', Γ') can be chosen so that Γ is obtained from Γ' by collapsing each connected component of a forest $\mathscr{F}$ in Γ' to a point, and s is the composition of s' with the collapsing map; we write $\Gamma = (\Gamma')_{\mathscr{F}}$. The vertices of K_n form a poset (partially ordered set) under forest collapse. The complex K_n is the geometric realization of this poset, i.e. vertices $\sigma_0, \ldots, \sigma_k$ of K_n span a k-simplex if σ_i is obtained from σ_{i-1} by a forest collapse.

B. The fixed-point subcomplex K_G and Culler's theorem

If ϕ is an automorphism of F_n, represent ϕ^{-1} by a map $f : R_n \to R_n$. Then $\phi(s, \Gamma) = (s \circ f, \Gamma)$ induces an action of Out (F_n) on K_n. For any pair (s, Γ), the map Aut $(\Gamma) \to$ Out (F_n) induced by $h \mapsto (s^{-1} \circ h \circ s)_*$, where s^{-1} is a homotopy inverse for s, induces an isomorphism from the finite group of graph automorphisms of Γ onto the stabilizer of $[s, \Gamma]$ (see [18], Proposition 1.6).

Let G be a finite subgroup of Out (F_n), and let $C(G)$ be the centralizer of G in Out (F_n). To obtain the homological finiteness properties that we want for $C(G)$, we need a contractible complex on which $C(G)$ acts with finite stabilizers and finite quotient. A natural candidate for such a complex is the subcomplex K_G of K_n fixed by G, since the action of Out (F_n) on K_n restricts to an action of $C(G)$ on K_G.

A vertex of K_G is a marked graph $\sigma = [s, \Gamma]$, where Γ comes equipped with a G-action $h : G \to$ Aut (Γ) (i.e. Γ is a *G-graph*), and s is a G-equivariant map in the following sense: represent $x \in G$ by a map $f_x : R_n \to R_n$; then $h(x) \circ s \simeq s \circ f_x$. If the action of G on Γ sends an edge e to its inverse e^{-1}, we subdivide e with a single bivalent vertex so that G acts without inversion on Γ. With this convention, vertices of K_G will be called *marked G-graphs*.

The complex K_G is the geometric realization of a sub-poset of the poset of vertices of K_n, where the poset relation is given by collapsing a G-invariant forest. A G-graph Γ is said to be *reduced* if it contains no G-invariant forest. Thus a minimal vertex of K_G is a marked G-graph $[s, \Gamma]$ where Γ is reduced.

A theorem of Culler [5] implies that any finite subgroup of Out (F_n) can be realized as a subgroup of the stabilizer of some vertex $[s, \Gamma]$ of K_n; thus K_G is not empty.

C. Inessential edges and essential graphs

In order to obtain the best homological finiteness results, we would like the smallest possible $C(G)$-complex with the required properties. The complex K_G has a natural equivariant deformation retract L_G, which we now describe. We will show in Section 9 that L_G is often the smallest possible complex.

An edge of a G-graph Γ is *inessential* if it is contained in every maximal invariant forest of Γ; otherwise it is *essential*. The set of all inessential edges of Γ is itself an invariant forest. If G is the trivial group, the maximal invariant forests are the maximal trees, and the inessential edges are the separating edges.

A G-graph is *essential* if all its edges are essential, all vertices have valence at least two, and the two edges terminating at a bivalent vertex are in the same G-orbit. Note that reduced G-graphs are essential.

EXAMPLE 3.1. Consider the $\mathbf{Z}_3$-graphs Γ_1 and Γ_2 given in Figure 2, the action of the generator of $\mathbf{Z}_3$ being given on either graph by rotation by $2\pi/3$. Neither is reduced. Γ_2 is essential, but Γ_1 is not.

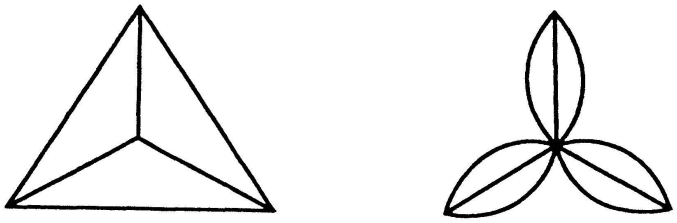

Figure 2

EXAMPLE 3.2. The graph Γ given in Example 2.4 is essential. The orbit of each edge of Γ spans a maximal invariant forest. There are three maximal invariant forests in Γ; each consists of four trees (parallel edges of the cube).

Let L_G denote the subcomplex of K_G spanned by the essential marked G-graphs. The action of the centralizer $C(G)$ on K_G restricts to an action on L_G.

PROPOSITION 3.3. *The subcomplex L_G of K_G spanned by essential marked G-graphs is a deformation retract of K_G.*

Proof. Let $\sigma_0 = [s_0, \Gamma_0]$ be the marked graph obtained from $\sigma = [s, \Gamma]$ by collapsing all inessential edges in Γ. Then the map $\phi : K_G \to L_G$ sending $[s, \Gamma]$ to $[s_0, \Gamma_0]$ is a poset map with $\sigma_0 = \phi(\sigma) \leq \sigma$ for all σ. The Proposition now follows by applying the following lemma, which follows easily from [15], Section 1.3. □

LEMMA 3.4 (*Poset Lemma*). *Let X be a poset and $f : X \to X$ be a poset map with the property that $f(x) \leq x$ for all $x \in X$ (or $f(x) \geq x$ for all $x \in X$). Then $f(X)$ is a deformation retract of X.* □

§4. Invariant forests and the Factorization Lemma

In order to work with the complex L_G, we need to understand the invariant forests in a marked G-graph $[s, \Gamma]$. In this section we characterize invariant forests both in Γ and in the quotient of Γ by the G-action, and prove a Factorization Lemma (Proposition 4.7) which relates two maximal invariant forests in Γ.

A. Characterization of inessential edges and invariant forests

DEFINITION. Let e be an edge of the G-graph Γ and let $v = \iota(e)$ or $\tau(e)$. The endpoint v is *weak* if stab (v) = stab (e). The edge e is *elliptic*, *parabolic* or *hyperbolic* if it has respectively zero, one or two weak endpoints. The edge e is *straight* if its endpoints are in different orbits, and *bent* if its endpoints are in the same orbit.

Let $\bar{\Gamma}$ denote the quotient graph $G\backslash\Gamma$, and $q : \Gamma \to \bar{\Gamma}$ the quotient map. An edge $\bar{e}$ in $\bar{\Gamma}$ is *hyperbolic*, *parabolic* or *elliptic* if any preimage e of $\bar{e}$ is. Straight edges in Γ have distinct endpoints in $\bar{\Gamma}$, and bent edges are loops in $\bar{\Gamma}$.

EXAMPLE 4.1. Let Γ be the $\mathbf{Z}_6$-graph in Figure 3, where the generator of $\mathbf{Z}_6$ acts by

$$a_1 \mapsto a_2 \mapsto a_3 \mapsto a_4 \mapsto a_5 \mapsto a_6 \mapsto a_1, \qquad b_1 \mapsto b_2 \mapsto b_3 \mapsto b_1.$$

The edges a_i are bent elliptic, while the b_j are bent hyperbolic. The free-by-finite group corresponding to this $\mathbf{Z}_6$-graph is $\mathbf{Z} * (\mathbf{Z} \times \mathbf{Z}_2)$.

EXAMPLE 4.2. Let Γ be the graph shown in Figure 4, corresponding to the edges of a truncated cube, with the action of $(\mathbf{Z}_2)^3$ defined by reflections in coordinate planes, as in Example 2.4. The edges are straight hyperbolic and parabolic. The corresponding free-by-finite group is $\mathbf{Z} * \mathbf{Z}_2 * \mathbf{Z}_2 * \mathbf{Z}_2$.

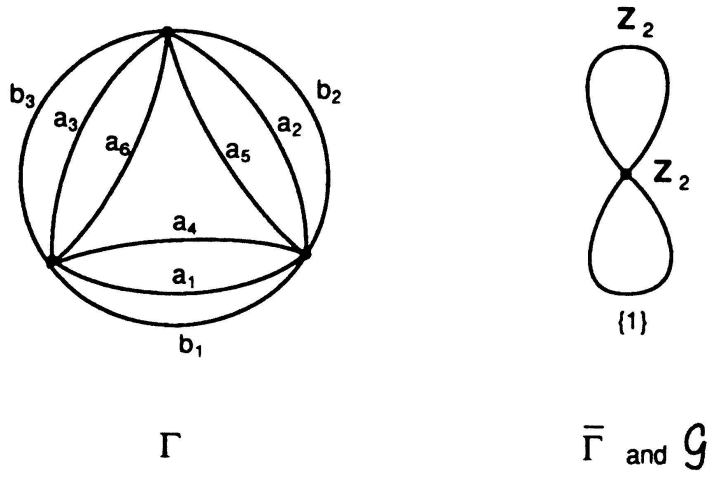

Figure 3

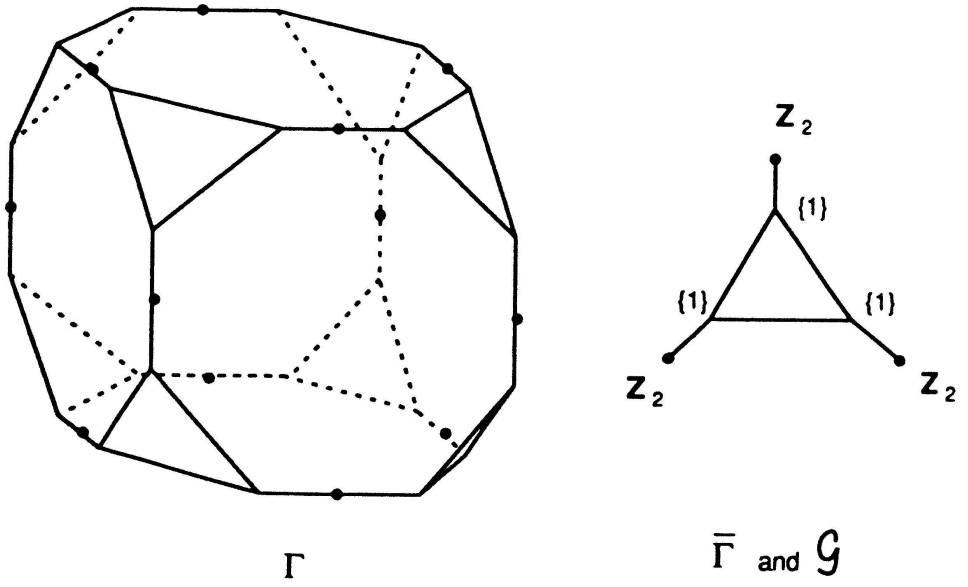

Figure 4

LEMMA 4.3. *Let e be an edge of a G-graph Γ. Then the orbit Ge spans a forest in Γ if and only if e is parabolic or straight hyperbolic.*

Proof. Let v and w be the endpoints of e. If v and w are in different G-orbits, the valence of v in the span of Ge is equal to the index of stab (e) in stab (v); otherwise it is twice the index. Since a forest must contain a free edge, the orbit Ge spans a forest if and only if some endpoint of e has valence one in the span of Ge. □

DEFINITION. A path in Γ is *level* if the stabilizers of all edges and all interior vertices are the same and its projection in $\bar{\Gamma}$ is simple (i.e. does not cross itself). An endpoint of a level path is *weak* if its stabilizer is equal to the stabilizer of the edges

of the path. A level path is *elliptic*, *parabolic* or *hyperbolic* depending on whether it has zero, one or two weak endpoints. A level path is *straight* if its endpoints are in different orbits and *bent* otherwise. A *strong path* is a level path which is either elliptic or bent hyperbolic.

Again, these notions descend to paths in $\bar{\Gamma}$.

PROPOSITION 4.4. *Let $\mathscr{E}$ be an invariant subgraph of Γ. Then $\mathscr{E}$ is a forest if and only if $\mathscr{E}$ contains no strong paths.*

Proof. If $\mathscr{E}$ is not a forest, let C be a simple cycle in $\mathscr{E}$. If all vertex and edge stabilizers of C have the same order, they are all equal and C contains a bent hyperbolic path. If the stabilizers do not have the same order, then C must contain a level elliptic sub-path. In either case, we have found a strong path in $\mathscr{E}$.

Now suppose $\mathscr{E}$ contains a strong path $\mathscr{P}$. We induct on the length of $\mathscr{P}$. If $\mathscr{P}$ has one edge e, then e is elliptic or bent hyperbolic. Thus the span of $G\mathscr{P}$ contains a cycle by Lemma 4.3, so $\mathscr{E}$ is not a forest. If $\mathscr{P}$ has more than one edge, then every edge of $\mathscr{P}$ is either parabolic or straight hyperbolic, so every edge-orbit spans a forest. Choose an edge e of $\mathscr{P}$. Since Ge spans a forest in $G\mathscr{P}$, the collapsing map from $G\mathscr{P}$ to $(G\mathscr{P})_{Ge}$ is a homotopy equivalence. If e is straight hyperbolic, the image of e is a vertex with stabilizer equal to the stabilizer of e. If e is parabolic, its image vertex in Γ_{Ge} has stabilizer equal to the stabilizer of its non-weak endpoint. In either case, the image of $\mathscr{P}$ in Γ_{Ge} is a strong path. By induction, the orbit $(G\mathscr{P})_{Ge}$ contains a cycle. Since $G\mathscr{P} \to (G\mathscr{P})_{Ge}$ is a homotopy equivalence, $G\mathscr{P}$ contains a cycle, so $\mathscr{E}$ is not a forest. □

COROLLARY 4.5. *An edge e of Γ is essential if and only if it belongs to a strong path.*

Proof. Suppose e belongs to a strong path $\mathscr{P}$. By Proposition 4.4, the span $\mathscr{F}$ of $G(\mathscr{P} - e)$ is a forest, but the span of $G\mathscr{P} = \mathscr{F} \cup Ge$ is not; thus e does not belong to every maximal invariant forest, i.e. e is essential.

Now suppose e is essential. Then there is an invariant forest $\mathscr{F}$ such that $\mathscr{F} \cup Ge$ is not a forest. By Proposition 4.4, there is a strong path $\mathscr{P}$ in $\mathscr{F} \cup Ge$ which is not contained in $\mathscr{F}$, i.e. $\mathscr{P}$ must contain a translate xe of e for some $x \in G$. Then $x^{-1}\mathscr{P}$ is a strong path containing e. □

PROPOSITION 4.6. *Let $\Gamma_{\mathscr{F}}$ be obtained from Γ by collapsing the invariant forest $\mathscr{F}$. An edge $e \in \Gamma_{\mathscr{F}}$ is essential if and only if its (unique) lift $\tilde{e}$ in Γ is essential.*

Proof. If $\tilde{e}$ is essential, $\tilde{e}$ is contained in a strong path $\tilde{\mathscr{P}}$. But the image $\mathscr{P}$ of $\tilde{\mathscr{P}}$ in $\Gamma_{\mathscr{F}}$ is strong (as in the proof of Proposition 4.4), so by Corollary 4.5, e is essential since e is an edge of $\mathscr{P}$.

If e is essential, choose a strong path $\mathscr{P}$ containing e. We need to find a strong path $\tilde{\mathscr{P}}$ containing $\tilde{e}$. For each vertex $v \in \Gamma_{\mathscr{F}}$, let T_v be the tree in Γ with image v under the collapsing map. Note that the stabilizer in Γ of v is the maximum of the stabilizers in $\Gamma_{\mathscr{F}}$ of the vertices of T_v. Write $\mathscr{P} = v_0 e_1 v_1 e_2 \cdots e_k v_k$, where $e_i \in E(\Gamma)$ is an edge from v_{i-1} to v_i. Then set $\tilde{\mathscr{P}} = p_0 \tilde{e}_1 p_1 \tilde{e}_2 \cdots \tilde{e}_k p_k$, where $\tilde{e}_i$ is the unique lift of e_i, and p_i is a path in T_{v_i} defined as follows. If $0 < i < k$, then p_i is the unique reduced path from $\tau(\tilde{e}_i)$ to $\iota(\tilde{e}_{i+1})$. If $v_0 = xv_k$ for some $x \in G$, then p_0 is the path from $\tau(x\tilde{e}_k)$ to $\iota(\tilde{e}_1)$, and p_k is constant. If stab (v_0) strictly contains stab (e_1), then p_0 is any shortest path from $\tau(\tilde{e}_1)$ to a vertex in T_{v_0} whose stabilizer properly contains stab (e_1), and p_k is defined similarly. It is straightforward to check that $\tilde{\mathscr{P}}$ is strong. □

B. Orientations and statement of the Factorization Lemma

In this section we state a proposition (the Factorization Lemma) which we will need in the proof that L_G is contractible. The Factorization Lemma finds a particularly nice bijection between the edges of two maximal invariant forests. For the proof it is convenient to put a "natural" orientation on invariant forests.

DEFINITION. An orientation of a tree is *confluent* if at most one edge emanates from any vertex. It follows that there is a unique sink in the tree, and all edges point to it. An orientation of an invariant forest $\mathscr{F}$ in Γ is *natural* if it is equivariant, confluent on each tree in $\mathscr{F}$, and every parabolic edge of $\mathscr{F}$ is oriented away from its weak endpoint.

In Section 5, we will show that every maximal invariant forest in an essential G-graph can be given a natural orientation.

Denote by $E^+\mathscr{F}$ the set of positively oriented edges in the forest $\mathscr{F}$. If $\mathscr{F}$ is an oriented forest whose orientation is confluent on each tree and $e \in E^+(\mathscr{F})$, let $\mathscr{F}_{<e}$ denote the connected component of $\mathscr{F} - \{e\}$ which contains $\iota(e)$. In other words, $\mathscr{F}_{<e}$ is the subtree of $\mathscr{F}$ spanned by all edges of $E^+(\mathscr{F})$ which "point towards" the initial vertex of e.

If $\mathscr{F}$ is a naturally oriented maximal invariant forest in Γ, and a is a positively oriented edge of $\mathscr{F}$, then define $D_{\mathscr{F}}(a)$ to be the set of $e \in E(\Gamma) - E(\mathscr{F})$ such that stab (e) = stab (a), $\tau(e) \in \mathscr{F}_{<a}$ and $\iota(e)$ is not in the orbit of $\mathscr{F}_{<a}$.

PROPOSITION 4.7 (*Factorization Lemma*). *Let $\mathscr{F}$ and $\mathscr{F}'$ be maximal invariant forests in the G-graph Γ, and suppose that $\mathscr{F}$ is naturally oriented. Then there is an orientation of $\mathscr{F}'$ (not necessarily natural) and an equivariant bijection $\phi : E^+(\mathscr{F}) \to E^+(\mathscr{F}')$ such that $\phi(e) \in D_{\mathscr{F}}(e)$.*

C. *Proof of the Factorization Lemma in the case* $G = 1$

A proof of this lemma for trivial G can be found in [7], Lemma 3.3.1. We present here a different, algorithmic proof for this case which helps to motivate the general proof.

Since $G = 1$, the maximal invariant forests $\mathcal{F}$ and $\mathcal{F}'$ are maximal trees in Γ, $\mathcal{F}$ is oriented confluently, $\mathcal{F}_{<a}$ is one of the two components of $\mathcal{F} - a$, and the set $D_{\mathcal{F}}(a)$ is the set of edges in Γ which terminate in $\mathcal{F}_{<a}$ but begin in the other component of $\mathcal{F} - a$.

Proof. The proof proceeds by induction on the number k of vertices of Γ, the case $k = 1$ being vacuous. If $k > 1$, let v be an extremal vertex of $\mathcal{F}$ which is not the sink of $\mathcal{F}$, and let a be the edge of $\mathcal{F}$ with initial vertex v. Draw the geodesic in $\mathcal{F}'$ from the terminal vertex of a to v, and define $\phi(a)$ to be the edge of this geodesic terminating at v.

Now collapse a and delete $\phi(a)$ to obtain a new graph Γ_a with $k - 1$ vertices. The images $\mathcal{F}_a$ and $\mathcal{F}'_a$ of $\mathcal{F}$ and $\mathcal{F}'$ are maximal trees in Γ_a, and $\mathcal{F}_a$ inherits a confluent orientation from $\mathcal{F}$. By induction, we can define ϕ_a on $\mathcal{F}_a$ with the required properties. Define ϕ on $E^+(\mathcal{F}) - \{a\}$ to be the map induced by ϕ_a. Since $\phi_a(e)$ terminates in $(\mathcal{F}_a)_{<e}$, $\phi(e)$ terminates in $\mathcal{F}_{<e}$. Since $\phi_a(e)$ does not begin in $(\mathcal{F}_a)_{<e}$, $\phi(e)$ does not begin in $\mathcal{F}_{<e}$, i.e. for all e, $\phi(e) \in D_{\mathcal{F}}(e)$. □

D. *Admissible forests and clusters*

Since we want the map ϕ in the Factorization Lemma to be equivariant, it is convenient to work in the quotient graph $\bar{\Gamma} = q(\Gamma)$ of Γ by the G-action. The next lemma shows how to recognize invariant forests in Γ by looking in $\bar{\Gamma}$.

DEFINITION. A forest of $\bar{\Gamma}$ is *admissible* if it contains no strong paths.

LEMMA 4.8. *If $\mathcal{F}$ is an invariant forest in Γ, then $q(\mathcal{F})$ is an admissible forest in $\bar{\Gamma}$. If $\mathcal{F}$ is an admissible forest in $\bar{\Gamma}$, then $q^{-1}(\mathcal{F})$ is an invariant forest in Γ.*

Proof. If $\mathcal{F}$ is an invariant forest in Γ, it contains no strong paths by Proposition 4.4. We claim that the image $q(\mathcal{F})$ contains no strong paths. Any strong path in $q(\mathcal{F})$ can be lifted to a strong path in $\mathcal{F}$, contradicting the fact that $\mathcal{F}$ is an invariant forest. In particular, $q(\mathcal{F})$ is a forest, since any cycle would contain a strong path.

If $\mathcal{F}$ is an admissible forest in $\bar{\Gamma}$, then $q^{-1}(\mathcal{F})$ cannot contain a strong path, since the image of a strong path in $q^{-1}(\mathcal{F})$ is a strong path in $\mathcal{F}$. By Proposition 4.4, $q^{-1}(\mathcal{F})$ is an invariant forest. □

We now decompose maximal forests in Γ and $\bar{\Gamma}$ into pieces to which we can apply the $G = 1$ case of the Factorization Lemma.

DEFINITION. Let $H(\Gamma)$ be the set of hyperbolic edges of Γ, and $W(\Gamma)$ the set of weak endpoints of edges in Γ. A *cluster* is a connected component of the subgraph $H(\Gamma) \cup W(\Gamma)$. In particular, a vertex which is the weak endpoint of a parabolic edge but not an endpoint of any hyperbolic edge is a one-point cluster.

For any cluster C, the *extended cluster* PC is the subgraph of Γ spanned by C and the set of parabolic edges whose weak endpoint is in C. A cluster is *bald* if $PC = C$.

The notions of cluster, extended cluster and bald cluster descend to the graph $\bar{\Gamma}$.

Let C be a cluster, and consider the quotient PC^* of PC obtained by identifying all the vertices of $PC - C$ to a single point.

PROPOSITION 4.9. *If $\mathcal{F}$ is a maximal admissible forest in $\bar{\Gamma}$ and C is a cluster in $\bar{\Gamma}$ then the image of $\mathcal{F} \cap PC$ in PC^* is a maximal tree.*

Proof. If the image is not a maximal tree, we can add an adge e of PC to $\mathcal{F}$ without creating any new strong paths; then $\mathcal{F} \cup e$ is an admissible forest, contradicting the maximality of $\mathcal{F}$. □

In Example 4.1, the only cluster is spanned by the orbit $\{b_1, b_2, b_3\}$ and is bald. In the truncated cube (Example 4.2) the clusters are the small triangles; the maximal admissible forests in the quotient are given in Figure 5.

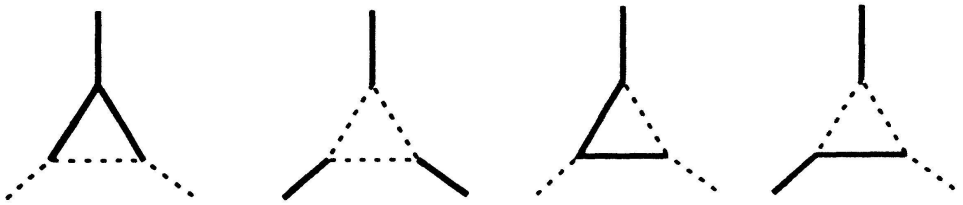

Figure 5

E. Proof of the Factorization Lemma in the general case

Let $\bar{\mathcal{F}}$ and $\bar{\mathcal{F}}'$ be the images of $\mathcal{F}$ and $\mathcal{F}'$ in $\bar{\Gamma}$, and let PC be an extended cluster of $\bar{\Gamma}$. By Proposition 4.9 the images T and T' of $\bar{\mathcal{F}} \cap PC$ and $\bar{\mathcal{F}}' \cap PC$ in PC^* are maximal trees. The confluent orientation on $\mathcal{F}$ induces a confluent

orientation on T. If C is not bald, the sink is the image of the vertices of $PC - C$; otherwise, the sink is a vertex in C. By the case $G = 1$ of the Factorization Lemma, we can define an orientation on T' and a bijection $\bar{\phi}_T : E^+(T) \to E^+(T')$ with $\bar{\phi}_T(e) \in D_T(e)$. The orientation and map lift to an orientation of $\bar{\mathscr{F}}' \cap PC$ and map $\bar{\phi} : E^+(\bar{\mathscr{F}} \cap PC) \to E^+(\bar{\mathscr{F}}' \cap PC)$. Since $\bar{\mathscr{F}} = \bigcup_C (\bar{\mathscr{F}} \cap PC)$, we can define $\bar{\phi}$ on all of $\bar{\mathscr{F}}$ by repeating this procedure.

We claim that $\bar{\phi}(e) \in D_{\bar{\mathscr{F}}}(e)$ for each $e \in \bar{\mathscr{F}}$. Fix an edge e of $\bar{\mathscr{F}}$, and let PC be the extended cluster containing e. Since the preimage of $T_{<e}$ is contained in $\bar{\mathscr{F}}_{<e}$, we have $\tau(\bar{\phi}(e)) \in \bar{\mathscr{F}}_{<e}$ and $\iota(\bar{\phi}(e))$ is not in $\bar{\mathscr{F}}_{<e}$.

Now lift the orientation of $\bar{\mathscr{F}}'$ to an orientation on $\mathscr{F}'$, and define a lift $\phi : E^+(\mathscr{F}) \to E^+(\mathscr{F}')$ of the map $\bar{\phi}$ as follows. If $e \in E^+(\mathscr{F})$, let C be the cluster of Γ containing the weak endpoint of e. Then $\mathscr{F}_{<e} \cap PC \subset C$. The quotient map $\Gamma \to \bar{\Gamma}$ is a covering map on C, since the stabilizer of each edge of C is equal to the stabilizer of each vertex in C. Therefore, $\mathscr{F}_{<e} \cap C$ maps homeomorphically onto $\bar{\mathscr{F}}_{<\bar{e}} \cap \bar{C}$. We know $\bar{\phi}(\bar{e})$ terminates in $\bar{\mathscr{F}}_{<\bar{e}} \cap \bar{C}$. Define $\phi(e)$ to be the unique lift of $\bar{\phi}(\bar{e})$ which terminates in $\mathscr{F}_{<e}$. Since $\bar{\phi}(\bar{e})$ does not begin in $\bar{\mathscr{F}}_{<\bar{e}}$, $\phi(e)$ does not begin in any translate of $\mathscr{F}_{<e}$. Thus $\phi(e) \in D_{\mathscr{F}}(e)$ as desired. □

§5. The star of a reduced marked G-graph

In this section we give a combinatorial description of the star of a minimal vertex in L_G, i.e. of a reduced marked G-graph. We define an operation called "blowing up" which produces a new marked G-graph, and determine when two blow-ups produce the same marked G-graph. We then show that a blow-up of a reduced marked G-graph is essential, and that every essential G-graph can be obtained by blowing up some reduced marked G-graph.

We will use this combinatorial description in the proof of contractibility and to compute the dimension of L_G.

A. Ideal edges and ideal pairs

DEFINITION. Let v be a vertex of the G-graph Γ, and let E_v be the set of oriented edges terminating at v. Let α be a subset of E_v, and let P be the subgroup of stab (v) generated by stabilizers of edges in α. Then α is an *ideal edge* if

(i) card $(\alpha) \geq 2$ and card $(E_v - \alpha) \geq 2$;

(ii) $Ge \cap \alpha = Pe$ for all $e \in \alpha$;

(iii) For some $a \in \alpha$, stab $(a) = P$ and $a^{-1} \notin \bigcup G\alpha$.

The subgroup P is the *stabilizer* of α. Denote by $D(\alpha)$ the set of all edges which satisfy condition (ii). A pair (α, a) is *ideal* if α is an ideal edge and $a \in D(\alpha)$.

Let α be an ideal edge of Γ at a vertex v. If $E_v - \alpha$ is an ideal edge which is not a translate $x\alpha$ for some $x \in G$, then $\alpha^{-1} = E_v - \alpha$ is the *inverse* of α. In all other cases, α is not invertible.

The following lemma characterizes invertible ideal edges in reduced G-graphs.

LEMMA 5.1. *Let Γ be a reduced G-graph, and let α be an ideal edge of Γ at a vertex v. Then α is invertible if and only if* $\text{stab}(\alpha) = \text{stab}(v)$.

Proof. Let $a \in D(\alpha)$. Since Γ is reduced, Ga is not a forest. By Lemma 4.3, a is either bent hyperbolic or elliptic. In the first case, $\text{stab}(\alpha) = \text{stab}(v)$, and we claim that α is invertible, i.e. $\beta = E_v - \alpha$ is an ideal edge. Conditions (i) and (ii) are clear. Since a is bent hyperbolic, $\tau(a^{-1}) = \iota(a) = x^{-1}v$ for some $x \in G$, so $\tau(xa^{-1}) = x(x^{-1}v) = v$, i.e. xa^{-1} terminates at v but is not in α. Since $\text{stab}(xa^{-1}) = \text{stab}(v) = \text{stab}(a)$, we have $xa^{-1} \in D(\beta)$. Now assume a is elliptic. If $\beta = E_v - \alpha$ is an ideal edge, choose $b \in D(\beta)$. Then b must also be elliptic by the argument above. Choose elements $x \in \text{stab}(v) - \text{stab}(b)$ and $y \in \text{stab}(v) - \text{stab}(a)$. Then $x\beta \subset \alpha$ and $y\alpha \subset \beta$; in particular, α and β have the same number of elements, so in fact $\beta = x\alpha$. □

COROLLARY 5.2. *Let α and β be ideal edges with $\alpha \subset \beta$. If α is invertible, then β is invertible.* □

B. Blowing up an ideal edge-orbit

If α is an ideal edge, then $x\alpha$ is also an ideal edge, for any $x \in G$. The set $G\alpha = \{x\alpha \mid x \in G\}$ is an *ideal edge-orbit*. If P is the stabilizer of α, then $G\alpha$ has $p = [G: P]$ elements, and the stabilizer of $x\alpha$ is xPx^{-1}.

Given a marked graph $\sigma = [s, \Gamma]$ and an ideal edge-orbit $G\alpha$, we obtain a new marked G-graph $\sigma^{G\alpha} = [s^{G\alpha}, \Gamma^{G\alpha}]$ by *blowing up the ideal edge-orbit* $G\alpha$ as follows. Let P be the stabilizer of α, with index $p = [G: P]$. The vertices and edges of $\Gamma^{G\alpha}$ are the same as the vertices and edges of Γ, with one additional vertex orbit $Gv(\alpha)$ and additional edge orbits $Ge(\alpha)$ and $Ge(\alpha)^{-1}$. The new vertex $v(\alpha)$ has stabilizer P. The new edge $e(\alpha)$ begins at $v(\alpha)$, terminates at v and also has stabilizer P. For $x \in G$, the edges which terminate at $xv(\alpha)$ are $xe(\alpha)^{-1}$ and the elements of $x\alpha$. The terminus of an edge which is not in $\bigcup G\alpha$ remains unchanged. The orbit $Ge(\alpha)$ spans an invariant forest in $\Gamma^{G\alpha}$, and Γ can be recovered from $\Gamma^{G\alpha}$ by collapsing this

invariant forest. Collapsing $Ge(\alpha)$ is an equivariant homotopy equivalence from $\Gamma^{G\alpha}$ to Γ. Choose an equivariant homotopy inverse $f: \Gamma \to \Gamma^{G\alpha}$ and define $s^{G\alpha}$ to be $f \circ s$. The marked graph $[s^{G\alpha}, \Gamma^{G\alpha}]$ is said to be obtained from $[s, \Gamma]$ by *blowing up* $G\alpha$.

EXAMPLE 5.3. Consider the $\mathbf{Z}_4$-graph Γ in Figure 6; the action of the generator x of $\mathbf{Z}_4$ is defined by $a_1 \mapsto a_2 \mapsto a_3 \mapsto a_4 \mapsto a_1$, $b_1 \leftrightarrow b_2$. The set $\alpha = \{a_1, a_3, b_1\}$ is an ideal edge with stab $(\alpha) = \mathbf{Z}_2$ and $D(\alpha) = \{b_1\}$. The figure presents also the graph obtained by blowing up the ideal edge-orbit $\mathbf{Z}_4\alpha = \{\alpha, x\alpha\} = \{\{a_1, a_3, b_1\}, \{a_2, a_4, b_2\}\}$.

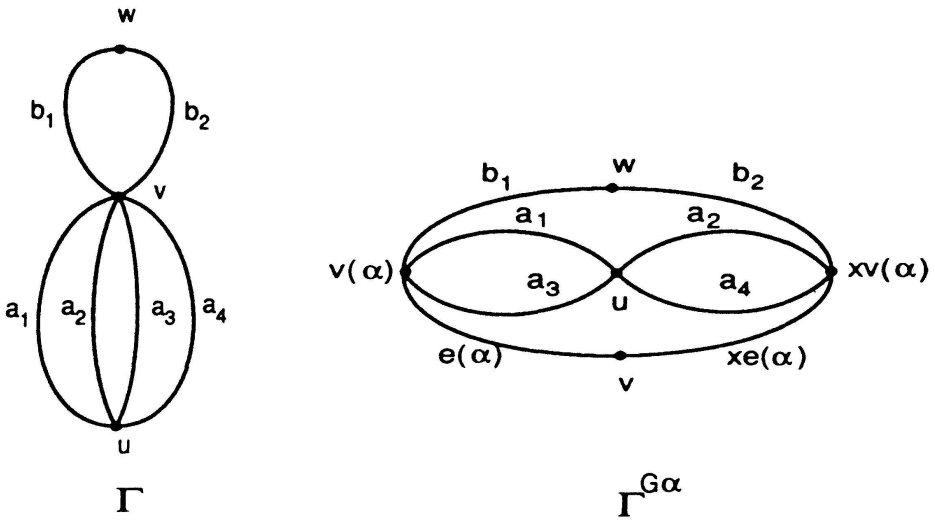

Figure 6

EXAMPLE 5.4. In Figure 7 we have a $\mathbf{Z}_3$-graph Γ; the action of the generator x is given by $a_1 \mapsto a_2 \mapsto a_3 \mapsto a_1$, $b_1 \mapsto b_2 \mapsto b_3 \mapsto b_1$. The set $\alpha = \{a_1, b_1\}$ is an invertible ideal edge with stab $(\alpha) = \mathbf{Z}_3 =$ stab (v), $D(\alpha) = \alpha$ and $\alpha^{-1} = \{a_2^{-1}, b_2^{-1}\}$. The figure also presents the blow-up $\Gamma^{\mathbf{Z}_3\alpha}$, where $\mathbf{Z}_3\alpha = \{\alpha, x\alpha, x^2\alpha\}$.

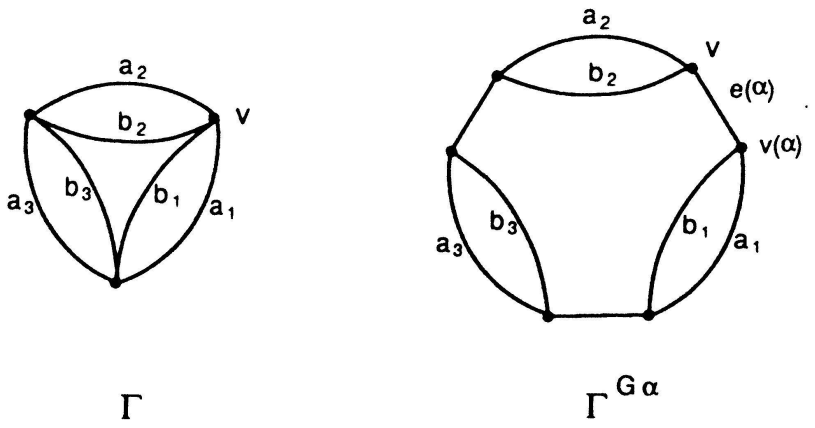

Figure 7

C. *Compatibility*

The operation of blowing up can be done with several ideal edge-orbits at once if they satisfy certain compatibility conditions.

DEFINITION. An ideal edge-orbit $G\alpha$ is *included* in the ideal edge-orbit $G\beta$ if α' is a subset of β' for some ideal edges $\alpha' \in G\alpha$ and $\beta' \in G\beta$. Ideal edge-orbits $G\alpha$ and $G\beta$ are *disjoint* if $\alpha' \cap \beta' = \varnothing$ for all $\alpha' \in G\alpha$ and $\beta' \in G\beta$, i.e. $(\bigcup G\alpha) \cap (\bigcup G\beta) = \varnothing$. They are *inverse* if an edge of $G\alpha$ is invertible, and its inverse is in $G\beta$.

Ideal edge-orbits $G\alpha$ and $G\beta$ are *compatible* if one is included in the other or if they are disjoint and not inverse.

Note that if $G\beta$ is either disjoint from $G\alpha$ or included in $G\alpha$, then $G\beta$ can be regarded as an ideal edge-orbit of $\Gamma^{G\alpha}$. Moreover, if $G\alpha$ and $G\beta$ are disjoint, one has $(\Gamma^{G\alpha})^{G\beta} = (\Gamma^{G\beta})^{G\alpha}$.

D. *Blowing up an oriented ideal forest*

Let $[s, \Gamma]$ be a marked G-graph, and let $\{G\alpha_1, \ldots, G\alpha_k\}$ be a pairwise compatible collection of ideal edge-orbits of Γ. The union $\Phi = G\alpha_1 \cup \cdots \cup G\alpha_k$ is called an *oriented ideal forest*. An oriented ideal forest Φ gives rise to a new marked G-graph $[s^\Phi, \Gamma^\Phi]$ as follows. Partially order the ideal edges in Φ by inclusion. Since orbits of Φ are pairwise compatible, maximal elements of Φ are disjoint, and the set of all maximal elements of Φ is a union of pairwise disjoint ideal edge-orbits. We may blow these ideal edge-orbits up, in any order, to obtain a new marked G-graph $[s', \Gamma']$. The remaining ideal edge-orbits can be regarded as ideal edge-orbits in Γ', so the operation can be repeated until all ideal edge-orbits in Φ are used; the final result is the marked G-graph $[s^\Phi, \Gamma^\Phi]$. We say $[s^\Phi, \Gamma^\Phi]$ is obtained from $[s, \Gamma]$ by *blowing up the oriented ideal forest* Φ. The graph Γ^Φ can be described explicitly as follows, where $+$ denotes disjoint union.

Vertices. $V(\Gamma^\Phi) = V(\Gamma) + \{v(\alpha) \mid \alpha \text{ is an ideal edge in } \Phi\}$.

Edges. $E(\Gamma^\Phi) = E(\Gamma) + \{e(\alpha)^{\pm 1} \mid \alpha \text{ is an ideal edge in } \Phi\}$.

Incidence maps. Let ι and τ be the initial and terminal vertex maps of Γ, and ι_Φ and τ_Φ the corresponding maps of Γ^Φ.

If $\alpha \subset E_v$ is an ideal edge in Φ, then

$$\iota_\Phi(e(\alpha)) = v(\alpha),$$

$$\tau_\Phi(e(\alpha)) = \begin{cases} v & \text{if } \alpha \text{ is maximal} \\ v(\beta) & \text{if } \beta \text{ is the smallest ideal edge in } \Phi \text{ with } \alpha \subset \beta. \end{cases}$$

If $e \in E(\Gamma)$, then

$$\tau_\Phi(e) = \begin{cases} \tau(e) & \text{if } e \notin \bigcup \Phi \\ v(\alpha) & \text{if } \alpha \text{ is the smallest ideal edge in } \Phi \text{ which contains e,} \end{cases}$$

$$\iota_\Phi(e) = \tau_\Phi(e^{-1}).$$

The edges $e(\alpha)$ for $\alpha \in \Phi$ form a forest $\mathscr{F}$ in Γ^Φ, which is naturally oriented, in the sense of Section 4B; the sink of a component T of $\mathscr{F}$ is a vertex whose stabilizer contains the stabilizers of all vertices of T.

E. Blowing up a reduced marked G-graph

The operation of blowing up an oriented ideal forest of Γ produces a new marked G-graph. In this section we show that for reduced marked G-graphs, blowing up always produces an essential marked G-graph.

PROPOSITION 5.5. *Let $[s, \Delta]$ be a reduced marked G-graph, and Φ an oriented ideal forest of Φ. Then the marked G-graph obtained from Δ by blowing up Φ is essential.*

Proof. Let $\Gamma = \Delta^\Phi$, and let $\mathscr{F}$ be the subgraph of Γ spanned by edges $e(\alpha)$, $\alpha \in \Phi$. Then, as remarked above, $\mathscr{F}$ is an invariant forest in Γ, and $\Gamma_{\mathscr{F}} = \Delta$. Since Δ is reduced, $\mathscr{F}$ is maximal in Γ.

Note that all edges of Δ are essential, since Δ has no invariant forests. Since $\Delta = \Gamma_{\mathscr{F}}$, Lemma 4.6 implies that all edges in $\Gamma - \mathscr{F}$ are essential in Γ. If $e \in E^+(\mathscr{F})$, then $e = e(\alpha)$ for some $\alpha \in \Phi$ and $\Gamma_{\mathscr{F} - Ge} = \Delta^{G\alpha}$; thus it suffices to check that $e(\alpha)$ is essential in $\Delta^{G\alpha}$. Let $a \in D(\alpha)$. Then Ga is a forest in $\Delta^{G\alpha}$. If e were inessential, then $Ga \cup Ge$ would span a forest in $\Delta^{G\alpha}$; thus Ga would span a forest in Δ, contradicting the assumption that Δ is reduced.

It remains to check the valence conditions on the vertices of Γ. We claim that the valence of $v(\alpha)$ is at least 3 for every $\alpha \in \Phi$. If $v(\alpha)$ is extremal in $\mathscr{F}$, all edges of α end at $v(\alpha)$ and there are at least two of them. If the valence of $v(\alpha)$ in $\mathscr{F}$ is equal to two, then $v(\alpha) = \tau(e(\beta))$ for some β, so $\beta \subset \alpha$ and edges of $\alpha - \beta$ end at $v(\alpha)$.

Now let $v \in V(\Delta)$ and let E_v and E_v^Φ be the sets of edges ending at v in Δ and Γ. Let $\alpha_1, \ldots, \alpha_k (k \geq 0)$ be all maximal ideal edges in Φ contained in E_v. If $k = 0$, then $E_v^\Phi = E_v$. If $k = 1$, then $E_v^\Phi = \{e(\alpha_1)\} \cup (E_v - \alpha_1)$, so the valence of v in Γ is at least three. If $k = 2$ and $E_v = \alpha_1 \cup \alpha_2$, then $\alpha_2 = x\alpha_1$ for some $x \in G$, so the valence of v in Γ is two and the two edges at v in Γ are in the same orbit. If $k = 2$ and

$E_v - (\alpha_1 \cup \alpha_2) \neq \varnothing$, then E_v^Φ contains $e(\alpha_1), e(\alpha_2)$ and at least one more element. Finally, if $k \geq 3$, then obviously the valence of v in Γ is at least three.

F. *Essential marked G-graphs*

Let $[s, \Gamma]$ be an essential marked G-graph, and $\mathscr{F}$ a maximal invariant forest in Γ. In this section we will find an ideal forest Φ in the reduced marked G-graph obtained from Γ by collapsing $\mathscr{F}$, so that blowing up Φ results in $[s, \Gamma]$.

Recall from Section 4B that an orientation of $\mathscr{F}$ is *natural* if it is invariant, confluent on each component and each parabolic edge is oriented away from its weak endpoint.

PROPOSITION 5.6. *Let $[s, \Gamma]$ be an essential marked G-graph, $\mathscr{F}$ a naturally oriented maximal invariant forest in Γ, and $[r, \Delta]$ the reduced marked G-graph obtained by collapsing $\mathscr{F}$. Then $\mathscr{F}$ determines an oriented ideal forest Φ of Δ with $[s, \Gamma] = [r^\Phi, \Delta^\Phi]$.*

Proof. For every $a \in E^+(\mathscr{F})$, define $\alpha(a) = \{e \in E(\Gamma) - E(\mathscr{F}) \mid \tau(e) \in \mathscr{F}_{<a}\}$. Since the edges of Δ are naturally identified with $E(\Gamma) - E(\mathscr{F})$, we may consider $\alpha(a)$ as a subset of edges of Δ.

CLAIM 1. *$\alpha(a)$ is an ideal edge of Δ.*

Proof. Since all edges in $\alpha(a)$ terminate in the same connected component of $\mathscr{F}$, their images all terminate at the same vertex v of Δ. An element of G stabilizes a if and only if it stabilizes $\mathscr{F}_{<a}$, if and only if it stabilizes $\alpha(a)$, i.e. the stabilizer of $\alpha(a)$ is equal to the stabilizer of a. We now check the conditions for $\alpha = \alpha(a)$ to be an ideal edge.

(i) card $(\alpha) \geq 2$ and card $(E_v(\Delta) - \alpha) \geq 2$.

Let T be the component of $\mathscr{F}$ which contains a, and let w be an extremal vertex of T. If w were bivalent in Γ, then both edges incident to w would be in $\mathscr{F}$ since they are in the same G-orbit. Thus at least three edges of Γ terminate at w, including at least two which are not in $\mathscr{F}$. By considering $w \in T_{<a}$, we see that card $(\alpha) \geq 2$, and taking $w \in T - T_{<a}$ shows that card $(E_v(\Delta) - \alpha) \geq 2$.

(ii) $Ge \cap \alpha = \text{stab}(\alpha)e$ for $e \in \alpha$.

Let x be an element of G which is not in stab (α). Then $xe \notin \alpha$ since xe terminates at $\mathscr{F}_{<xa}$ and the trees $\mathscr{F}_{<a}$ and $\mathscr{F}_{<xa}$ are equal or disjoint depending on whether x is in stab (α) or not.

(iii) $D(\alpha) \neq \varnothing$.

We must show that there is an edge $e \in \alpha$ with stab $(e) =$ stab (a), $\tau(e) \in \mathscr{F}_{<a}$ and $\iota(e) \notin G\mathscr{F}_{<a}$.

Since Γ is essential, there is a strong path $\mathscr{P}$ in Γ which contains a (Proposition 4.4). Let $\bar{\mathscr{P}}$ be the image of $\mathscr{P}$ in $\bar{\Gamma}$. Since $\mathscr{P}$ is strong, $\bar{\mathscr{P}}$ is also strong. Starting at $\iota(\bar{a})$ and travelling along $\bar{\mathscr{P}}$ in the direction opposite to the orientation of $\bar{a}$ we must come across a vertex which is not in $\bar{\mathscr{F}}_{<\bar{a}}$. Indeed, if $\mathscr{P}$ is hyperbolic, then $\bar{\mathscr{P}}$ is a cycle passing through $\tau(\bar{a})$, which is not in $\bar{\mathscr{F}}_{<\bar{a}}$. If $\mathscr{P}$ is elliptic, then its endpoints have stabilizers which contain stab (a), so their images in $\bar{\Gamma}$ are not in $\bar{\mathscr{F}}_{<\bar{a}}$. It follows that there must be an edge $\bar{e}$ in $\bar{\mathscr{P}}$ with $\iota(\bar{e}) \notin \bar{\mathscr{F}}_{<\bar{a}}$ and $\tau(\bar{e}) \in \bar{\mathscr{F}}_{<\bar{a}}$. Let e be a lift of $\bar{e}$ which terminates in $\mathscr{F}_{<a}$. Clearly $\iota(e) \notin G\mathscr{F}_{<a}$. Since $\tau(e) \in \mathscr{F}_{<a}$, we have stab $(e) \subseteq$ stab $(\mathscr{F}_{<a}) =$ stab (a). Since e is a translate of an edge of $\mathscr{P}$ and $\mathscr{P}$ is a strong path, it follows that stab (e) is a conjugate of stab (a), so stab $(e) =$ stab (a).

CLAIM 2. *Let* $\Phi = \{\alpha(a) \mid a \in E^+(\mathscr{F})\}$. *Then* Φ *is an oriented ideal forest of* Δ.

Proof. We have to check that $\alpha(a)$ and $\alpha(b)$ are compatible for all edges a and b in $E^+(\mathscr{F})$. For any such a and b, either $\mathscr{F}_{<a} \cap x\mathscr{F}_{<b} = \mathscr{F}_{<a} \cap \mathscr{F}_{<xb}$ is empty for every $x \in G$, or, for some $x \in G$, there exists an oriented path in $\mathscr{F}$ which contains both a and xb. It follows that $\alpha(a) \cap x\alpha(b) = \varnothing$ for every $x \in G$ or, for some $x \in G$, $\alpha(a) \subset x\alpha(b)$ or $\alpha(a) \supset x\alpha(b)$.

Finally, if for some $a, b \in E^+(\mathscr{F})$ and $v \in V(\Delta)$ we have $\alpha(a) \cup \alpha(b) = E_v(\Delta)$, then it follows that a and b terminate at the same bivalent vertex of Γ. But then a and b are in the same G-orbit, so $\alpha(a) = x\alpha(b)$ for some $x \in G$. □

Finally, we show $\Delta^\Phi = \Gamma$.

Proof. We define a map from Γ to Δ^Φ by sending a to itself if $a \notin \mathscr{F}$, and to $e(\alpha(a))$ if $a \in \mathscr{F}$. The description of Δ^Φ given in Section 5D shows that this map is an isomorphism. □

G. *Ideal forests and the star of a reduced marked G-graph*

Blowing up two different oriented ideal forests may result in the same marked G-graph. We now determine exactly when this happens.

Ideal edge-orbits A and B are called *pre-compatible* if either A and B are compatible or A is invertible and A^{-1} is included in B. In the second case, B is necessarily invertible, and B^{-1} is included in A as well. Note that an ideal edge-orbit is compatible with itself, but not with its inverse; however, it is pre-compatible with its inverse. If A and B are not pre-compatible they are said to *cross*.

The union of a collection of ideal edge-orbits is an *ideal forest* if its elements are pairwise pre-compatible and if it contains the inverse of each of its invertible edge-orbits.

If Φ is an oriented ideal forest in Γ, it can be completed to an ideal forest $\Phi^{\pm}$ by adding the inverses of all invertible edges in Φ.

LEMMA 5.7. *Let Φ and Ψ be oriented ideal forests in the reduced marked G-graph $[s, \Gamma]$. If $\Phi^{\pm} = \Psi^{\pm}$ then $[s^{\Psi}, \Gamma^{\Psi}] = [s^{\Phi}, \Gamma^{\Phi}]$.*

Proof. Since $\Phi^{\pm} = \Psi^{\pm}$, there is a bijection $f: \Phi \to \Psi$ sending each ideal edge α either to itself or to its inverse. Let $I = \{\alpha \in \Phi \mid f(\alpha) = \alpha^{-1}\}$. Let α be a maximal element of I. We claim that α is maximal in Φ. If $\alpha \subset \beta$ for some $\beta \in \Phi$, then $G\alpha^{-1}$ is not compatible with $G\beta$. But $Gf(\alpha) = G\alpha^{-1}$ is compatible with $Gf(\beta)$, so we must have $f(\beta) = \beta^{-1}$, i.e. $\beta \in I$, contradicting the maximality of α.

The proof now proceeds by induction on the number k of G-orbits in I. If $k = 1$ then $I = G\alpha$ for some ideal edge α. Let v be the terminus of edges in α. Since α is maximal in Φ, the edge $e(\alpha)$ in Γ^{Φ} begins at $v(\alpha)$ and terminates at v. Define a map from Γ^{Φ} to Γ^{Ψ} by sending $xv \mapsto xv(\alpha^{-1})$, $xv(\alpha) \mapsto xv$, $xe(\alpha)^{\pm 1} \mapsto xe(\alpha^{-1})^{\mp 1}$ for all $x \in G$, and fixing all other vertices and edges. This induces a graph isomorphism giving $[s^{\Phi}, \Gamma^{\Phi}] = [s^{\Psi}, \Gamma^{\Psi}]$.

Now assume $k > 1$, and let β be any ideal edge in Φ different from α. Since α is maximal in Φ, we must have $G\beta$ included in $G\alpha$ or disjoint from $G\alpha$. In either case, $G\beta$ is compatible with $G\alpha^{-1}$, so the set $\Phi_1 = \Phi - G\alpha + G\alpha^{-1}$ is an oriented ideal forest. By induction, $[s^{\Phi}, \Gamma^{\Phi}] = [s^{\Phi_1}, \Gamma^{\Phi_1}] = [s^{\Psi}, \Gamma^{\Psi}]$. □

LEMMA 5.8. *Every invariant forest $\mathscr{F}$ in an essential G-graph Γ can be given a natural orientation.*

Proof. Let T be a connected component of $\mathscr{F}$. Let v be a vertex of T whose stabilizer has maximal order among stabilizers of vertices of T. If w is any other vertex of T, let $\mathscr{P}$ be the geodesic path from w to v. Since $\mathscr{P}$ contains no strong sub-paths (in particular, each edge of $\mathscr{P}$ is either hyperbolic or parabolic), the stabilizers of vertices on $\mathscr{P}$ must be linearly ordered by inclusion. In particular, all vertex stabilizers of T are contained in stab (v), and vertices of T with this stabilizer span a subtree of T.

Orient T confluently towards v. For parabolic edges, "towards v" means "away from its weak endpoint", so the orientation satisfies the conditions for being natural. Since stab (v) = stab (T), we can extend equivariantly to get a natural orientation on the invariant forest GT. Continue to orient orbits of trees in $\mathscr{F}$ in this way until all of $\mathscr{F}$ is oriented naturally. □

Let $\rho = [r, \Delta]$ be a reduced marked G-graph. Recall that the star of ρ in L_G is the geometric realization of the poset of marked G-graphs which collapse to ρ, with forest collapse as the poset relation. The set of ideal forests of Δ also forms a poset, under inclusion.

PROPOSITION 5.9. *Let $\rho = [r, \Delta]$ be a reduced marked G-graph. Then blowing up induces a poset isomorphism from the poset of ideal forests of Δ to the star of ρ.*

Proof. Denote the poset of ideal forests of Δ by $IF(\Delta)$, and the vertices of the star of ρ in L_G by $ST(\rho)$.

Let Φ be an oriented ideal forest of Δ. Define a map from the set of oriented ideal forests in Δ to the star of ρ in K_G by sending Φ to $[r^\Phi, \Delta^\Phi]$. By Proposition 5.5 the image lies in $ST(\rho)$, and by Lemma 5.7, the map induces a map $\phi : IF(\Delta) \to ST(\rho)$.

We define an inverse $\psi : ST(\rho) \to IF(\Delta)$ to ϕ as follows. If $[s, \Gamma]$ is in the star of ρ, then for some maximal invariant forest $\mathscr{F}$ of Γ, ρ is obtained from $[s, \Gamma]$ by collapsing $\mathscr{F}$. By Lemma 5.8, we may put a natural orientation on $\mathscr{F}$; by Proposition 5.6, $\mathscr{F}$ then determines an oriented ideal forest Φ of Δ, with $[r^\Phi, \Delta^\Phi] = [s, \Gamma]$. Different choices of natural orientation on $\mathscr{F}$ correspond to different choices of orientation on the ideal forest $\Phi^\pm$, so the map $\psi([s, \Gamma]) = \Phi^\pm$ is a well-defined inverse to ϕ. □

H. Whitehead moves and connectivity of L_G

In this section we apply results of Krstić [11] to show that L_G is connected.

LEMMA 5.10. *Let $\alpha \subset E_v$ be an ideal edge of the G-graph Γ and let $a \in D(\alpha)$. Then the orbits Ga and $Ge(\alpha)$ each span an invariant forest in $\Gamma^{G\alpha}$.*

Proof. The vertex $v(\alpha)$ is an endpoint of both $e(\alpha)$ and a in $\Gamma^{G\alpha}$, and $v(\alpha)$, $e(\alpha)$ and a all have stabilizer equal to the stabilizer of α. The endpoints v and $v(\alpha)$ of $e(\alpha)$ are in different G-orbits by construction. Since $a \in D(\alpha)$, the endpoints of a are also in distinct G-orbits. The result follows by Lemma 4.3. □

Since Ga is an invariant forest in $\Gamma^{G\alpha}$, we may collapse it to obtain a new reduced marked G-graph $[s', \Gamma']$. The path $[s, \Gamma] - [s^{G\alpha}, \Gamma^{G\alpha}] - [s', \Gamma']$ in K_G is called a *Whitehead move*, and the marked G-graph $[s', \Gamma']$ will be denoted $[s, \Gamma](G\alpha, Ga)$. On the level of edge-path groupoids, a Whitehead move is expressed by formulas which generalize Whitehed automorphisms of free groups (see [11, 12] for details). The following proposition can be deduced from Corollary 1 and Proposition 4 of [11].

PROPOSITION 5.11. *Any two reduced marked G-graphs are connected by a sequence of Whitehead moves.* □

Since any marked G-graph is connected to a reduced marked G-graph, we have

COROLLARY 5.12. *The complex L_G is connected.* □

§6. The norm of a reduced graph

In this section we define a norm on reduced marked G-graphs. We will use this norm to filter the complex L_G by subcomplexes. We will prove L_G is contractible inductively by showing the subcomplexes are contractible.

A. Dot product and absolute value

Let $\sigma = [s, \Gamma]$ be a marked G-graph, and w a conjugacy class of elements of F_n. Represent w by a cyclically reduced word in the edges of Γ, and define the *star graph* of w to be the graph with one vertex for each oriented edge of Γ and one edge from the vertex e to the vertex f for each occurrence of ef^{-1} in the word representing w. The *G-star graph* of w (with respect to σ) is the graph formed by superimposing the star graphs of xw for each element $x \in G$.

Let $\mathcal{W} = \{w_1, w_2, \ldots\}$ be the set of conjugacy classes of elements of F_n. Then $\Lambda = \mathbf{Z}^{\mathcal{W}}$ is an ordered abelian group with the lexicographical ordering, which we will denote by the symbol "$\preceq$".

If A and B are two subsets of $E(\Gamma)$, the *dot product* $A.B$ of A and B is the element of Λ whose ith coordinate $(A.B)_i$ is the number of edges of the G-star graph for w_i with one vertex in A and one vertex in B. Note that $xA.xB = A.B$ for any $x \in G$, and $A.(B + C) = A.B + A.C$, where "$+$" denotes disjoint union. The *absolute value* of A is the dot product $|A| = A.(E(\Gamma) - A)$, i.e. $|A|$ is the element of Λ whose ith coordinate $|A|_i$ is the number of edges of the G-star graph of w_i with exactly one endpoint in A. Note that $|A| = |E(\Gamma) - A|$. If A consists of a single edge e of Γ, then $|e|_i$ is the valence of the vertex of the G-star graph for w_i corresponding to e, i.e. it is the number of times one of the reduced paths xw_i crosses e in either direction. If we represent oriented edges of Γ by points in the plane, and edges of star graphs by arcs in the plane which miss these points, then we can separate A from $E(\Gamma) - A$ by a simple closed curve which intersects the G-star graph for w_i in exactly $|A|_i$ points.

The following elementary properties of dot product and absolute value follow easily from the definitions.

PROPOSITION 6.1. *If A and B are disjoint subsets of $E(\Gamma)$, then* $|A + B| = |A| + |B| - 2(A.B)$. □

PROPOSITION 6.2. *Let K be a subgroup of G, let A be a K-invariant subset of $E(\Gamma)$, and let e be an edge of Γ with* stab (e) *contained in K. Then* $(Ke).A = [K\colon \text{stab}\,(e)](e.A)$. □

B. Norm of a marked graph

If $\sigma = [s, \Gamma]$ is a marked G-graph, define its *norm* $\|\sigma\|$ to be the element of Λ whose ith coordinate is given by

$$\|\sigma\| = \frac{1}{2} \sum_{e \in E(\Gamma)} |e|_i.$$

Equivalently, the ith coordinate $\|\sigma\|_i$ is equal to the sum of the lengths of the cyclically reduced edge-paths representing the conjugacy classes xw_i for $x \in G$.

Note that the image of a reduced edge-path in Γ under a forest collapse is still reduced. In particular, the absolute value of an edge e of Γ which is not in the forest is unchanged by the forest collapse. From this it follows that collapsing an invariant forest in a marked G-graph strictly decreases the norm.

If we think of a marked G-graph σ as an action of F_n on a $\mathbf{Z}$-tree which is fixed by the subgroup G of Out (F_n), then the coordinate $\|\sigma\|_i$ is the (hyperbolic) length of w_i in the action, multiplied by the order of G. Results of Alperin and Bass [1] or Culler and Morgan [6] imply that an action of F_n on a $\mathbf{Z}$-tree is determined by its norm. In particular, the norm gives a total ordering to the set of marked G-graphs for a fixed G.

PROPOSITION 6.3. *The set of marked G-graphs is well-ordered by the norm.*

Proof. Let U be any set of marked graphs. We must find a least element of U. Let $U_0 = U$, and define subsets U_i of U, for $i \geq 1$, inductively as follows. Let $\ell_i = \min\{\|\sigma\|_i \mid \sigma \in U_{i-1}\}$. Then U_i is the subset of U_{i-1} consisting of elements σ with $\|\sigma\|_i = \ell_i$. We have

$$U = U_0 \supseteq U_1 \supseteq U_2 \cdots.$$

Define a function $f: F_n \to \mathbf{Z}$ by $f(w_i) = \ell_i$. The axioms for hyperbolic length functions (see [6] or [16]) are satisfied by f, so f is the length function of some action

of F_n on a $\mathbf{Z}$-tree [16]. Since $\ell_i \neq 0$ for all i, this action is free, and since f is constant on G-orbits, this action is in K_G; we will call the vertex of K_G corresponding to this action σ_0.

CLAIM. *A minimal free action of F_n on a $\mathbf{Z}$-tree is determined by the lengths of finitely many conjugacy classes (which depend on the action), i.e. there is a finite set of conjugacy classes such that no other action has the same lengths on those classes.*

Proof. Given any finite set of conjugacy classes $\mathscr{A} = \{a_1, \ldots, a_k\}$ in F_n, we can define an $\mathscr{A}$-norm on the set of free actions on $\mathbf{Z}$-trees by taking the $\mathscr{A}$-norm of an action to be the sum of the lengths of the conjugacy classes a_i in that action. By Proposition 6.2.5 of [7], there is a finite set of words $\mathscr{A}_0$ so that there is a unique rose ρ_0 of minimal $\mathscr{A}_0$-norm.

Fix a minimal free action τ. Then τ is in the ball B_N consisting of graphs of $\mathscr{A}_0$-norm less than N, for some N. We claim that B_N contains only finitely many marked graphs. By the Existence Theorem of [7], every rose ρ in B_N can be connected to ρ_0 by a Whitehead path (i.e. a path of elementary Whitehead moves) which reduces $\mathscr{A}_0$-norm at each step, so has length at most N; since there are only a finite number of Whitehead paths of a given length from ρ_0, this shows that B_N contains only a finite number of roses. If σ is any action in B_N, then collapsing a maximal tree in σ produces a rose with $\mathscr{A}_0$-norm smaller than the $\mathscr{A}_0$-norm of σ, so σ is in the star of some rose in B_N. Since the star of a rose contains only finitely many marked graphs, there are only a finite number of marked graphs in B_N.

The finitely many marked graphs in B_N can be distinguished by a finite set $\mathscr{B}$ of conjugacy classes of F_n. Therefore the lengths of conjugacy classes in $\mathscr{A}_0 \cup \mathscr{B}$ distinguish τ from any other action. □

Now choose N large enough so that σ_0 is determined by the lengths of $w_1, \ldots, w_N$ (i.e. $\{w_1, \ldots, w_N\}$ contains all the conjugacy classes which are needed to determine σ_0). Then $U_N = \{\sigma_0\}$, so σ_0 is an element of U which is smaller than any other element of U. □

We now compute the effect of blowing up and blowing down on the norm of a reduced G-graph.

Let $\sigma = [s, \Gamma]$ be any marked G-graph, and suppose that σ_{Ga} is obtained from σ by collapsing an edge-orbit Ga which is a forest in Γ. Since there are $[G\colon \text{stab}(a)]$ edges in Ga, all with the same absolute value, we have

$$\|\sigma_{Ga}\| = \|\sigma\| - \sum_{e \in Ga} |e| = \|\sigma\| - [G\colon \text{stab}(a)]|a|. \qquad (*)$$

If α is an ideal edge of Γ, then σ is obtained from $\sigma^{G\alpha}$ by collapsing the edge-orbit $Ge(\alpha)$; thus equation (*) gives that $\|\sigma^{G\alpha}\| = \|\sigma\| + [G\colon \text{stab}(\alpha)]|e(\alpha)|$. But $|e(\alpha)|$ in $\Gamma^{G\alpha}$ is equal to $|\alpha|$ in Γ, and stab $(\alpha) =$ stab $(e(\alpha))$, so

$$\|\sigma^{G\alpha}\| = \|\sigma\| + \sum_{\beta \in G\alpha} |\beta| = \|\sigma\| + [G\colon \text{stab}(\alpha)]|\alpha|. \qquad (**)$$

Combining equalities (*) and (**) gives

PROPOSITION 6.4. *Let (α, a) be an ideal pair in the reduced G-graph ρ, and suppose ρ' is obtained from ρ by the Whitehead move $(G\alpha, Ga)$. Then $\|\rho'\| = \|\rho\| + p(|\alpha| - |a|)$, where $p = [G\colon \text{stab}(\alpha)] = [G\colon \text{stab}(a)]$.* □

The quantity $[G\colon \text{stab}(\alpha)](|a| - |\alpha|)$ will be called the *reductivity* of (α, a), denoted red (α, a). The Whitehead move $(G\alpha, Ga)$ reduces the norm if and only if red $(\alpha, a) > 0$. We will be particularly interested in Whitehead moves which reduce the norm, which we call *reductive* Whitehead moves. The reductivity of an ideal edge α is the maximum, over all $a \in D(\alpha)$, of red (α, a). If α is invertible and $a \in D(\alpha)$, then $xa^{-1} \in D(\alpha^{-1})$ for a unique $x \in G$, and red $(\alpha, a) =$ red (α^{-1}, xa^{-1}).

The following proposition is an application of the Factorization Lemma 4.7, which will be needed in the proof of the contractibility of L_G.

PROPOSITION 6.5. *Let ρ' be a reduced marked G-graph, which is obtained from the reduced marked G-graph ρ by first blowing up the oriented ideal forest Φ of ρ, and then collapsing the maximal invariant forest $\mathcal{F}$ of ρ^{Φ}. If $\|\rho'\| < \|\rho\|$, then some ideal edge in Φ is reductive.*

Proof. The oriented ideal forest Φ gives a maximal invariant forest in Γ^{Φ} which is oriented naturally; thus the Factorization Lemma 4.7 gives an orientation on $\mathcal{F}$ and a bijection $f\colon \Phi \to E^+(\mathcal{F})$ such that $(\alpha, f(\alpha))$ is an ideal pair for each ideal edge α of Φ. Applying equations (*) and (**) repeatedly, we obtain

$$\|\rho'\| - \|\rho\| = \sum_{\alpha \in \Phi} |\alpha| - \sum_{a \in \mathcal{F}} |a| = \sum_{\alpha \in \Phi} (|\alpha| - |f(\alpha)|).$$

Since $\|\rho'\| - \|\rho\| < 0$, we must have $|\alpha| - |f(\alpha)| < 0$ for some $\alpha \in \Phi$, i.e. the Whitehead move $(G\alpha, Gf(\alpha))$ is reductive. □

§7. Combinatorial Lemmas: Finding new reductive ideal edges

In this section we prove two combinatorial lemmas which we will need in the proof that L_G is contractible. Each lemma shows that one of a small number of candidates is a reductive ideal edge. The method is to compare the reductivity of the candidates to the reductivity of ideal edges already known to be reductive, using two basic counting inequalities.

In previous sections we have written $G\alpha$ for the set of ideal edges of the form $x\alpha$ for $x \in G$. We abuse notation in this section by writing $G\alpha$ for $\bigcup G\alpha = \bigcup_{x \in G} x\alpha$ as well. It should be clear from the context which is meant in each case.

A. Intersection components and crossing number

Let α and β be two ideal edges. Let P denote the stabilizer of α, with index p in G, and let Q be the stabilizer of β, with index q in G. Choose double coset representatives $x_1 = 1, x_2, \ldots, x_k$ for $P\backslash G/Q$, i.e.

$$G = PQ + Px_2Q + \cdots + Px_kQ,$$

The intersections $\gamma = \alpha \cap G\beta$ and $\gamma' = \beta \cap G\alpha$ decompose as disjoint unions

$$\gamma = \gamma_1 + \cdots + \gamma_k$$

$$\gamma' = \gamma'_1 + \cdots + \gamma'_k$$

where $\gamma_i = \alpha \cap Px_i\beta = P(\alpha \cap x_i\beta)$ and $\gamma'_i = \beta \cap Qx_i^{-1}\alpha = Q(x_i^{-1}\alpha \cap \beta)$ (see Figure 8). The γ_i are the *intersection components of* α *with* β.

Note that $G\gamma = G\gamma'$ and $G\gamma_i = G\gamma'_i$. The map $e \mapsto x_i^{-1}e$ sending $\alpha \cap x_i\beta$ to $x_i^{-1}\alpha \cap \beta$ induces a bijection between the P-orbits in γ_i and the Q-orbits in γ'_i.

DEFINITION. The *crossing number* $N(G\alpha, G\beta)$ is the number of non-empty intersection components $\gamma_i = P(\alpha \cap x_i\beta)$.

If $N(G\alpha, G\beta) = 0$, the orbits $G\alpha$ and $G\beta$ are disjoint, hence pre-compatible. If $G\alpha$ and $G\beta$ are not disjoint, we replace β by some translate $x\beta$ so that $\alpha \cap \beta \neq \varnothing$, so $\gamma_1 \neq \varnothing$. If $G\alpha$ and $G\beta$ cross and $N(G\alpha, G\beta) = 1$, we say $G\alpha$ and $G\beta$ cross *simply*; in this case $\gamma = \gamma_1 = P(\alpha \cap \beta)$ and $\gamma' = \gamma'_1 = Q(\alpha \cap \beta)$.

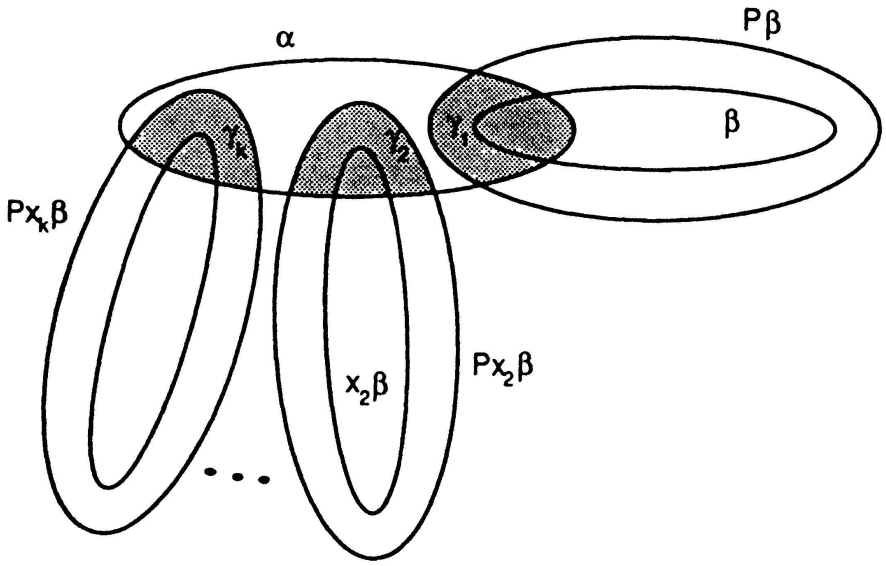

Figure 8

B. Basic inequalities

LEMMA 7.1. *If* $G\alpha$ *and* $G\beta$ *cross simply, and* stab $(\alpha) \leq$ stab (β), *then*

$$p|\alpha \cap \beta| + q|\beta \cup Q\alpha| \preceq p|\alpha| + q|\beta|.$$

Proof. Recall that $\gamma = \alpha \cap \beta$, let $A = \alpha - \gamma$ and $B = \beta + QA$, and let $E = E(\Gamma)$ (Figure 9).

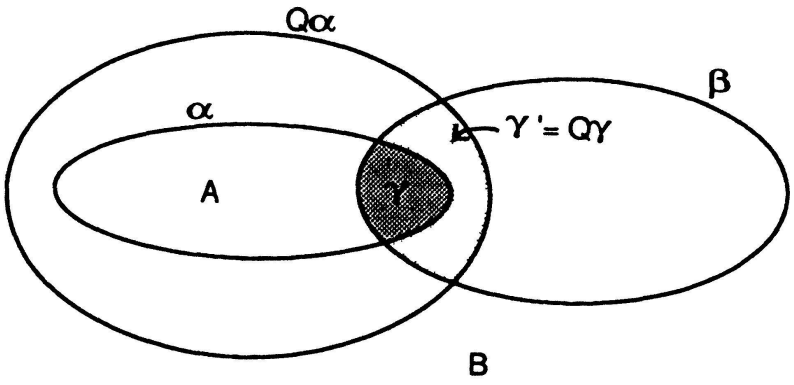

Figure 9

By Proposition 6.1,

$$|\alpha| - |\gamma| = |A| - 2A.\gamma$$

and $|\beta| - |B| = |E - \beta| - |E - B|$

$$= |QA| - 2(QA).(E - B).$$

Thus the inequality in the statement of the lemma reduces to showing $p|A| + q|QA| - 2pA.\gamma - 2q(QA).(E - B) \geq 0$. It suffices to show this on each orbit Pt of A separately, i.e.

$$pPt.(E - A) + qQt.(E - QA) - 2pPt.\gamma - 2qQt.(E - B) \geq 0$$

for each $t \in A$.

By Proposition 6.2, we have $pPt.(E - A) = p[P\colon \text{stab}(t)]t.(E - A) = [G\colon \text{stab}(t)]t.(E - A)$. Similarly, $[G\colon \text{stab}(t)]$ is a factor in each term of the inequality, so we can cancel to get the equivalent inequality

$$t.(E - A) + t.(E - QA) - 2t.\gamma - 2t.(E - B) \geq 0.$$

Since $t.(E - A) \geq t.(E - QA)$, it suffices to show

$$t.(E - QA) - t.\gamma - t.(E - B) \geq 0.$$

But this is immediate since QA, γ and $E - B$ are disjoint subsets of E. □

LEMMA 7.2. *If $G\alpha$ and $G\beta$ cross, then for all i,*

$$p|\alpha - \gamma_i| + q|\beta - \gamma_i'| \leq p|\alpha| + q|\beta|.$$

Proof. Let $E = E(\Gamma)$, $A = \alpha - \gamma_i$ and $B = \beta - \gamma_i'$ (Figure 10). By Proposition 6.1 we have

$$|\alpha| = |\gamma_i| + |A| - 2\gamma_i.A,$$

$$|\beta| = |\gamma_i'| + |B| - 2\gamma_i'.B.$$

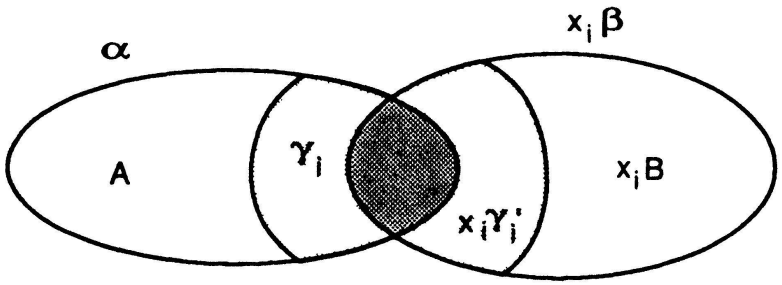

Figure 10

Thus the inequality in the statement of the lemma reduces to showing

$$p|\gamma_i| + q|\gamma_i'| - 2\gamma_i.A - 2\gamma_i'.B \geq 0.$$

We do this separately on the orbits in γ_i and γ_i' as follows. For each P-orbit in γ_i, choose an edge t representing it which lies in $\alpha \cap x_i\beta$; then the associated Q-orbit in γ_i' is represented by $t' = x_i^{-1}t$. We will show

$$pPt.(E - \gamma_i) + qQt'.(E - \gamma_i') - 2pPt.A - 2qQt'.B \geq 0$$

By Proposition 6.2,

$$pPt.(E - \gamma_i) = p[P\text{: stab }(t)]t.(E - \gamma_i) = [G\text{: stab }(t)]t.(E - \gamma_i).$$

Similarly, $[G\text{: stab }(t)] = [G\text{: stab }(t')]$ is a factor in each term of the inequality, so we can cancel to obtain the equivalent inequality

$$t.(E - \gamma_i) + t'.(E - \gamma_i') - 2t.A - 2t'.B \geq 0.$$

Now $t.(E - \gamma_i) \geq t.(E - G\gamma_i)$ and $t'.(E - \gamma_i') \geq t'.(E - G\gamma_i') = t'.(E - G\gamma_i) = t.(E - G\gamma_i)$. Thus it suffices to show

$$t.(E - G\gamma_i) - t.A - t.x_iB \geq 0.$$

But this is immediate since $G\gamma_i$, A and x_iB are disjoint subsets of E. □

C. The Pushing and Shrinking Lemmas

PROPOSITION 7.3 (*Pushing Lemma*). *Let (μ, m) be a maximally reductive ideal pair of a reduced marked G-graph, and let α be a reductive ideal edge which simply crosses μ, with $m \in \alpha \cap \mu$. Then one of the sets $\mu - \alpha$ or $\alpha \cup \text{stab }(\alpha)\mu$ is a reductive ideal edge.*

Proof. Set $P = \text{stab }(\alpha)$, and fix $a \in D(\alpha)$, with $\text{red }(\alpha, a) > 0$. We divide the proof into cases according to the position of a^{-1} and m^{-1}, illustrated in Figure 11. Since $m \in \alpha$, we have $\text{stab }(\mu) \leq \text{stab }(\alpha)$.

Case 1. $a^{-1} \notin G\mu$.
An application of inequality 7.1 shows that

$$\text{red }(\alpha \cup P\mu, a) + \text{red }(\alpha \cap \mu, m) \geq \text{red }(\alpha, a) + \text{red }(\mu, m).$$

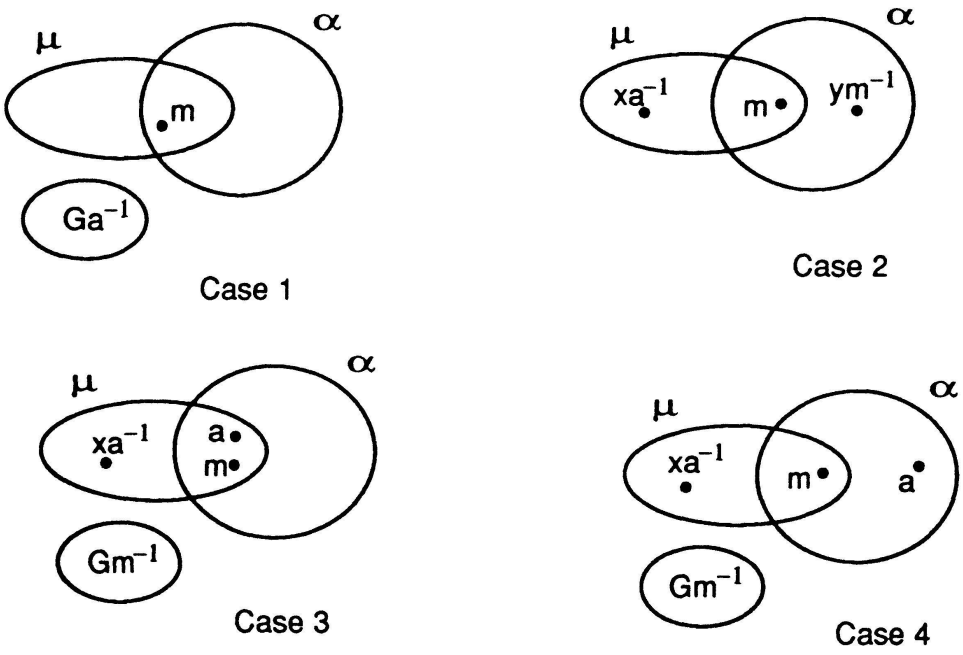

Figure 11

Since (μ, m) is maximally reductive and (α, a) is reductive, both $(\alpha \cup P\mu, a)$ and $(\alpha \cap \mu, m)$ must be reductive.

To complete the proof, we must check that $(\alpha \cup P\mu, a)$ is an ideal pair; in particular, we must check condition (i) of the definition of ideal edge, which says that card $(E_v - (\alpha \cup P\mu)) \geq 2$. Since a is an edge of a reduced G-graph, a must be either elliptic or bent hyperbolic. If a is bent hyperbolic, some translate xa^{-1} is in E_v, but is not in $G\alpha$ or in $G\mu$. If xa^{-1} is the only edge of $E_v - (\alpha \cup P\mu)$, then $|\alpha \cup P\mu| = |a^{-1}| = |a|$, contradicting reductivity of $(\alpha \cup P\mu, a)$. If a is elliptic, then P is a proper subgroup of stab (v). Choose an element $x \in \text{stab}(v) - P$, an edge $b \in \alpha - P\mu$ and an edge $c \in \mu - \alpha$. Then xb and xc are both in $E_v - (\alpha \cup P\mu)$.

In what follows, we omit similar cardinality checks.

For the remaining cases we may assume that $a^{-1} \in G\mu$. Choose x with $xa^{-1} \in \mu$. Since stab (xa^{-1}) is a conjugate of stab $(\alpha) = P$ and a subgroup of stab $(\mu) \leq P$, we must have stab $(\mu) = P$.

Case 2. $a^{-1} \in G\mu$ and $m^{-1} \in G\alpha$.

Choose $y \in G$ with $ym^{-1} \in \alpha$. Then $(\alpha - \mu, ym^{-1})$ and $(\mu - \alpha, xa^{-1})$ are ideal pairs and inequality 7.2 shows that

$$\text{red}\,(\alpha - \mu, ym^{-1}) + \text{red}\,(\mu - \alpha, xa^{-1}) \geq \text{red}\,(\alpha, a) + \text{red}\,(\mu, m).$$

Since (μ, m) is maximally reductive and (α, a) is reductive, both $(\alpha - \mu, ym^{-1})$ and $(\mu - \alpha, xa^{-1})$ must be reductive.

Case 3. $a^{-1} \in G\mu$, $m^{-1} \notin G\alpha$ and $a \in \mu$.
Then $(\alpha \cup \mu, m)$ and $(\alpha \cap \mu, a)$ are ideal pairs and inequality 7.1 shows that

$$\text{red}\,(\alpha \cup \mu, m) + \text{red}\,(\alpha \cap \mu, a) \succeq \text{red}\,(\alpha, a) + \text{red}\,(\mu, m).$$

As above, both $(\alpha \cup \mu, m)$ and $(\alpha \cap \mu, a)$ must be reductive.

Case 4. $a^{-1} \in G\mu$, $m^{-1} \notin G\alpha$ and $a \notin \mu$.
Then $xa^{-1} \in D(\mu)$, so $|m| \succeq |a|$. Thus

$$\text{red}\,(\alpha \cup \mu, m) + \text{red}\,(\alpha \cap \mu, m) \succeq \text{red}\,(\alpha, a) + \text{red}\,(\mu, m).$$

As above, both $(\alpha \cup \mu, m)$ and $(\alpha \cap \mu, m)$ must be reductive. □

PROPOSITION 7.4 (*Shrinking Lemma*). *Let (μ, m) be a maximally reductive ideal pair of a reduced marked G-graph, and let α be an ideal edge which crosses μ. Let $\gamma_{i_1}, \ldots, \gamma_{i_k}$ be the intersection components of α with μ which contain no translate of m, and let $\beta = \alpha - \bigcup \gamma_{i_j}$. Then β or one of the sets γ_{i_j} is a reductive ideal edge.*

Proof. Let $P = \text{stab}\,(\alpha)$, with index p in G, and $Q = \text{stab}\,(\mu)$ with index q in G.

Suppose $m \in G\alpha$, i.e. some intersection component $\gamma_i = \alpha \cap Px_i\mu$ contains a translate xm of m. Since α is an ideal edge, $G(xm) \cap \alpha = P(xm)$. Since γ_i is a union of P-orbits, $P(xm)$ is in γ_i. Thus at most one intersection component γ_i contains a translate of m. Replacing μ by $x\mu$, we may assume this intersection component is γ_1, and $m \in \gamma_1$.

Let ϵ be the union of all intersection components γ_i which contain neither a nor m, and ϵ' the corresponding union of γ_i'. If $\epsilon \neq \varnothing$, repeated use of the inequality 7.2 gives

$$p|\alpha - \epsilon| + q|\mu - \epsilon'| \preceq p|\alpha| + q|\mu|.$$

From this it follows that

$$\text{red}\,(\alpha - \epsilon, a) + \text{red}\,(\mu - \epsilon', m) \succeq \text{red}\,(\alpha, a) + (\mu, m).$$

Since (μ, m) is maximally reductive, we have $\text{red}\,(\alpha - \epsilon, a) \succ 0$, i.e. $\alpha - \epsilon$ is reductive.

If a is not in γ_i for any i, or if a and m are both in γ_1, then $\alpha - \epsilon = \beta$ and we are done. By renumbering, we may therefore assume that either $m \in \gamma_1$ and $a \in \gamma_2$ or $a \in \gamma_1$ and $m \notin G\alpha$. We may also replace α by $\alpha - \epsilon$; then either $N(G\alpha, G\mu) = 1$, $a \in \gamma = \gamma_1$ and $m \notin G\alpha$ or $N(G\alpha, G\mu) = 2$, $m \in \gamma_1$ and $\alpha \in \gamma_2$.

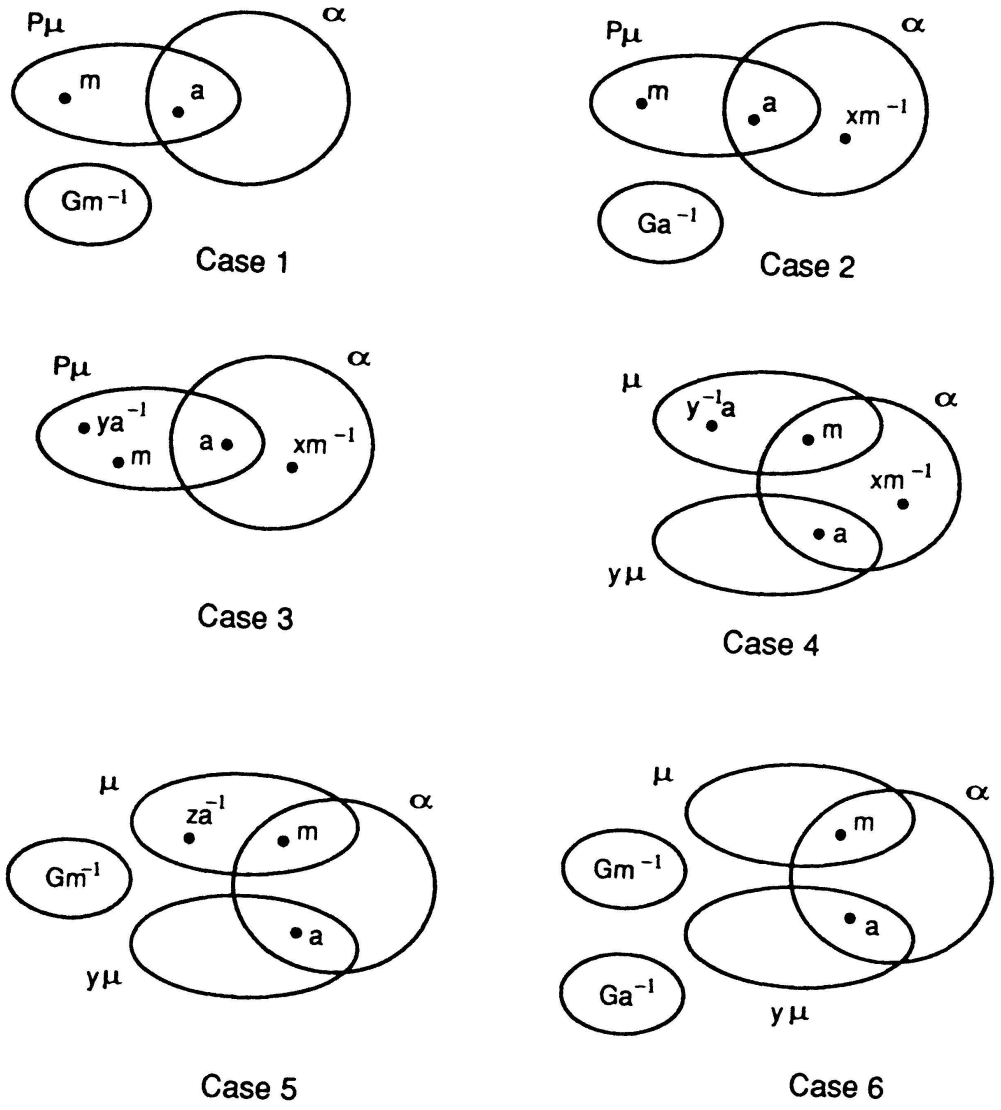

Figure 12

We divide the proof into cases depending on $N(G\alpha, G\mu)$ and the location of a^{-1} and m^{-1} (see Figure 12).

Case 1. $N(G\alpha, G\mu) = 1, m^{-1} \notin G\alpha$.
Then (γ, a) and $(\mu \cup Q\alpha, m)$ are ideal pairs, and the inequality 7.1 gives

$$\text{red}\,(\gamma, a) + \text{red}\,(\mu \cup Q\alpha, m) \geq \text{red}\,(\alpha, a) + \text{red}\,(\mu, m).$$

Since (μ, m) is maximally reductive, (γ, a) is reductive.

Case 2. $N(G\alpha, G\mu) = 1, m^{-1} \in G\alpha$ but $a^{-1} \notin G\mu$.
Fix x with $xm^{-1} \in \alpha$. Then $P = Q$, and $(\alpha - \gamma, xm^{-1})$ and $(\mu - \gamma', m)$ are ideal

pairs. Since $a \in D(\mu)$ and (μ, m) is maximally reductive, we have $|m| \geq |a|$; this together with inequality 7.2 gives

$$\begin{aligned}\operatorname{red}(\alpha - \gamma, xm^{-1}) + \operatorname{red}(\mu - \gamma', m) &\geq p(|m| - |\alpha| + |m| - |\mu|) \\ &\geq \operatorname{red}(\alpha, a) + \operatorname{red}(\mu, m).\end{aligned}$$

Maximal reductivity of (μ, m) now shows that $(\alpha - \gamma, xm^{-1})$ is reductive.

Case 3. $N(G\alpha, G\mu) = 1$, $m^{-1} \in G\alpha$ and $a^{-1} \in G\mu$.

Again $P = Q$. Choose $y \in G$ with $ya^{-1} \in \mu$. Then $(\alpha - \gamma, xm^{-1})$ and $(\mu - \gamma', ya^{-1})$ are ideal edges. Inequality 7.2 gives

$$\operatorname{red}(\alpha - \gamma, xm^{-1}) + \operatorname{red}(\mu - \gamma', ya^{-1}) \geq \operatorname{red}(\alpha, a) + \operatorname{red}(\mu, m).$$

Thus $(\alpha - \gamma, xm^{-1})$ is reductive.

In the next three cases, we have $N(G\alpha, G\mu) = 2$. Since $m \in \alpha$, we have $P = Q$ and $\gamma_1 = \gamma_1'$. Fix $y \in G$ so that $\gamma_2 = \alpha \cap y\mu$ and $\gamma_2' = y^{-1}\gamma_2$.

Case 4. $N(G\alpha, G\mu) = 2$, $m^{-1} \in G\alpha$.

Choose $x \in G$ with $xm^{-1} \in \alpha$. The inequality 7.2 gives

$$\operatorname{red}(\alpha - \gamma_2, xm^{-1}) + \operatorname{red}(\mu - \gamma_1, y^{-1}a) \geq \operatorname{red}(\alpha, a) + \operatorname{red}(\mu, m).$$

Inequality 7.1 gives

$$\begin{aligned}&\operatorname{red}(\gamma_2, a) + \operatorname{red}((a - \gamma_1) \cup y(\mu - \gamma_1), xm^{-1}) \\ &\quad \geq \operatorname{red}(\alpha - \gamma_1, xm^{-1}) + \operatorname{red}(\mu - \gamma_1, y^{-1}a).\end{aligned}$$

The pairs (γ_2, a) and $((a - \gamma_1) \cup y(\mu - \gamma_1), xm^{-1})$ are ideal pairs. Since (μ, m) is maximally reductive and (α, a) is reductive, these two inequalities together show that (γ_2, a) is reductive.

Case 5. $N(G\alpha, G\mu) = 2$, $m^{-1} \notin G\alpha$, $a^{-1} \in G\mu$.

Choose $z \in G$ with $za^{-1} \in \mu$. Then $(\alpha - \gamma_2, m)$ and $(\mu - \gamma_2', za^{-1})$ are ideal pairs, and inequality 7.2 gives

$$\operatorname{red}(\alpha - \gamma_2, m) + \operatorname{red}(\mu - \gamma_2', za^{-1}) \geq \operatorname{red}(\alpha, a) + \operatorname{red}(\mu, m),$$

so $\alpha - \gamma_2$ is reductive.

Case 6. $N(G\alpha, G\mu) = 2, m^{-1} \notin G\alpha, a^{-1} \notin G\mu$.

In this case, $(\alpha - \gamma_2, m)$ and $(\mu - \gamma_2', m)$ are ideal pairs and $|m| \geq |a|$ since $y^{-1}a \in D(\mu)$. An application of inequality 7.2 gives

$$\text{red}\,(\alpha - \gamma_2, m) + \text{red}\,(\mu - \gamma_2', m) \geq \text{red}\,(\alpha, a) + \text{red}\,(\mu, m),$$

so that $\alpha - \gamma_2$ is reductive. □

§8. Proof of Contractibility

THEOREM 8.1. *The complex L_G is contractible.*

Proof. By Proposition 6.3, the set of all reduced marked G-graphs is well-ordered by the norm. The complex L_G is the union of the stars of the reduced marked G-graphs ρ. Set

$$L_{\prec\rho} = \bigcup_{\|\rho'\| \prec \|\rho\|} \text{st}\,(\rho')$$

and let $S_\rho = \text{st}\,(\rho) \cap L_{\prec\rho}$. We will show that S_ρ is contractible whenever it is non-empty; an easy transfinite induction argument then proves the statement "For every ρ, all components of $L_{\prec\rho}$ are contractible." Since L_G is connected (Proposition 5.11), it must then be contractible.

Fix a reduced marked G-graph $\rho = [s, \Gamma]$. By Proposition 5.9, the star of ρ in L_G is identified with the geometric realization of the poset of all ideal forests of ρ. We will prove that S_ρ is contractible by repeated retraction using the Poset Lemma (3.4). The proof is contained in the sequence of lemmas below.

For every invariant set $\mathscr{C}$ of ideal edges of ρ, let $S(\mathscr{C})$ denote the subcomplex of st (ρ) spanned by the ideal forests of ρ, all of whose edges are in $\mathscr{C}$. Let $\mathscr{C}^{\pm}$ be the set obtained by adjoining to $\mathscr{C}$ the inverses of its invertible elements, and let $\mathscr{R}$ be the set of all reductive ideal edges of ρ. Notice that $\mathscr{R} = \mathscr{R}^{\pm}$.

Let μ be a maximally reductive ideal edge of ρ, and choose $m \in \mu$ so that red (μ, m) is maximal. Define

$$\mathscr{C}_0 = \{\alpha \in \mathscr{R} \mid \alpha \text{ is compatible with } \mu\},$$

and

$$\mathscr{C}_1 = \mathscr{C}_0 \cup \{\alpha \in \mathscr{R} \mid \alpha \text{ is invertible}\}$$
$$\cup \{\alpha \in \mathscr{R} \mid m \in G\alpha \text{ and } N(G\alpha, G\mu) = 1\}.$$

LEMMA 8.2. *S_ρ deformation retracts onto $S(\mathscr{R})$.*

Proof. By Proposition 6.5, we can identify S_ρ with the geometric realization of the subposet consisting of ideal forests of ρ which contain a reductive ideal edge. The map which sends an ideal forest to the subforest consisting of its reductive ideal edges is thus a poset map. The lemma follows by applying the Poset Lemma (3.4). □

LEMMA 8.3. *Let $\mathscr{C} = \mathscr{C}^{\pm}$ be an invariant subset of $\mathscr{R}$ which contains $\mathscr{C}_1$. Then $S(\mathscr{C})$ deformation retracts to $S(\mathscr{C}_1)$. In particular, $S(\mathscr{R})$ deformation retracts onto $S(\mathscr{C}_1)$.*

Proof. The proof is by induction on the cardinality card $(\mathscr{C} - \mathscr{C}_1)$. If this cardinality is greater than 0, choose $\alpha \in \mathscr{C} - \mathscr{C}_1$ such that

1. card $(\alpha \cap G\mu)$ is minimal;
2. α is minimal with respect to property (1), i.e. if $\alpha' \in \mathscr{C} - \mathscr{C}_1$, $\alpha' \subset \alpha$, and $\alpha' \cap G\mu = \alpha \cap G\mu$, then $\alpha' = \alpha$.

Since $\alpha \in \mathscr{C} - \mathscr{C}_1$, α is not invertible, and either $N(G\alpha, G\mu) \geq 2$ or $N(G\alpha, G\mu) = 1$ and $m \notin G\alpha$. By the Shrinking Lemma (7.4) and with its notation, one of the sets $\gamma_1, \ldots, \gamma_k$ or $\alpha - (\gamma_1 + \cdots + \gamma_k)$ is a reductive ideal edge, which we call α_0. If $N(G\alpha, G\mu) = 1$ or $\alpha_0 = \gamma_i$, then α_0 is compatible with μ; otherwise, $N(G\alpha_0, G\mu) = 1$ and $m \in G\alpha_0$. In either case, $\alpha_0 \in \mathscr{C}_1$.

CLAIM. *For every $\beta \in \mathscr{C}$, if $G\beta$ is pre-compatible with $G\alpha$, then $G\beta$ is pre-compatible with $G\alpha_0$.*

Proof. If $G\alpha$ is included in $G\beta$ or is disjoint from it, then the same is true for $G\alpha_0$. The case when $G\beta^{-1}$ is included in $G\alpha$ does not occur because $G\alpha$ is not invertible (Corollary 5.2). There only remains the case when $G\beta$ is included in $G\alpha$. We may safely assume that β is a subset of α, so that card $(\beta \cap G\mu) \leq$ card $(\alpha \cap G\mu)$. If $\beta \notin \mathscr{C}_1$, then, by the first condition on our choice of α, these cardinalities must be equal, and then by the second condition, we must have $\beta = \alpha$. If $\beta \in \mathscr{C}_1$ then, since β cannot be invertible, we have either $\beta \in \mathscr{C}_0$ or $m \in G\beta$ and $N(G\beta, G\mu) = 1$. The intersection $\beta \cap G\mu$ is empty in the former, and is a subset of γ_0 in the latter case. So β is a subset of $\alpha - (\gamma_1 + \cdots + \gamma_k)$ and thus compatible with every choice of α_0. □

The claim shows that if Φ is an ideal forest which contains α, then $\Phi \cup G\alpha_0$ is also an ideal forest. Thus there is a well-defined poset map f from $S(\mathscr{C})$ to itself which sends an ideal forest Φ to $\Phi \cup G\alpha_0$ if Φ contains α, and to itself if Φ does not contain α. Since $f(\Phi) \supseteq \Phi$, the image $f(S(\mathscr{C}))$ is a deformation retract of $S(\mathscr{C})$. The

map g from $f(S(\mathscr{C}))$ to itself sending a forest Ψ to $\Psi - G\alpha$ if Ψ contains α, and to itself if Ψ does not contain α is then a well-defined poset map with image $S(\mathscr{C} - \{G\alpha\})$. Since $g(\Psi) \subseteq \Psi$, $S(\mathscr{C} - \{G\alpha\})$ is a deformation retract of the image of f, and hence of $S(\mathscr{C})$. Since $\alpha_0 \in \mathscr{C}_1$, card $(\mathscr{C} - \{G\alpha\} - \mathscr{C}_1) <$ card $(\mathscr{C} - \mathscr{C}_1)$, so $S(\mathscr{C} - \{G\alpha\})$ deformation retracts to $S(\mathscr{C}_1)$ by induction. □

LEMMA 8.4. *Let $\mathscr{C} = \mathscr{C}^{\pm}$ be an invariant subset of $\mathscr{C}_1$ which contains $\mathscr{C}_0$. Then $S(\mathscr{C})$ deformation retracts to $S(\mathscr{C}_0)$. In particular, $S(\mathscr{C}_1)$ deformation retracts to $S(\mathscr{C}_0)$.*

Proof. This proof is similar to the proof of the previous lemma, but uses the Pushing Lemma (7.3) instead of the Shrinking Lemma. Again, the proof is by induction, now on the cardinality card $(\mathscr{C} - \mathscr{C}_0)$. If this cardinality is greater than 0, choose $\alpha \in \mathscr{C} - \mathscr{C}_0$ such that $m \in \alpha$ and

1. card $(\alpha \cap G\mu)$ is maximal;
2. α is maximal with respect to property (1), i.e. if $\alpha' \in \mathscr{C} - \mathscr{C}_0$, $\alpha' \supset \alpha$, and $\alpha' \cap G\mu = \alpha \cap G\mu$, then $\alpha' = \alpha$.

If an ideal edge belongs to $\mathscr{C} - \mathscr{C}_0$, then a translate of it or of its inverse contains m; so α is well-defined. Since $G\alpha$ and $G\mu$ cross simply (for α invertible this is automatic), the Pushing Lemma applies. Thus one of the sets $\mu - \alpha$ or $\alpha \cup P\mu$ is a reductive ideal edge, where P is the stabilizer of α. Call this edge α_0 and note that $\alpha_0 \in \mathscr{C}_0$.

CLAIM. *For every $\beta \in \mathscr{C}$, if $G\beta$ is pre-compatible with $G\alpha$, then $G\beta$ is pre-compatible with $G\alpha_0$.*

Proof. If $G\beta$ is included in $G\alpha$, then $G\beta$ is included in $G(\alpha \cup P\mu)$ and is disjoint from $G(\mu - \alpha)$, so is compatible with either possibility for $G\alpha_0$. The same argument settles also the case when $G\beta^{-1}$ is included in $G\alpha$. If $G\alpha$ is included in $G\beta$, then, assuming without loss of generality that α is a subset of β, we have card $(\beta \cap G\mu) \geq$ card $(\alpha \cap G\mu)$. If $\beta \notin \mathscr{C}_0$, then by the first condition on our choice of α, these cardinalities must be equal, and then by the second condition, we must have $\beta = \alpha$. If $\beta \in \mathscr{C}_0$, then β must contain μ, so $G\alpha_0$ is included in $G\beta$.

Finally, suppose $G\alpha$ and $G\beta$ are disjoint. We may assume β is not invertible; otherwise $G\alpha$ is included in $G\beta^{-1}$ and we are in the previous case. If $\beta \notin \mathscr{C}_0$, then $m \in G\beta$ and so $\alpha \cap G\beta \neq \varnothing$ – a contradiction. So $G\beta$ is compatible with $G\mu$. If $G\beta$ is disjoint from $G\mu$ then it is disjoint from either possibility for $G\alpha_0$. If $G\beta$ is included in $G\mu$ then it is included in $G(\mu - \alpha)$ and so in $G(\alpha \cup P\mu)$ too. The possibility that $G\mu$ is included in $G\beta$ does not occur because $G\beta$ and $G\alpha$ are disjoint. □

The claim shows that if Φ is an ideal forest which contains α, then $\Phi \cup (G\alpha_0)^{\pm}$ is also an ideal forest. Thus there is a well-defined poset map f from $S(\mathscr{C})$ to itself

which sends an ideal forest Φ to $\Phi \cup (G\alpha_0)^{\pm}$ if Φ contains α, and to itself if Φ does not contain α. Since $f(\Phi) \supseteq \Phi$, the image $f(S(\mathscr{C}))$ is a deformation retract of $S(\mathscr{C})$. The map g from $f(S(\mathscr{C}))$ to itself sending a forest Ψ to $\Psi - (G\alpha)^{\pm}$ if Ψ contains α, and to itself if Ψ does not contain α is then a well-defined poset map with image $S(\mathscr{C} - \{G\alpha\}^{\pm})$. Since $g(\Psi) \subseteq \Psi$, $S(\mathscr{C} - \{G\alpha\}^{\pm})$ is a deformation retract of the image of f, and hence of $S(\mathscr{C})$. Since $\alpha_0 \in \mathscr{C}_0$, $\operatorname{card}(\mathscr{C} - \{G\alpha\}^{\pm} - \mathscr{C}_0) < \operatorname{card}(\mathscr{C} - \mathscr{C}_0)$, so $S(\mathscr{C} - \{G\alpha\}^{\pm})$ deformation retracts to $S(\mathscr{C}_0)$ by induction. □

The following lemma completes the proof of the theorem.

LEMMA 8.5. *$S(C_0)$ is contractible.*

Proof. The poset map $\Phi \mapsto \Phi \cup \mu^{\pm} \mapsto \mu^{\pm}$ gives a deformation retraction of $S(\mathscr{C}_0)$ to a point, by the Poset Lemma (3.4). □

§9. Virtual cohomological dimension

A. Dimension of L_G

In this section we compute the dimension of the complex L_G, thereby giving an upper bound to the virtual cohomological dimension of $C(G)$. We first note that any maximal simplex in L_G contains a reduced marked G-graph ρ, which is a minimal element in the poset ordering on the vertices of L_G, and a maximal element σ. The vertex ρ is obtained from σ by collapsing a maximal invariant forest in σ, and σ is obtained from ρ by blowing up a maximal oriented ideal forest in ρ. The dimension of the simplex is the number of edge-orbits in the maximal (oriented) forest of σ or, equivalently, the number of ideal edge-orbits in the maximal oriented ideal forest of ρ. It follows from the Factorization Lemma (4.7) that any two maximal oriented invariant forests in σ contain the same number of edge-orbits (and, in fact, the same number of edges). We show below (Proposition 9.3) that any two maximal oriented ideal forests in ρ have the same number of elements. Since any two reduced marked G-graphs can be connected by a sequence of blow-ups and collapses, it follows that every maximal simplex in L_G has the same dimension. We compute this dimension in terms of the quotient graph of groups determined by a fixed reduced marked G-graph ρ.

B. The quotient graph of groups

Fix a reduced marked G-graph $\rho = [s, \Gamma]$, with quotient map $q : \Gamma \to \bar{\Gamma} = G\backslash\Gamma$. The action of G on Γ gives rise to a graph of groups $\mathscr{G} = \{\bar{\Gamma}, \{\mathscr{G}_e\}, \{\mathscr{G}_v\}\}$ based on

$\bar{\Gamma}$, where the edge and vertex groups are stabilizers of edges and vertices in Γ, and the injection $\mathcal{G}_e \to \mathcal{G}_{\tau(e)}$ is an inclusion or is conjugation by an element of G followed by an inclusion (see [17] for details). Write $\mathcal{G}(e)$ for the image of $\mathcal{G}_e$ in $\mathcal{G}_{\tau(e)}$.

The following elementary lemmas show how to recognize ideal edges and ideal forests of ρ by their images in $\mathcal{G}$.

If a and e are edges of $\bar{\Gamma}$ terminating at v, write $e \leq a$ if $\mathcal{G}(e)$ is conjugate to a subgroup of $\mathcal{G}(a)$ by an element of $\mathcal{G}_v$, and $e < a$ if $\mathcal{G}(e)$ is conjugate to a proper subgroup of $\mathcal{G}(a)$. Note that we may have $e \leq a$ and $a \leq e$ for distinct edges a and e, so this fails to be a partial ordering.

LEMMA 9.1. *Let v be a vertex of $\bar{\Gamma}$, E_v the set of edges of $\bar{\Gamma}$ terminating at v, and $\alpha \subset E_v$. Then α is the image of an ideal edge of Γ if and only if*

(i) $\text{card}(\alpha) \geq 2$;

(ii) *there is an edge $a \in \alpha$ such that $a^{-1} \notin \alpha$ and, for all $e \in \alpha$, $e \leq a$;*

(iii) *if $\mathcal{G}(a) = \mathcal{G}_v$ for some a satisfying condition* (ii), *then* $\text{card}(E_v - \alpha) \geq 2$. □

The image of an ideal edge of Γ will be called an *ideal edge* of $\bar{\Gamma}$. All translates of an ideal edge of Γ project onto the same ideal edge of $\bar{\Gamma}$, but the converse is not true: an ideal edge of $\bar{\Gamma}$ does not determine an ideal edge-orbit of Γ. In Example 5.3, there is only one ideal edge of $\bar{\Gamma}$, but two ideal edge-orbits in Γ.

If α is the image of $\tilde{\alpha}$, define $D(\alpha)$ to be the image of $D(\tilde{\alpha})$; this is precisely the set of edges in α which satisfy condition (ii) above. If $a, b \in D(\alpha)$, then $\mathcal{G}(a)$ is conjugate to $\mathcal{G}(b)$ in $\mathcal{G}_v$; we let $P(\alpha)$ denote the corresponding conjugacy class of subgroups of $\mathcal{G}_v$. If $\tilde{a}$ and $\tilde{\beta}$ are inverse ideal edges in Γ, with images α and β in $\bar{\Gamma}$, then $P(\alpha) = P(\beta) = \mathcal{G}_v$ and $\alpha \cup \beta = E_v$; in this case we say α and β are *inverse*.

LEMMA 9.2. *A set Φ of ideal edges of $\bar{\Gamma}$ is the image of an oriented ideal forest of Γ if and only if, for each α and β in Φ, we have $\alpha \subseteq \beta$, $\beta \subseteq \alpha$, or $\alpha \cap \beta = \varnothing$ and α and β are not inverse.* □

The image of an oriented ideal forest of Γ will be called an *oriented ideal forest* of $\bar{\Gamma}$.

We compute the dimension of L_G by computing the number of elements in an oriented ideal forest of $\bar{\Gamma}$. This number is computed in terms of the numbers of particular types of edges, vertices and conjugacy classes of groups in $\mathcal{G}$, which we now specify.

DEFINITION. An edge e of $\bar{\Gamma}$ is *active* if it belongs to an ideal edge of $\bar{\Gamma}$. An edge e is *critical* if e is active and $e \not< a$ for any a, i.e. $\mathcal{G}(e)$ is not conjugate in $\mathcal{G}_{\tau(e)}$ to a proper subgroup of $\mathcal{G}(a)$ for any a terminating at $\tau(e)$.

A vertex v of $\bar{\Gamma}$ is *critical* if there is a critical edge e terminating at v with $\mathcal{G}(e) = \mathcal{G}_v$.

A conjugacy class P of subgroups of a vertex group $\mathcal{G}_v$ is *active* or *critical* if $\mathcal{G}(e) \in P$ for some active or critical edge e terminating at v.

A conjugacy class P of $\mathcal{G}_v$ is a *rose class* if $\{e \mid \mathcal{G}(e) \in P\}$ is a rose.

PROPOSITION 9.3. *Let G be a finite subgroup of* Out (F_n), *and let $[s, \Gamma]$ be a reduced marked G-graph, with quotient graph of groups $\mathcal{G} = \{\bar{\Gamma}, \{\mathcal{G}_e\}, \{\mathcal{G}_v\}\}$. Then the number of elements in any maximal oriented ideal forest in $\bar{\Gamma}$ is equal to*

$$\text{card}\,\{\textit{active edges of } \bar{\Gamma}\} - \text{card}\,\{\textit{critical vertices of } \bar{\Gamma}\}$$
$$-\text{card}\,\{\textit{critical conjugacy classes of } \mathcal{G}\} - \text{card}\,\{\textit{critical rose classes of } \mathcal{G}\}.$$

LEMMA 9.4. *If Φ is a maximal oriented ideal forest of $\bar{\Gamma}$ and $\alpha \in \Phi$, then there are* card $(\alpha) - 2$ *ideal edges in Φ which are proper subsets of α.*

Proof. Let I be the set of ideal edges in Φ which are properly contained in α. Decompose α as the disjoint union $\alpha = A_1 + \cdots + A_k$, where each A_i is either a maximal element of I or a single edge $e \in \alpha - \bigcup I$. Choose an element $a \in D(\alpha)$, which we may assume lies in A_1. If $k > 2$, then $A_1 \cup A_2$ is an ideal edge which is compatible with every ideal edge of Φ but not contained in Φ, contradicting maximality of Φ. Thus $k = 2$, and the lemma follows easily by induction on card (α). □

In particular, Lemma 9.4 implies the following.

LEMMA 9.5. *Let Φ be a maximal oriented ideal forest of $\bar{\Gamma}$. Then*

$$\text{card}\,(\Phi) = \sum_{\alpha \in \Phi^o} (\text{card}\,(\alpha) - 1),$$

where Φ^o is the set of all maximal elements of Φ. □

Proof of Proposition 9.3. If Φ is an oriented ideal forest in $\bar{\Gamma}$ and v a vertex of $\bar{\Gamma}$, let Φ_v be the oriented ideal forest consisting of all ideal edges in Φ which are subsets of E_v. It suffices to show that, for every vertex v of $\bar{\Gamma}$ and every maximal oriented ideal forest Φ,

$$\begin{aligned}\text{card}\,(\Phi_v) = \text{card}\,&\{\text{active edges in } E_v\}\\ &-\text{card}\,\{\text{critical conjugacy classes in } \mathcal{G}_v\}\\ &-\text{card}\,\{\text{critical rose classes in } \mathcal{G}_v\}\\ &-\epsilon\end{aligned}$$

where $\epsilon = 1$ if v is a critical vertex, and $\epsilon = 0$ if v is not critical.

Decompose the active edges of E_v into a disjoint union $A_1 + \cdots + A_{k(v)}$, where each A_i is either a maximal element of Φ_v or a singleton edge in $E_v - \bigcup \Phi_v$. The following local version of Lemma 9.5 also follows immediately from Lemma 9.4:

$$\text{card}\,(\Phi_v) = \sum_{i=1}^{k(v)} (\text{card}\,(A_i) - 1).$$

Therefore it suffices to prove

$$\begin{aligned} k(v) = {} & \text{card}\,\{\text{critical conjugacy classes in } \mathcal{G}_v\} \\ & + \text{card}\,\{\text{critical rose classes in } \mathcal{G}_v\} \\ & + \epsilon \end{aligned} \qquad (*)$$

where $\epsilon = 1$ if v is a critical vertex, and $\epsilon = 0$ if v is not critical.

We prove this formula separately for critical and non-critical vertices.

If v is critical, the only critical conjugacy class in $\mathcal{G}_v$ is $\mathcal{G}_v$ itself. Since Γ is reduced, the critical edges at v are all loops, so we have a critical rose at v. All edges in E_v are active. Thus equation $(*)$ reduces to showing that $k(v) = 3$. Choose an edge a with $\mathcal{G}(a) = \mathcal{G}_v$; we may assume $a \in A_1$. If A_1 is an ideal edge, then for any $b \in D(A_1)$ we must also have $\mathcal{G}(b) = \mathcal{G}_v$ and, in addition, $b^{-1} \notin A_1$. Thus we may assume $a \in A_1$ and $a^{-1} \in A_2$. If $k = k(v) \geq 4$, then $A_1 \cup A_3$ is an ideal edge which is compatible with every element of Φ_v, contradicting the maximality of Φ. Thus $k \leq 3$. It is clear that $k > 1$. Suppose $k = 2$. If either A_1 or A_2 is a singleton, this contradicts the definition of ideal edge. If A_1 and A_2 are ideal edges, then $A_2 = A_1^{-1}$, contradicting the definition of oriented ideal forest. Thus $k = 3$.

Now assume v is not critical.

CLAIM 1. *Every active edge in E_v which is not critical belongs to an ideal edge of Φ.*

Proof. If e is an active edge at v which is not critical, then $e \leq a$ for some $a \in E_v$. Suppose $e \notin \bigcup \Phi_v$. Then $\{e\} = A_i$ for some i. If $a \in A_j$ then $A_i + A_j$ is an ideal edge which is compatible with every element of Φ but is not in Φ_v, contradicting maximality of Φ. □

CLAIM 2. *Every maximal element of Φ_v contains a critical edge.*

Proof. Let $\alpha = A_i$ be a maximal ideal edge in Φ_v, and $a \in D(\alpha)$. If a is not critical, then there is some $b \in E_v$ with $a < b$, say $b \in A_j$. Then $e < b$ for all $e \in \alpha$, so $A_i + A_j$ is an ideal edge, contradicting the maximality of Φ. □

Thus each A_i is either a single critical edge, or is a maximal ideal edge α with $P(\alpha)$ critical.

Fix a critical conjugacy class P, and let $k_P(v) = \text{card}\,\{A_i \mid P(A_i) = P\}$. To prove $(*)$ it suffices to show that

$$k_P(v) = \begin{cases} 2 & \text{if } P \text{ is a rose class} \\ 1 & \text{otherwise.} \end{cases}$$

Let $B_1, \ldots, B_r$ denote the sets A_i with $\mathcal{G}(A_i) \in P$. Extend the definition of $D(\alpha)$ to those A_i which are singletons by setting $D(\{e\}) = \{e\}$. If $D(B_i) \neq \{e^{-1} \mid e \in D(B_j)\}$ then $B_i + B_j$ is an ideal edge compatible with every ideal edge in Φ, contradicting maximality of Φ. In particular, we must have $r \leq 2$, and $r = 2$ if and only if P is a rose class. □

C. Free abelian subgroups of $C(G)$ and the vcd

Again, fix a reduced marked G-graph $[s, \Gamma]$, and let $q : \Gamma \to \bar{\Gamma}$ be the quotient map, with quotient graph of groups $\mathcal{G} = \{\bar{\Gamma}, \{\mathcal{G}_e\}, \{\mathcal{G}_v\}\}$.

DEFINITION. An edge e of Γ is *flexible* if the subgraph spanned by the orbit Ge contains a non-trivial cycle stabilized by stab (e).

We can check this condition in $\bar{\Gamma}$ using the following:

(i) If e is a bent edge, then e is flexible;

(ii) if e is a straight elliptic edge, and $\mathcal{G}(e)$ has nontrivial normalizer in both $\mathcal{G}_{\iota(e)}$ and $\mathcal{G}_{\tau(e)}$, then e is flexible.

Note that since $[s, \Gamma]$ is reduced, all edges are either hyperbolic or elliptic, and all hyperbolic edges are bent.

PROPOSITION 9.6. *If all critical edges of Γ are flexible, then the virtual cohomological dimension of $C(G)$ is equal to the dimension of L_G.*

Proof. Let Φ be a maximal oriented ideal forest in $\bar{\Gamma}$, and $\alpha_1, \ldots, \alpha_k$ the maximal ideal edges in Φ. For each $i = 1, \ldots, k$, choose an ideal edge $\tilde{\alpha}_i$ in Γ with image α_i and an edge $a_i \in D(\tilde{\alpha}_i)$. Finally, for each element a of α_i other than the image of a_i, choose a representative $b(a)$ in $\tilde{\alpha}_i$. Set $B_i = \{b(a) \mid a \in \alpha_i, a \neq q(a_i)\}$, and $B = B_1 + \cdots + B_k$.

By Lemma 9.5, $\dim(L_G) = \Sigma_{i=1}^k \text{card}\,(\alpha_i) - 1 = \Sigma_{i=1}^k \text{card}\,(B_i) = \text{card}\,(B)$. In order to prove the proposition, we will find a free abelian subgroup in $C(G)$ with one generator ϕ_b for each element $b \in B$. Since such a free abelian subgroup has

cohomological dimension equal to its rank, dim L_G is a lower bound for the vcd of $C(G)$; but we showed it is an upper bound in the previous section.

It is easiest to define the automorphisms ϕ_b on the level of the edge-path groupoid $\pi(\Gamma)$. Since the ideal edges α_i are maximal, the edges $a_i \in D(\alpha_i)$ are critical (Claim 2 of Proposition 9.3). By hypothesis, all critical edges are flexible; thus for each i we may choose a simple closed cycle p_i in the orbit of a_i which starts at $\tau(a_i)$ and is stabilized by stab (a_i). For $b \in B$, there is a G-equivariant automorphism ϕ_b of $\pi(\Gamma)$ which sends b to bp_i if $b \in B_i$ and fixes every edge which is not in the orbit of b or b^{-1}. It is easy to check that these automorphisms commute, so generate a free abelian subgroup H of the group $\mathrm{Aut}_G(\pi(\Gamma))$ of G-equivariant automorphisms of $\pi(\Gamma)$.

An automorphism of $\pi(\Gamma)$ gives an outer automorphism of $\pi_1(\Gamma)$ by restriction. To complete the proof, it suffices to show that the resulting map

$$f: \mathrm{Aut}_G(\pi(\Gamma)) \to \mathrm{Out}_G(\pi_1(\Gamma)) \cong C(G)$$

is injective on H. By Corollary 2 of [11], the kernel of f is equal to the subgroup of G-equivariant inner automorphisms, where $\phi \in \mathrm{Aut}_G(\pi(\Gamma))$ is *inner* if there are paths λ_v starting at each vertex v with $\lambda_{xv} = x\lambda_v$ for all $x \in G$, and $\phi(e) = \lambda_{\iota(e)}^{-1} e \lambda_{\tau(e)}$. We will show that no non-trivial element of H is inner.

Fix an element $\phi = \Pi_{b \in B} \phi_b^{m_b} \in H$. For each edge e of Γ, we have

$$\phi(e) = (\mu_{e^{-1}})^{-1} e \mu_e, \tag{1}$$

where

$$\mu_e = \begin{cases} (x^{-1}p_i)^{m_e} & \text{if } xe \in B_i \\ 1_{\tau(e)} & \text{if } e \notin \bigcup GB. \end{cases}$$

Suppose that ϕ is inner. Then $\phi(e)$ can also be expressed as

$$\phi(e) = \lambda_{\iota(e)}^{-1} e \lambda_{\tau(e)} \tag{2}$$

for appropriate paths $\lambda_{\iota(e)}$ and $\lambda_{\tau(e)}$.

CLAIM 1. *If e is elliptic, with initial vertex u and terminal vertex v, then $\mu_e = \lambda_v$ and $\mu_{e^{-1}} = \lambda_u$.*

Proof. Since $x\lambda_v = \lambda_{xv}$ for $x \in G$, λ_v is fixed by stab (v). If e is elliptic with $\tau(e) = v$, then neither e nor e^{-1} is in λ_v; in fact neither λ_u nor λ_v contain edges in the orbit of e or e^{-1}. The same is true for μ_e and $\mu_{e^{-1}}$. □

CLAIM 2. *If μ_e is not trivial, then $\lambda_{\tau(e)}$ is not trivial.*

Proof. This follows for e elliptic by Claim 1. If e is hyperbolic, its initial vertex is a translate xv of its terminal vertex v. Thus $\phi(e) = (x\lambda_v^{-1})e\lambda_v$ and $\phi(e) = (\mu_{e^{-1}})^{-1}e\mu_e$. If $\lambda_v = 1$, then $\phi(e) = e$, so $\mu_e = 1$. □

Now suppose that ϕ is non-trivial, i.e. there is some $b \in B$ with $m_b \neq 0$. If $v = \tau(b)$, then $\mu_b \neq 1$, so by Claim 2, $\lambda_v \neq 1$. Consider the edge a_i, where $b \in B_i$. Using equation (1), $\phi(a_i) = a_i$; using (2), $\phi(a_i) = \lambda_{\iota(a_i)}a_i\lambda_v$. By Claim 1, a_i cannot be elliptic. Thus a_i is hyperbolic, so stab (a_i) = stab (v) and $\iota(a_i)$ is a translate xv for some $x \in G$. Equations (1) and (2) give

$$\phi(a_i) = a_i = (x\lambda_v^{-1})a_i\lambda_v.$$

It follows that all edges in λ_v are in the orbit of a_i or a_i^{-1}, and λ_v is a cycle. Since stab (a_i) = stab (v), $\tilde{\alpha}_i$ is invertible, so there are at least two edges in $E_v - \tilde{a}_i$, at least one of which is not a translate of a_i^{-1}. Choose $e \in E_v$ which is not in the orbit of $\tilde{\alpha}_i$ or of Ga^{-1}. By equation (1), $\phi(e) = (\mu_{e^{-1}})^{-1}e\mu_e$ ends in e or in an edge $xa_j^{\pm 1}$ for some $j \neq i$ and some $x \in G$. But by equation (2), $\phi(e) = \lambda_{\iota(e)}^{-1}e\lambda_v$, which ends in a translate of $a_i^{\pm 1}$, a contradiction. Thus ϕ is trivial, so ker (f) and H intersect trivially, i.e. H injects into $\mathrm{Out}_G(\pi(\Gamma))$. □

COROLLARY 9.7. *If all edge groups are normal in their vertex groups, then the* vcd *of $C(G)$ is equal to the dimension of L_G.* □

Remark. The technique used in the proof of Proposition 9.6 can be used to find free abelian subgroups in arbitrary graphs of finite groups, whose ranks depend on conjugacy information present in the graph of groups; this can be used to give a lower bound on the vcd of $C(G)$.

We end this section with an example which shows that the dimension of L_G can sometimes be strictly larger than the vcd of $C(G)$.

EXAMPLE 9.8. Let H_1, H_2 and H_3 be finite groups, and let $P \to H_i$ be monomorphisms such that the image of P is proper and has trivial normalizer for each i. Let E be the amalgamated free product $H_1 *_P H_2 *_P H_3$; we can represent E as the fundamental group of the graph of groups $\mathcal{G}$ shown in Figure 13.

Let F_n be a free subgroup of finite index in E. As in Section 2, there is a subgroup G of Out (F_n) and a G-graph Γ with quotient graph of groups $\mathcal{G}$. From $\mathcal{G}$ one sees that Γ is reduced, since every edge is elliptic.

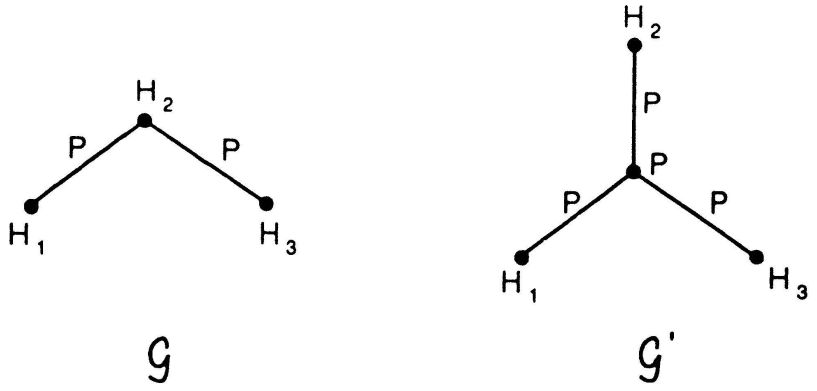

Figure 13

The quotient graph of groups $\mathscr{G}$ has only one ideal edge, consisting of both edges which terminate at the vertex with stabilizer H_2. The condition that the normalizer of P in H_2 is trivial implies that there is in fact only one ideal edge-orbit α in Γ.

Let $s : R_n \to \Gamma$ be a marking of Γ, and let $[s', \Gamma']$ be the (essential) marked graph obtained by blowing up $G\alpha$; Γ' has quotient graph groups $\mathscr{G}'$, shown in Figure 13.

All edges in $\mathscr{G}'$ are parabolic, so the three edge-orbits in Γ' are all (maximal) invariant forests. Collapsing any one results in a reduced marked G-graph similar to $[s, \Gamma]$, but with the H_i permuted; each has a unique ideal edge-orbit which can be blown up to recover $[s', \Gamma']$. Thus the complex L_G is a finite one-dimensional complex with central vertex $[s', \Gamma']$ and three "spokes" leading to the three other vertices. The centralizer $C(G)$ must fix $[s', \Gamma']$, so $\text{vcd}\,(C(G)) = 0$, while $\dim\,(L_G) = 1$.

§10. Corollaries

A. Trivial edge stabilizers

If E is the fundamental group of a finite graph of groups in which the vertex groups are finite and the edge groups are trivial, then E is a free product of the form $G_1 * \cdots * G_m * F_n$, where $m, n \geq 0$, the groups $G_1, \ldots, G_m$ are finite, and F_n is a free group of rank n. Consider the graph of groups representation for E given in Figure 14.

For very small m and n (both less than two or $m = 2, n = 0$), there are no ideal edges at all; in these cases $\text{vcd}\,(\text{Out}\,(E)) = 0$. If there is at least one ideal edge, then the set of active edges is all of E_t, there is only one critical conjugacy class,

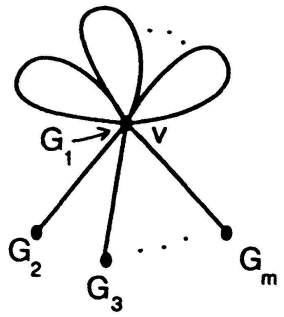

Figure 14

corresponding to the trivial subgroup, and this conjugacy class is a rose class if and only if $m \leq 1$. In view of Proposition 9.3, we have

PROPOSITION 10.1. *If* $E = G_1 * \cdots * G_m * F_n$ *with* G_i *finite and* F_n *free of rank* n, *then*

$$\text{vcd}\,(\text{Out}\,(E)) = \begin{cases} 0 & \text{if } m, n \leq 1 \\ 2n + m - 3 & \text{if } m \leq 1 \text{ and } n > 1\,. \\ 2n + m - 2 & \text{if } m > 1 \end{cases}$$ □

This proposition contains the result of Culler and Vogtmann [7] that vcd (Out $(F_n)) = 2n - 3$ and also has the following corollary, which was conjectured by Collins in [3, 4].

COROLLARY 10.2. *The* vcd *of the outer automorphism group of a free product of* m *finite groups is* $m - 2$, *for* $m \geq 2$. □

Remark. This corollary also follows from work of McCullough and Miller [14].

B. Fixed point set of Outer Space is contractible

The complex K_n is an equivariant deformation retract of the full outer space X_n. It can be shown that X_n does not admit a piecewise-Euclidean metric of non-positive curvature [2], although it has properties reminiscent of non-positively curved metric spaces. We have the following property to add to the list.

COROLLARY 10.3. *Let* G *be a finite subgroup of* Out (F_n). *Then the fixed-point set* X_G *of the action of* G *on outer space* X_n *is contractible.*

Proof. The ideal triangulation of X_n described in [7] restricts to an ideal triangulation of X_G: an ideal simplex is given by assigning lengths to the edges of a marked graph, so that the sum of the lengths of all edges is equal to one. The complex K_n (resp. K_G) is obtained by taking the barycentric subdivision of the ideal trangulation of X_n (resp. X_G), then removing all vertices which are at infinity. The entire ideal simplex collapses linearly onto the subcomplex of the barycentric subdivision spanned by the remaining vertices, giving the equivariant deformation retraction of X_n onto K_n (resp. X_G onto K_G). □

Remark. This has been proved independently by White [19].

C. VFL

COROLLARY 10.4. *The group of outer automorphisms of a free-by-finite group is VFL. In particular, it has finitely-generated homology in all dimensions.*

Proof. $C(G)$ acts transitively on the set of marked G-graphs based on the same G-graph. There are only finitely many homeomorphism types of essential G-graphs. Thus the quotient of L_G by $C(G)$ is finite. □

REFERENCES

1. ALPERIN, R. and BASS, H. *Length functions of group actions on Λ-trees.* Gersten, S. M., et al. *Combinatorial Group Theory and Topology.* Princeton: Princeton University Press; 1987: 265–378.
2. BRIDSON, M. *Geodesics and curvature in metric simplicial complexes.* Ghys, E., Haefliger, A. and Verjovsky, A. *Group Theory from a Geometrical Viewpoint.* Singapore: World Scientific; 1991: 373–463.
3. COLLINS, D. J. *The automorphism group of a free product of finite groups.* Arch. Math. (Basel); 1988; *50* (5): 385–390.
4. COLLINS, D. J. *Cohomological dimension and symmetric automorphisms of a free group.* Comment. Math. Helv.: 1989; *64* (1): 44–61.
5. CULLER, M. *Finite groups of outer automorphisms of a free group.* Appel, K. I., Ratcliffe, J. G. and Schupp, P. E. *Contributions to Group Theory.* Providence, R.I.: American Math. Soc.; 1984; *33*: 197–207 (Contemporary Math.).
6. CULLER, M. and MORGAN, J. W. *Group actions on* **R***-trees.* Proc. Lond. Math. Soc.; 1987; *55*: 571–604.
7. CULLER, M. and VOGTMANN, K. *Moduli of graphs and automorphisms of free groups.* Invent. Math.; 1986; *84*: 91–119.
8. DICKS, W. and DUNWOODY, M. J. *Groups acting on graphs.* Cambridge: Cambridge University Press; 1989.
9. KALAJDŽIEVSKI, S. *Automorphism group of a free group: centralizers and stabilizers.* J. Algebra; 1992; *150* (2): 435–502
10. KARRASS, A., PIETROWSKI, A. and SOLITAR, D. *Finite and infinite cyclic extensions of free groups.* J. Austral. Math. Soc.; 1973; *16*: 458–466.

11. KRSTIĆ, S. *Actions of finite groups on graphs and related automorphisms of free groups.* J. of Algebra; 1989; *24* (1): 119–138.
12. KRSTIĆ, S. *Finitely generated virtually free groups have finitely presented automorphism group.* Proc. London Math. Soc.; 1992; *64* (1): 49–69.
13. MCCOOL, J. *The automorphism groups of finite extensions of free groups.* Bull. London Math. Soc.; 1988; *20*: 131–135.
14. MCCULLOUGH, D. and MILLER, A. In preparation.
15. QUILLEN, D. *Homotopy properties of the poset of p-subgroups of a finite group.* Advances in Math.; 1978; *28*: 101–128.
16. PARRY, W. *Axioms for translation length functions.* Alperin, R. C. *Arboreal Group Theory.* New York: Springer-Verlag; 1991: 295–330.
17. SERRE, J.-P. *Trees.* Berlin: Springer-Verlag; 1980. Translation of "*Arbres, Amalgames,* SL_2", Asterisque *46*, 1977.
18. SMILLIE, J. and VOGTMANN, K. *A generating function for the Euler characteristic of* $Out(F_n)$. J. Pure and Applied Algebra; 1987; *44*: 329–348.
19. WHITE, T. *Fixed points of finite groups of free group automorphisms*, to appear in Proc. A.M.S.

Matematički Institut
Knez-Mihailova 35
11000 Belgrade
Yugoslavia

and

Department of Mathematics
Cornell University
Ithaca, NY 14853-7901

Received February 25, 1992