

AG
T

*Algebraic & Geometric
Topology*

Volume 25 (2025)

Finite groups of untwisted outer automorphisms of RAAGs

COREY BREGMAN

RUTH CHARNEY

KAREN VOGTMANN

Finite groups of untwisted outer automorphisms of RAAGs

COREY BREGMAN

RUTH CHARNEY

KAREN VOGTMANN

For any right-angled Artin group A_Γ , Charney, Stambaugh and Vogtmann showed that the subgroup $U^0(A_\Gamma) \leq \text{Out}(A_\Gamma)$ generated by Whitehead automorphisms and inversions acts properly and cocompactly on a contractible space K_Γ . We show that any finite subgroup of $U^0(A_\Gamma)$ fixes a point of K_Γ . This generalizes the fact that any finite subgroup of $\text{Out}(F_n)$ fixes a point of outer space, and implies that there are only finitely many conjugacy classes of finite subgroups in $U^0(A_\Gamma)$.

20F28, 20F36, 20F65

1 Introduction

A right-angled Artin group (RAAG) is a finitely generated group whose only defining relations are that some of the generators commute. This can be encoded by forming a finite simplicial graph Γ with one vertex for each generator and an edge between each pair of commuting generators; the associated RAAG is then called A_Γ . The extreme examples are the free group F_n (if Γ has no edges) and the free abelian group $\mathbb{Z}^n$ (if Γ is a complete graph). We are interested in studying finite subgroups of the group $\text{Out}(A_\Gamma)$ of outer automorphisms of A_Γ .

For $A_\Gamma = \mathbb{Z}^n$, it follows from the classical Jordan–Zassenhaus theorem that there are only finitely many conjugacy classes of finite subgroups in $\text{Out}(A_\Gamma) = \text{GL}(n, \mathbb{Z})$ (see eg Curtis and Reiner [6]). Since $\text{GL}(n, \mathbb{Z})$ acts on the symmetric space $\text{GL}(n, \mathbb{R})/O(n)$ preserving a CAT(0) metric, any finite subgroup fixes a point. Since $\text{GL}(n, \mathbb{R})/O(n)$ can be identified with the space of marked lattices $\Lambda \subset \mathbb{R}^n$, where a *marking* is a choice of basis B , which gives an isomorphism $\Lambda \cong \mathbb{Z}^n$, it follows that any finite subgroup $G < \text{GL}(n, \mathbb{Z})$ acts by isometries on a lattice Λ . Equivalently, any finite subgroup $G < \text{GL}(n, \mathbb{Z})$ can be embedded in the isometry group of a flat torus T , so that the induced action on π_1 agrees with G .

For $A_\Gamma = F_n$, there is a realization theorem that says any finite subgroup G of $\text{Out}(F_n)$ can be realized as automorphisms of a finite graph X ; see Culler [5], Khramtsov [10] and Zimmermann [15]. This means one can *mark* the graph by an isomorphism $\pi_1(X) \cong F_n$ so that automorphisms of X induce the elements of G on π_1 . Furthermore, one may assume that all vertices of X have valence at least three. Since there

are only finitely many such graphs, this implies that there are only finitely many conjugacy classes of finite subgroups of $\text{Out}(F_n)$. Thus one can study finite subgroups of $\text{Out}(F_n)$ by studying symmetries of such graphs (see eg Levitt and Nicolas [12] and Smillie and Vogtmann [14]). An equivalent way to state the realization theorem is that the action of the finite group $G \leq \text{Out}(F_n)$ on outer space CV_n has a fixed point.

In previous work we constructed an *outer space* $\mathbb{O}_\Gamma$ for an arbitrary RAAG A_Γ that combines features of both CV_n and symmetric spaces [1]. The group $\text{Out}(A_\Gamma)$ acts on $\mathbb{O}_\Gamma$ with finite stabilizers, and it is proved in [1] that $\mathbb{O}_\Gamma$ is contractible. The group $\text{Out}(A_\Gamma)$ contains a natural *untwisted* subgroup $U^0(A_\Gamma)$, which is the whole group in some cases, including when $A_\Gamma = F_n$. The results in [1] build on the fact that $\mathbb{O}_\Gamma$ contains a subspace K_Γ on which the subgroup $U^0(A_\Gamma)$ acts with compact quotient. The space K_Γ was first defined by Charney, Stambaugh and Vogtmann [4], who proved it to be contractible. Points in K_Γ are special types of cube complexes called Γ -complexes, with special types of markings called *untwisted markings*. We prove the following theorem.

Theorem 8.1 *Let Γ be a simplicial graph, G a finite group and $\rho: G \rightarrow U^0(A_\Gamma)$ a homomorphism. Then there is a Γ -complex X with an untwisted marking $h: X \rightarrow \mathbb{S}_\Gamma$ on which ρ is realized by isometries.*

The following corollary is immediate.

Corollary 8.2 *Any finite subgroup of $U^0(A_\Gamma)$ has a fixed point in K_Γ (and therefore in $\mathbb{O}_\Gamma$).*

We conjecture that the entire fixed-point set is contractible, ie that K_Γ is an $\underline{\text{EG}}$ for $G = U^0(A_\Gamma)$. Corollary 8.2 is a necessary first step towards this goal.

It is easy to see that there are only a finite number of combinatorial types of Γ -complexes, generalizing the fact that there are only a finite number of combinatorial types of graphs in CV_n . This gives us the following information about finite subgroups of $U^0(A_\Gamma)$.

Corollary 8.3 *The group $U^0(A_\Gamma)$ contains only finitely many conjugacy classes of finite subgroups.*

Extending these theorems to all of $\text{Out}(A_\Gamma)$ presents subtle difficulties that we do not address in this paper. Among these is the problem of including outer automorphisms of A_Γ that are induced by graph automorphisms of Γ . More serious is the fact that the full group $\text{Out}(A_\Gamma)$ may contain finite subgroups of $\text{GL}(n, \mathbb{Z})$ which do not preserve any Γ -complex structure, so that understanding these will require additional techniques involving the action of $\text{GL}(n, \mathbb{Z})$ on the symmetric space $\text{GL}(n, \mathbb{R})/O(n)$.

To prove Theorem 8.1, we use an inductive approach which starts from the realization theorem for $\text{Out}(F_n)$. This was inspired by work of Hensel and Kielak [8], who proved that a finite subgroup G of $U^0(A_\Gamma)$ can be realized on some cube complex, but it is not clear whether this can be taken to be a Γ -complex. We borrow a number of ideas from [8], but our proof is shorter. In particular, much of our proof is independent of the specific group G being considered, depending rather on the combinatorial structure of the defining graph Γ .

Structure of the paper and outline of the proof

In Sections 2 and 3, we review the group $U^0(A_\Gamma)$ and the definition and basic properties of Γ -complexes. The strategy of the proof is to build a marked Γ -complex realizing a finite $G < U^0(A_\Gamma)$ by gluing together marked Δ -complexes for certain subgraphs $\Delta \subseteq \Gamma$. The subgraphs we use are those whose associated special subgroup A_Δ is invariant (up to conjugacy) under $U^0(A_\Gamma)$, which we will call *U^0 -invariant subgraphs*. The argument is inductive, and the induction parameter is the *chain length* of Γ , ie the longest length of a chain of U^0 -invariant subgraphs contained in Γ . In Section 4, we study U^0 -invariant subgraphs Δ , show there is a restriction homomorphism $r_\Delta : U^0(A_\Gamma) \rightarrow U^0(A_\Delta)$, and show that minimal U^0 -invariant subgraphs are discrete, providing a base case for our induction.

In Section 5, we show that a marked Γ -complex that realizes a finite subgroup $G < U^0(A_\Gamma)$ contains a subcomplex associated to each U^0 -invariant subgraph Δ with empty link, and that the restriction of the marking to this subcomplex realizes the restriction of G to $U^0(A_\Delta)$. In Section 6, we address the opposite problem, establishing a necessary condition for extending a Δ -complex realizing the restriction of G to a Γ -complex realizing G .

In Section 7, we show how to build marked Γ -complexes when Γ is a simplicial join or a disjoint union of subgraphs Δ for which we already have marked Δ -complexes. We also show that, if Γ is a join or disjoint union, and one can realize the restriction of a finite subgroup $G < U^0(A_\Gamma)$ on each component, then one can realize all of G .

Finally, in Section 8, we induct on the length of a maximal chain of $U^0(A_\Gamma)$ -invariant subgraphs to construct a marked Γ -complex that realizes G .

Acknowledgements

We would like to thank Dawid Kielak for several helpful discussions and the referee for giving thoughtful comments. Bregman is supported by NSF grant DMS-2052801.

2 Review of RAAGs and the untwisted subgroup of $\text{Out}(A_\Gamma)$

Let Γ be a finite simplicial graph. The *right-angled Artin group* (RAAG) A_Γ is the group generated by the vertices V of Γ with defining relations given by declaring that adjacent vertices commute.

By a subgraph of Γ we will always mean a full (induced) subgraph, unless otherwise specified. Given a subgraph $\Delta \subseteq \Gamma$, we write $x \in \Delta$ if x is a vertex of Δ .

For $x \in \Gamma$, the *link* $\text{lk}(x)$ is the subgraph spanned by vertices adjacent to x . The link of a subgraph $\Delta \subset \Gamma$ is the intersection of the links of all vertices of Δ . The *double link* $\text{dlk}(x)$ is the link of $\Delta = \text{lk}(x)$. The *star* $\text{st}(x)$ is the subgraph spanned by x and $\text{lk}(x)$.

Recall from [4] that a Γ -*Whitehead partition* $\mathcal{P}$ based at $x \in \Gamma$ is a partition of $V^\pm = V \cup V^{-1}$ into three sets $\text{lk}^\pm(x)$, P_1 and P_2 satisfying certain conditions. The sets P_1 and P_2 are called the *sides* of $\mathcal{P}$. A

Γ -Whitehead partition can be most easily described using the *double* $\Gamma^\pm$ of Γ , where the vertices of $\Gamma^\pm$ are $V^\pm$ and two vertices are joined by an edge if they commute but are not inverses of each other. If $\mathcal{P}$ is based at x , $\text{lk}^\pm(x)$ consists of all vertices adjacent to x in $\Gamma^\pm$, and each of P_1 and P_2 is a union of (the vertices in) some connected components of $\Gamma^\pm \setminus \text{lk}^\pm(x)$. Furthermore, we require x and x^{-1} to be in different sides of $\mathcal{P}$, and each side must contain at least one additional element. We will abbreviate Γ -Whitehead partition to simply Γ -partition.

A vertex $y \in V$ is *split* by a Γ -partition $\mathcal{P}$ if y is in one side and y^{-1} is in the other. If y is split by $\mathcal{P}$, then y and y^{-1} must lie in different components of $\Gamma^\pm \setminus \text{lk}^\pm(x)$; hence, $\text{lk}(y) \subseteq \text{lk}(x)$.

A Γ -partition $\mathcal{P}$ based at x determines a *Whitehead automorphism* $\varphi(\mathcal{P}, x)$, defined as follows. Let P_i be the side of $\mathcal{P}$ containing x . If $\mathcal{P}$ splits y , then $\varphi(\mathcal{P}, x)$ sends $y \mapsto yx^{-1}$ if $y \in P_i$, and $y \mapsto xy$ if $y^{-1} \in P_i$. If both y and y^{-1} are in P_i , then $\varphi(\mathcal{P}, x)(y) = xyx^{-1}$. For all other y , $\varphi(\mathcal{P}, x)(y) = y$. The simplest Whitehead automorphisms are the *folds* sending $y \mapsto yx^{-1}$ or $y \mapsto xy$ for some x and y (and fixing all generators other than y), and the *partial conjugations* sending $y \mapsto xyx^{-1}$ for all y in some component C of $\Gamma \setminus \text{lk}(x)$. These correspond to partitions $\mathcal{P} = (\text{lk}^\pm(x) \mid P_1 \mid P_2)$ with $P_1 = \{x, y\}$ or $\{x, y^{-1}\}$ (for a fold) or $P_1 = \{x, C^\pm\}$ (for a partial conjugation). Every Whitehead automorphism is a product of folds and partial conjugations.


The subgroup of $\text{Out}(A_\Gamma)$ generated by Whitehead automorphisms and by inversions of the generators is denoted by $U^0(A_\Gamma)$. If $A_\Gamma = F_n$, this is the whole group, ie $U^0(F_n) = \text{Out}(F_n)$. If $A_\Gamma = \mathbb{Z}^n$, there are no Whitehead automorphisms, and $\text{Out}(\mathbb{Z}^n) = \text{GL}(n, \mathbb{Z})$ is generated by inversions and *twists*, where a twist sends a generator y to xy for some y with $\text{st}(y) \subseteq \text{st}(x)$ and fixes all other generators. By a theorem of Laurence and Servatius [11; 13], for a general RAAG the group $\text{Out}(A_\Gamma)$ is generated by Whitehead automorphisms, inversions, twists and automorphisms of Γ .

The subgroup generated by $U^0(A_\Gamma)$ and graph automorphisms was called the *untwisted subgroup* and denoted by $U(A_\Gamma)$ in [4; 1].

3 Blowups and Γ -complexes

Let $\mathcal{P} = (\text{lk}^\pm(x) \mid P_1 \mid P_2)$ be a Γ -partition based at x . If $\text{lk}(x) = \text{lk}(y)$ and $\mathcal{P}$ splits y , then y can also serve as a base for $\mathcal{P}$. Specifying a choice of base specifies the corresponding Whitehead automorphism, but we will often use Γ -partitions without specifying a base, in which case we write $\mathcal{P} = (\text{lk}(\mathcal{P}) \mid P_1 \mid P_2)$.

We say Γ -partitions $\mathcal{P}, \mathcal{Q}$ are *adjacent* if some (hence any) base of $\mathcal{P}$ commutes with some (any) base of $\mathcal{Q}$, and they are *compatible* if either they are adjacent or some side of $\mathcal{Q}$ is disjoint from some side of $\mathcal{P}$. A collection Π of Γ -partitions is a *compatible collection* if its elements are distinct and pairwise compatible. In [4], the authors constructed a labeled cube complex $\mathbb{S}_\Gamma^\Pi$ called a *blowup* from a compatible collection Π of Γ -partitions. The underlying (unlabeled) cube complex is called a Γ -complex. In this section we review some facts about special cube complexes and Γ -complexes that we will need.

Figure 1: Hyperplane collapse $X \rightarrow X//H$.

3.1 Special cube complexes, collapsing and duplicating hyperplanes

Recall that a cube complex is called a *special cube complex* if it is locally CAT(0) and has no hyperplanes that selfintersect or are one-sided, selfosculating or interosculating. We refer to the original article by Haglund and Wise [7] for the basic definitions.

Let X be a special cube complex. If H is a hyperplane in X , the *collapse map* $c: X \rightarrow X//H$ collapses the carrier $\kappa(H)$ of H orthogonally onto H . We say the result $X//H$ is obtained from X by a *hyperplane collapse* (see Figure 1).

The edges that intersect a hyperplane H are said to be *dual* to H , and by an *orientation* on H we mean a consistent choice of orientation of the edges dual to H .

If $\mathcal{S}$ is a collection of hyperplanes, we write $X//\mathcal{S}$ for the space obtained by collapsing all hyperplanes in $\mathcal{S}$ (in any order). The collection $\mathcal{S}$ is *acyclic* if the collapse map $X \rightarrow X//\mathcal{S}$ is a homotopy equivalence.

If H is a hyperplane in X with carrier $\kappa(H)$, we can obtain a new cube complex by doubling $\kappa(H)$ (see Figure 2). We will refer to this as *duplicating the hyperplane* H . The resulting cube complex has two




Figure 2: Duplicating a hyperplane.

new hyperplanes H' and H'' , and collapsing either recovers the original complex X . We say H' and H'' are *parallel*. A hyperplane is called a *duplicate* if it is parallel to another hyperplane.

3.2 Blowups

The blowup $\mathbb{S}_\Gamma^\Pi$ associated to a compatible collection Π of Γ -partitions is a special cube complex with no separating hyperplanes and with some extra structure.

If Γ is discrete, then $\mathbb{S}_\Gamma^\Pi$ is a finite connected graph with no separating edges or bivalent vertices, and the extra structure consists of a maximal tree T and an orientation and label on each edge in $\mathbb{S}_\Gamma^\Pi \setminus T$, where the labels are the vertices of Γ . Each edge of T corresponds to a partition in Π , determined by the labels and orientations of the edges not in T .

If Π is empty, then $\mathbb{S}_\Gamma^\Pi$ is the *Salvetti complex* $\mathbb{S}_\Gamma$ associated to A_Γ . Recall that this is a cube complex with a single 0-cell, one oriented 1-cell for each vertex of Γ and one k -torus for each k -clique in Γ . The orientations on the 1-cells, which are labeled by vertices of Γ , determine an isomorphism $\pi_1(\mathbb{S}_\Gamma) \cong A_\Gamma$, and the cubical isomorphisms of $\mathbb{S}_\Gamma$ can be identified with the automorphisms of the graph Γ .

In general, $\mathbb{S}_\Gamma^\Pi$ has

- one hyperplane $H_\mathcal{P}$ for each partition $\mathcal{P} \in \Pi$, and
- one hyperplane H_v for each vertex $v \in \Gamma$.

The hyperplanes H_v are oriented, but the hyperplanes $H_\mathcal{P}$ are not.

The set of hyperplanes labeled by partitions is acyclic, and the complex obtained by collapsing all hyperplanes in this set is isomorphic to the Salvetti complex $\mathbb{S}_\Gamma$.

Collapsing a single hyperplane labeled by $\mathcal{P} \in \Pi$ is equivalent to removing $\mathcal{P}$ from the collection Π . In particular, collapsing every hyperplane in Π other than $\mathcal{P}$ results in a *single blowup* $\mathbb{S}^\mathcal{P}$. This has exactly two vertices x_1 and x_2 , and one can recover $\mathcal{P} = (\text{lk}(\mathcal{P}) \mid P_1 \mid P_2)$ from $\mathbb{S}^\mathcal{P}$ by looking at the (oriented!) edges dual to the hyperplanes H_v : if there is only one edge dual to H_v and it terminates at x_i , then $v \in P_i$; if it originates at x_i , then $v^{-1} \in P_i$. If there are two edges dual to H_v , then $v \in \text{lk}(\mathcal{P})$. The carrier of $H_\mathcal{P}$ in $\mathbb{S}^\mathcal{P}$ is isomorphic to the product of an interval with the Salvetti for $\text{lk}(\mathcal{P})$.

3.3 Γ -complexes

Definition 3.1 A cube complex is called a Γ -complex if it is isomorphic to the underlying cube complex of a blowup $\mathbb{S}_\Gamma^\Pi$. A *blowup structure* on a Γ -complex X is a labeling of its hyperplanes that identifies X with a blowup $\mathbb{S}_\Gamma^\Pi$, ie hyperplanes are labeled by Γ -partitions or by vertices of Γ , and the hyperplanes labeled by vertices are oriented. A blowup structure determines a *collapse map* $c_\pi: X \rightarrow \mathbb{S}_\Gamma$ that collapses all hyperplanes labeled by partitions. If $v \in \Gamma$, a *characteristic cycle* for v is a closed edge path which crosses each hyperplane at most once, and whose image under c_π is the loop labeled v .

In general, a Γ -complex may have several different blowup structures. For example, if Γ is discrete, a Γ -complex is a graph, which may have several different maximal trees, and the remaining edges may be oriented and labeled with the vertices of Γ in any way.

Definition 3.2 A set $\mathcal{T}$ of hyperplanes in a Γ -complex is called *treelike* if collapsing $\mathcal{T}$ gives a cube complex isomorphic to $\mathbb{S}_\Gamma$.

The following proposition says that any treelike set of hyperplanes in a Γ -complex is the set of hyperplanes labeled by partitions in at least one blowup structure. The only ambiguity comes from the assignment of labels and orientations to the hyperplanes not in the treelike set, which can be permuted by any automorphism of the graph Γ . If this assignment changes, the partitions labeling the hyperplanes in the tree also change, by the same (signed) permutation of vertices.

Proposition 3.3 Let $\mathcal{T}$ be a treelike set of hyperplanes in a Γ -complex X . Then there is a compatible set of Γ -partitions Π and an isometry $X \cong \mathbb{S}_\Gamma^\Pi$, such that $\mathcal{T}$ is the set of hyperplanes associated to the partitions in Π .

Proof We recall the construction. For complete details see [4, Section 4].

Label the edges dual to each hyperplane $H \in \mathcal{T}$ by H . Choose an isomorphism of $X//\mathcal{T}$ with $\mathbb{S}_\Gamma$; this orients each hyperplane that is not in $\mathcal{T}$ and labels its dual edges by a vertex of Γ . The set of cubes in X with all edge labels in $\mathcal{T}$ forms a CAT(0) subcomplex $\mathbb{C}$ that contains all vertices of X . A hyperplane $H \in \mathcal{T}$ cuts $\mathbb{C}$ into two pieces, so partitions the vertices of X into two sets, $v_1(H)$ and $v_2(H)$. Now form a partition $(\text{lk}(H) \mid U_1 \mid U_2)$ of $V^\pm$ as follows:

- (1) If the hyperplane H_v labeled by v intersects H , then v and v^{-1} are in $\text{lk}(H)$.
- (2) If $H_v \cap H = \emptyset$ and the terminal vertex of an edge dual to H_v is in $v_i(H)$, then $v \in U_i$.
- (3) If $H_v \cap H = \emptyset$ and the initial vertex of an edge dual to H_v is in $v_i(H)$, then $v^{-1} \in U_i$.

Then the partition $(\text{lk}(H) \mid U_1 \mid U_2)$ is a Γ -partition, the set of Γ -partitions for all $H \in \mathcal{T}$ is a compatible collection Π , and X is isomorphic to $\mathbb{S}_\Gamma^\Pi$. $\square$

We note that condition (1) in the proof of Proposition 3.3 is equivalent to saying that, in the universal cover $\widetilde{\mathbb{S}_\Gamma^\Pi}$, some lift of H contains an axis for v . Saying that v and v^{-1} are in different P_i is equivalent to saying that an axis for v in $\widetilde{\mathbb{S}_\Gamma^\Pi}$ intersects some lift of H transversally; in this case we say the axis *skewers* H . Saying v and v^{-1} are in the same P_i is equivalent to saying that no axis for v intersects any lift of H .

If we are given a special cube complex X which we do not know a priori is a Γ -complex, then, to prove that it is, we first need to find an acyclic collection $\mathcal{T}$ of hyperplanes which collapses to give a cube complex isomorphic to $\mathbb{S}_\Gamma$. Choosing an isomorphism $X//\mathcal{T} \cong \mathbb{S}_\Gamma$ gives a labeling and orientation on all of the remaining hyperplanes. We then need to check that each hyperplane $H \in \mathcal{T}$ determines a Γ -partition. We can do this by collapsing all hyperplanes other than H to get a complex with two vertices,

then checking whether the location of the initial and terminal vertices of edges labeled by $v \in \Gamma$ gives a valid Γ -partition. By Proposition 3.3, the partitions for one treelike set $\mathcal{T}$ are all Γ -partitions if and only if this holds for every treelike set.

3.4 Subdividing blowups

In a blowup $\mathbb{S}_\Gamma^\Pi$, no two hyperplanes are parallel, ie there are no duplicate hyperplanes. However, in the arguments that follow we will need to allow cubical subdivisions of blowups that result in duplicate hyperplanes. Duplicating $H_\mathcal{P}$ can be thought of as subdividing its carrier $\kappa(H_\mathcal{P})$, and is equivalent to adding a duplicate copy of $\mathcal{P}$ to Π . We want both of the new hyperplanes we have created to be in the treelike set since we must collapse both to recover $\mathbb{S}_\Gamma$. Subdividing the carrier of H_v is a little subtler; here we want only one of the two new hyperplanes to be added to the treelike set, so that collapsing the treelike set still gives $\mathbb{S}_\Gamma$. In other words, when we duplicate H_v , we want one of the two resulting hyperplanes to be labeled H_v , and the other to correspond to a partition. We also need the new H_v to have the orientation induced from the old H_v . This is accomplished by adding a “singleton partition” to Π ; this is a partition based at v with one side containing only v (if we want the initial segment of the dual edge to retain the v label) or v^{-1} (if we want the terminal segment to retain the v label). To make this a canonical operation, we can consistently use the singleton partition $\mathcal{S}_{v^{-1}} = (\text{lk}^\pm(v) \mid \{v^{-1}\} \mid (V \setminus \text{lk}(v))^\pm \setminus \{v^{-1}\})$, so that the terminal segment always retains the v label.

Note that duplicate partitions fit the definition of “compatible with each other”, and a singleton partition is compatible with every Γ -partition. A set of pairwise compatible partitions that is allowed to have singletons and duplicates will be called a *compatible multiset*. By the above remarks, compatible multisets correspond to subdivided blowups.

4 U^0 -invariant subgraphs

Recall that a *marking* on a Γ -complex is a homotopy equivalence $h: X \rightarrow \mathbb{S}_\Gamma$. Let G be a finite group and $\rho: G \rightarrow \text{Out}(A_\Gamma)$ a homomorphism.

Definition 4.1 A marked Γ -complex (X, h) *realizes* ρ if there is an action $f: G \rightarrow \text{Aut}(X)$ of G on X by cubical automorphisms such that $h \circ f(g) \circ h^{-1}$ induces $\rho(g)$ on $\pi_1(\mathbb{S}_\Gamma) = A_\Gamma$ for all $g \in G$.

Our goal is to build a marked Γ -complex that realizes (the inclusion of) a finite subgroup $G < U^0(A_\Gamma)$. Our approach is inductive. Specifically, we will build our Γ -complex by gluing together marked Δ -complexes for subgraphs Δ which are U^0 -invariant, in the sense that elements of $U^0(A_\Gamma)$ preserve the special subgroup A_Δ (up to conjugacy). We will induct on the length of a maximal chain of U^0 -invariant subgraphs. In this section we prepare for this by establishing some basic facts about U^0 -invariant subgraphs.

Definition 4.2 Given $\phi \in \text{Out}(A_\Gamma)$, a subgraph Δ of Γ is ϕ -invariant if $\hat{\phi}(A_\Delta)$ is conjugate to A_Δ for some lift $\hat{\phi}$ of ϕ to $\text{Aut}(A_\Gamma)$. Since any two such lifts differ by an inner automorphism, this is well defined. A subgraph Δ is U^0 -invariant if it is ϕ -invariant for every $\phi \in U^0(A_\Gamma)$.

The next lemma gives a criterion for U^0 -invariance.

Lemma 4.3 *Let Δ be a subgraph of Γ . Then Δ is U^0 -invariant if and only if the following two conditions hold for all $x, y \in \Gamma$:*

- (i) *If $x \in \Delta$ and $\text{lk}(x) \subseteq \text{lk}(y)$, then $y \in \Delta$.*
- (ii) *If Δ intersects more than one component of $\Gamma \setminus \text{st}(y)$, then $y \in \Delta$.*

Proof Since every subgraph $\Delta \subseteq \Gamma$ is invariant under inversions, a subgraph Δ is U^0 -invariant if and only if it is invariant under the remaining generators of $U^0(A_\Gamma)$, ie all folds and partial conjugations. There is a fold $\tau: x \mapsto xy$ if and only if $\text{lk}(y) \supseteq \text{lk}(x)$, so τ maps $x \in \Delta$ to A_Δ if and only if $y \in \Delta$. If Δ intersects two different components C and C' of $\Gamma \setminus \text{st}(y)$, then A_Δ is sent to a conjugate of itself under the partial conjugation $C \mapsto yCy^{-1}$ if and only if $y \in \Delta$. $\square$

Proposition 4.4 *Let Γ be a simplicial graph.*

- (1) *If Σ is a subgraph of Γ , then $\Delta = \text{lk}(\Sigma)$ is U^0 -invariant.*
- (2) *If Δ_1 and Δ_2 are two U^0 -invariant subgraphs of Γ , then $\Delta_1 \cap \Delta_2$ is U^0 -invariant.*
- (3) *If Δ_1 and Δ_2 are two U^0 -invariant subgraphs of Γ whose join $\Delta_1 * \Delta_2$ is also a subgraph, then $\Delta_1 * \Delta_2$ is U^0 -invariant.*
- (4) *If Σ is a nonsingleton connected component of a U^0 -invariant subgraph Δ , then Σ is U^0 -invariant.*
- (5) *If Δ is U^0 -invariant and $N(\Delta)$ is the subgraph spanned by Δ and all vertices adjacent to Δ , then $N(\Delta)$ is U^0 -invariant.*

Proof In each case we check conditions (i) and (ii) of Lemma 4.3:

- (1) (i) $x \in \Delta = \text{lk}(\Sigma)$ if and only if $\Sigma \subseteq \text{lk}(x)$. If $\text{lk}(x) \subseteq \text{lk}(y)$, then $\Sigma \subseteq \text{lk}(y)$, so $y \in \text{lk}(\Sigma)$.
- (ii) Suppose that $\Delta = \text{lk}(\Sigma)$ intersects two different components C_1 and C_2 of $\Gamma \setminus \text{st}(y)$, say $x_1 \in C_1 \cap \Delta$ and $x_2 \in C_2 \cap \Delta$. If $z \in \Sigma$, then x_1 and x_2 are both connected to z , so z must be in $\text{lk}(y)$. Thus, $\Sigma \subseteq \text{lk}(y)$, so $y \in \text{lk}(\Sigma)$.
- (2) (i) Let $x \in \Delta_1 \cap \Delta_2$. If $\text{lk}(y) \supseteq \text{lk}(x)$, then $y \in \Delta_1$ by invariance of Δ_1 ; similarly, $y \in \Delta_2$, so y is in the intersection.
- (ii) If $\Delta_1 \cap \Delta_2$ intersects two components of $\Gamma \setminus \text{st}(y)$, then the same is true of both Δ_1 and Δ_2 , so y is in both.
- (3) (i) If $x \in \Delta_1$ and $\text{lk}(y) \supseteq \text{lk}(x)$, then $y \in \Delta_1$, and similarly if $x \in \Delta_2$.
- (ii) Suppose $\Delta = \Delta_1 * \Delta_2$ intersects two different components C and C' of $\Gamma \setminus \text{st}(y)$. Let $x \in C \cap \Delta$ and $x' \in C' \cap \Delta$. Then x and x' must be in the same Δ_i , since otherwise there is an edge connecting them. But Δ_i is U^0 -invariant, so $y \in \Delta_i \subseteq \Delta$.
- (4) (i) Suppose $x \in \Sigma$, and $\text{lk}(y) \supset \text{lk}(x)$. Since Σ is not a singleton, $\text{lk}(x) \cap \Sigma$ contains a vertex z . Since z is in the links of both x and y , y is also in Σ .

- (ii) Suppose $x, z \in \Sigma$ are in different components of $\Gamma \setminus \text{st}(y)$. Since Δ is U^0 -invariant, this implies $y \in \Delta$, but then $\text{st}(y)$ cannot separate x from z unless $y \in \Sigma$.
- (5) (i) If $x \in N(\Delta)$, then either $x \in \Delta$ or $\text{lk}(x) \cap \Delta \neq \emptyset$. If the distance from y to Δ is at least 2, then $\text{lk}(y) \cap \Delta = \emptyset$, so $\text{lk}(x) \not\subseteq \text{lk}(y)$.
- (ii) If $y \notin N(\Delta)$, then $\Delta \cap \text{st}(y) = \emptyset$. Since Δ is U^0 -invariant, it lies in a single component of $\Gamma \setminus \text{st}(y)$, so all $x \in N(\Delta) \setminus \text{st}(y)$ must lie in the same component. $\square$

Recall from [1] that two vertices are called *fold-equivalent* if they have the same link in Γ , and we order the set of fold-equivalence classes by inclusion of their links. The following two propositions will allow us to establish the base case of our induction.

Proposition 4.5 *Let $\Delta \subseteq \Gamma$ be a minimal U^0 -invariant subgraph. Then Δ is a maximal fold-equivalence class. In particular, Δ is discrete.*

Proof If $[u]$ is a maximal fold-equivalence class, then it is easy to check using Lemma 4.3 that $[u]$ is U^0 -invariant. On the other hand, if $\Delta \subseteq \Gamma$ is any U^0 -invariant subgraph, then Δ contains a maximal equivalence class $[u]$ by condition (1) of Lemma 4.3, so $[u] = \Delta$ by the minimality of Δ . $\square$

Proposition 4.6 *Let Δ be a U^0 -invariant subgraph of Γ . Then there is a restriction homomorphism $r_\Delta: U^0(A_\Gamma) \rightarrow U^0(A_\Delta)$.*

Proof Let ϕ be an element of $U^0(A_\Gamma)$. Since Δ is U^0 -invariant, there is a lift $\hat{\phi}$ of ϕ to $\text{Aut}^0(A_\Gamma)$ with $\hat{\phi}(A_\Delta) = A_\Delta$. Define $r_\Delta(\phi)$ to be the image in $\text{Out}(A_\Delta)$ of the restriction of $\hat{\phi}$ to A_Δ .

To check that r_Δ is well defined, suppose $\hat{\phi}'$ is another lift of ϕ sending A_Δ to itself. Then $\hat{\phi}' = \iota_g \circ \hat{\phi}$, where ι_g is conjugation by some $g \in A_\Gamma$ that normalizes A_Δ . By [3, Lemma 2.2], the normalizer of A_Δ is $A_\Delta \times A_{\text{lk}(\Delta)}$. Since elements of $A_{\text{lk}(\Delta)}$ act trivially by conjugation, $\hat{\phi}' = \iota_h \circ \hat{\phi}$ for some $h \in A_\Delta$, ie the images of $\hat{\phi}$ and $\hat{\phi}'$ in $\text{Out}(A_\Delta)$ are equal. Moreover, if ϕ_1 and ϕ_2 are two elements of $U^0(A_\Gamma)$ and $\phi = \phi_1 \circ \phi_2$, then $\hat{\phi} = \hat{\phi}_1 \circ \hat{\phi}_2$ is a lift of ϕ which preserves A_Δ . Thus, $r_\Delta: U^0(A_\Gamma) \rightarrow \text{Out}(A_\Delta)$ is a homomorphism.

To see that this lands in $U^0(A_\Delta)$, we check that this is the case for the generators of $U^0(A_\Gamma)$. Let ϕ be a generator that lifts to a fold $\hat{\phi}: x \mapsto xy$, so $\text{lk}(x) \subset \text{lk}(y)$. If $x \notin \Delta$, then $\hat{\phi}$ restricts to the identity on A_Δ . If $x \in \Delta$, then U^0 -invariance of Δ implies that y is also in Δ , so the restriction is the lift of a fold in $U^0(A_\Delta)$. Now suppose $\hat{\phi}$ is a partial conjugation by x . If $x \notin \Delta$, then the fact that Δ is U^0 -invariant implies that Δ lies entirely in one component of $\Gamma \setminus \text{st}_\Gamma(x)$, so the action of $\hat{\phi}$ to A_Δ is trivial. If $x \in \Delta$, then $\text{st}_\Delta(x) \subseteq \text{st}_\Gamma(x)$, so the components of $\Delta \setminus \text{st}_\Delta(x)$ are contained in components of $\Gamma \setminus \text{st}_\Gamma(x)$. Thus, $\hat{\phi}$ restricts to a (product of) partial conjugation(s) by x on A_Δ . $\square$

Definition 4.7 If Δ is a U^0 -invariant subgraph and $f: H \rightarrow U^0(A_\Gamma)$ is any homomorphism, we call $f_\Delta = r_\Delta \circ f: H \rightarrow U^0(A_\Delta)$ the *restriction* of f to Δ .

5 U^0 -invariant subcomplexes of marked Γ -complexes

Throughout this section we assume $\Delta \subseteq \Gamma$ is a U^0 -invariant subgraph with $\text{lk}(\Delta) = \emptyset$ and (X, h) is a Γ -complex with an untwisted marking. This means that, for any blowup structure $\mathbb{S}_\Gamma^\Pi$ on X with associated collapse map $c_\pi: X \rightarrow \mathbb{S}_\Gamma$, the composition $c_\pi h^{-1}$ induces an untwisted automorphism of A_Γ , that is, an element of $U(A_\Gamma)$. The aim is to identify a subcomplex $X_\Delta \subset X$ which is invariant under the action of isometries that induce elements of $U^0(A_\Gamma)$. The reason for the restriction that $\text{lk}(\Delta) = \emptyset$ is that in this case the subcomplex X_Δ is unique. [Remark 5.8](#) discusses the general case.

The following lemma deals with the discrepancy between the subgroups $U(A_\Gamma)$ and $U^0(A_\Gamma)$.

Lemma 5.1 *If X is a Γ -complex with an untwisted marking h , then X has a blowup structure $\mathbb{S}_\Gamma^\Pi$ with collapse map $c_\pi: X = \mathbb{S}_\Gamma^\Pi \rightarrow \mathbb{S}_\Gamma$ such that $(c_\pi h^{-1})_* \in U^0(A_\Gamma)$.*

Proof $U^0(A_\Gamma)$ is normal in the untwisted subgroup $U(A_\Gamma)$ of $\text{Out}(A_\Gamma)$, and the quotient Q is a subgroup of $\text{Aut}(\Gamma)$, the group of graph automorphisms of Γ . The short exact sequence

$$1 \rightarrow U^0(A_\Gamma) \rightarrow U(A_\Gamma) \rightarrow Q \rightarrow 1$$

splits, so $U(A_\Gamma) = U^0(A_\Gamma) \rtimes Q$.

Let $\mathbb{S}_\Gamma^\Pi$ be any blowup structure on X such that $(c_\pi h^{-1})_*$ is untwisted. By the above observation, if $(c_\pi h^{-1})_*$ is not in $U^0(A_\Gamma)$, we can compose it with a graph automorphism α to produce an element of $U^0(A_\Gamma)$. Realize α by an isometry f_α of $\mathbb{S}_\Gamma$. Composing c_π with f_α is equivalent to changing the labels and orientations of the vertex-labeled hyperplanes in $\mathbb{S}_\Gamma^\Pi$ by α ; this changes the partitions by the same relabeling, giving a new set of partitions $\alpha\Pi$. In other words, this gives a blowup structure $\mathbb{S}_\Gamma^{\alpha\Pi}$ on X such that $(h^{-1}c_{\alpha\Pi})_* \in U^0(A_\Gamma)$. $\square$

Now fix a blowup structure $\mathbb{S}_\Gamma^\Pi$ on X such that $(c_\pi h^{-1})_* \in U^0(A_\Gamma)$. In any Γ -partition, all bases have the same link. Since Δ is U^0 -invariant, [Lemma 4.3](#) implies that either all bases of a partition in Π are in Δ , or none are. So we may write $\Pi = \{\mathfrak{D}_1, \dots, \mathfrak{D}_k, \mathcal{P}_1, \dots, \mathcal{P}_l\}$, where the $\mathfrak{D}_i$ are based in Δ and the $\mathcal{P}_j$ are based in $\Gamma \setminus \Delta$.

Lemma 5.2 *Each $\mathcal{P}_i$ has a unique side P_i^Δ such that $\Delta^\pm \subset P_i^\Delta \cup \text{lk}(\mathcal{P})$ and $\Delta^\pm \cap P_i^\Delta \neq \emptyset$. Furthermore, if $\mathcal{P}_i$ and $\mathcal{P}_j$ are not adjacent, then $P_i^\Delta \cap P_j^\Delta \neq \emptyset$.*

Proof Let $y_i \in \Gamma \setminus \Delta$ be a base for $\mathcal{P}_i$. Recall that we have assumed $\text{lk}(\Delta)$ is empty; this implies that $\Delta \setminus \text{st}(y_i)$ is nonempty. Since Δ is U^0 -invariant, it intersects at most one component of $\Gamma \setminus \text{st}(y_i)$. Each side of $\mathcal{P}_i$ is a union of components of $\Gamma \setminus \text{st}(y_j)$. Thus, there is a unique side P_i^Δ such that $\Delta^\pm \cap P_i^\Delta \neq \emptyset$ and $\Delta^\pm \subseteq P_i^\Delta \cup \text{lk}(\mathcal{P}_i)^\pm$.

If $\mathcal{P}_i$ and $\mathcal{P}_j$ are compatible but not adjacent and $P_i^\Delta \cap P_j^\Delta = \emptyset$, then $P_i^\Delta \cap \text{lk}(\mathcal{P}_j) = P_j^\Delta \cap \text{lk}(\mathcal{P}_i) = \emptyset$ as well (see [\[1, Lemma 2.9\]](#)), forcing $\Delta^\pm \subset \text{lk}(\mathcal{P}_i) \cap \text{lk}(\mathcal{P}_j)$. This contradicts our assumption that $\Delta^\pm \cap P_i^\Delta \neq \emptyset$. $\square$




Figure 3: Left: a graph Γ and U^0 -invariant subgraph Δ . Right: four different Γ -partitions with the same link and their restrictions to $\Delta^\pm$. In the restriction, the innermost and outermost partitions become trivial and the middle two become equal.

Recall from [4] that vertices of $\mathbb{S}_\Gamma^\Pi$ correspond to collections $\{Q_1^\times, \dots, Q_k^\times, P_1^\times, \dots, P_l^\times\}$, where the superscript $\times$ indicates a choice of side. Each pair of sides in the collection must be *consistent*, which means either they intersect nontrivially or their bases commute. We define $\mathbb{K}_\Delta$ to be the subcomplex of $\mathbb{S}_\Gamma^\Pi$ consisting of vertices of the form $\{Q_1^\times, \dots, Q_k^\times, P_1^\Delta, \dots, P_l^\Delta\}$, edges obtained by switching a side of some $\mathcal{Q}_i$ or labeled by some $v \in \Delta$, and all higher-dimensional cubes spanned by these edges.

Proposition 5.3 *The subcomplex $\mathbb{K}_\Delta$ is a subdivided blowup of $\mathbb{S}_\Delta$.*

Proof Let $\{\mathcal{Q}_1, \dots, \mathcal{Q}_k\}$ be the partitions of Π based in Δ , let Ω be the multiset of Δ -partitions obtained by intersecting the $\mathcal{Q}_i$ with $\Delta^\pm$ (see Figure 3), and let $\mathbb{S}_\Delta^\Omega$ be the corresponding subdivided blowup. If the bases of $\mathcal{Q}_i$ and $\mathcal{Q}_j$ don't commute, then exactly one pair of sides has empty intersection (see [4, Lemma 3.6]). The $\mathcal{Q}_i \cap \Delta^\pm$ are Δ -partitions, and the corresponding pair of sides still has empty intersection. This means that they are compatible, and all other pairs of sides must intersect. Since the $\mathcal{Q}_i$ are based in Δ , this means that $\{Q_1^\times, \dots, Q_k^\times\}$ is consistent in $\Gamma^\pm$ if and only if $\{Q_1^\times \cap \Delta^\pm, \dots, Q_k^\times \cap \Delta^\pm\}$ is consistent in $\Delta^\pm$, ie defines a vertex of $\mathbb{S}_\Delta^\Omega$. Note that, for each $j = 1, \dots, k$, either $\mathcal{P}_i$ is adjacent to $\mathcal{Q}_j$ or P_i^Δ intersects both sides of $\mathcal{Q}_j$ nontrivially, since $\mathcal{Q}_j$ splits its own base, which is in $\Delta \setminus \text{lk}(\mathcal{P}_i)$. Therefore, the map sending

$$\{Q_1^\times \cap \Delta^\pm, \dots, Q_k^\times \cap \Delta^\pm\} \rightarrow \{Q_1^\times, \dots, Q_k^\times, P_1^\Delta, \dots, P_l^\Delta\}$$

is well defined, and induces an isomorphism of $\mathbb{S}_\Delta^\Omega$ with $\mathbb{K}_\Delta$. □

A priori, our marking $h: X \rightarrow \mathbb{S}_\Gamma$ maps $\mathbb{K}_\Delta$ to $\mathbb{S}_\Gamma$, but we can adjust it within its homotopy class to map $\mathbb{K}_\Delta$ to $\mathbb{S}_\Delta$. This is because $h = u \circ c_\pi$, where $u: \mathbb{S}_\Gamma \rightarrow \mathbb{S}_\Gamma$ induces an element $\phi \in U^0(A_\Gamma)$. We can choose a representative $\hat{\phi} \in \text{Aut}(A_\Gamma)$ that sends A_Δ to itself, so, adjusting u by a homotopy, we get a map sending $\mathbb{S}_\Delta$ to itself. Since c_π sends $\mathbb{K}_\Delta$ to $\mathbb{S}_\Delta$, we may assume the composition $h = u \circ c_\pi$ restricts to a marking on $\mathbb{K}_\Delta$.

We constructed the subcomplex $\mathbb{K}_\Delta \subset X$ using the blowup structure that we chose on X . We now show that $\mathbb{K}_\Delta$ is essentially independent of this choice.

Proposition 5.4 *The subcomplex $\mathbb{K}_\Delta$ is independent of the blowup structure $\mathbb{S}_\Gamma^\Pi$ as long as this blowup structure satisfies $(c_\pi h^{-1})_* \in U^0(A_\Gamma)$.*

Proposition 5.4 can be proved combinatorially by keeping track of how the partitions change when we change the treelike set of hyperplanes or the labelings and orientations on the hyperplanes not in the treelike set. Being based in Δ turns out to be a property of hyperplanes in the treelike set, independent of the partitions used to describe them. The same is true for the property of being the Δ -side of a hyperplane based in $\Gamma \setminus \Delta$. Since these properties are what is used to define $\mathbb{K}_\Delta$, the subcomplex $\mathbb{K}_\Delta$ itself is independent of the blowup structure. Furthermore, an isometry $f: X \rightarrow X$ satisfying $(hf h^{-1})_* \in U^0(A_\Gamma)$ preserves these properties, so preserves $\mathbb{K}_\Delta$.

Below we give a different proof of **Proposition 5.4** in terms of the action of A_Γ on $\tilde{X}$ determined by h . This proof is more in the spirit of our previous paper [1] and more amenable to generalization. We continue to assume that $\text{lk}(\Delta) = \emptyset$.

The universal cover $\tilde{X}$ is a CAT(0) cube complex, and we will take advantage of the following facts about isometries of CAT(0) cube complexes. We say that a hyperbolic automorphism g of a CAT(0) cube complex *skewers* a hyperplane H if some axis for g crosses H transversely.

Lemma 5.5 *Let g be a hyperbolic automorphism of a CAT(0) cube complex $\tilde{X}$ and let H be a hyperplane in $\tilde{X}$. If some axis for g skewers H , then every axis for g skewers H . If no axis for g skewers H , then either all axes for g are on the same side of H , or H contains an axis for g .*

Proof Let α_1 and α_2 be two axes for g . Suppose α_1 crosses H but not α_2 does not. Let H^+ denote the half-space containing α_2 and H^- the complementary half-space. Set $\alpha_1^\pm = H^\pm \cap \alpha_1$. Let γ be a geodesic connecting a point in α_1^- to a point in α_2 . Any such path must cross H .

The action of g preserves both axes, so either g or g^{-1} maps α_1^+ into itself. Without loss of generality, assume $g(\alpha_1^+) \subset \alpha_1^+$ or, equivalently, $\alpha_1^- \subset g(\alpha_1^-)$. Consider the action on $\tilde{X}$ by positive powers of g . Since H does not intersect α_2 , the hyperplanes $g^k H$ also do not intersect α_2 . Thus, $g^k H$ separates α_1^- from α_2^+ for all $k > 0$. But this means that the path γ must cross infinitely many hyperplanes, which is impossible.

Now suppose α_1 and α_2 lie on opposite sides of H . The min set of g decomposes as an orthogonal product of an axis α and a convex subspace of $\tilde{X}$. It follows that α_1 and α_2 span a strip $\alpha \times I$ for some interval I . This strip intersects H in a convex set which separates these two axes. Any such set must contain $\alpha \times y$ for some point $y \in I$. $\square$

Proof of Proposition 5.4 Fix a basepoint at a vertex $x_0 \in X$, and let p_0 denote the base vertex of $\mathbb{S}_\Gamma$. We may assume the marking $h: X \rightarrow \mathbb{S}_\Gamma$ sends x_0 to p_0 , so induces an isomorphism $\pi_1(X, x_0) \cong \pi_1(\mathbb{S}_\Gamma, p_0) \equiv A_\Gamma$. The collapse map $c_\pi: X = \mathbb{S}_\Gamma^\Pi \rightarrow \mathbb{S}_\Gamma$ gives another marking, and by construction the composition $\phi = (h \circ c_\pi^{-1})_*: A_\Gamma \rightarrow A_\Gamma$ lies in $U^0(A_\Gamma)$ (as an automorphism). Since A_Δ is U^0 -invariant, ϕ sends A_Δ to a conjugate of itself. We may therefore homotope h , by dragging the basepoint p_0 around a loop in $\mathbb{S}_\Gamma$, so that ϕ sends A_Δ isomorphically to A_Δ . If we choose a lift $\tilde{x}_0$ of x_0 to the universal cover $\tilde{X}$, we obtain two actions of A_Γ on $\tilde{X}$, one from h and the other from c_π . Since ϕ preserves A_Δ , this means that the axes for elements of A_Δ under both actions coincide setwise.

Our goal will now be to characterize $\mathbb{K}_\Delta$ as subcomplex of X purely in terms of the U^0 -marking $h: X \rightarrow \mathbb{S}_\Gamma$, or equivalently, in terms of the action of A_Γ on $\tilde{X}$. The preceding paragraph implies that, if we can characterize $\mathbb{K}_\Delta$ in terms of the set of axes of elements of A_Δ , it does not matter whether we use the action from h or c_π .

Now let $\mathcal{H}$ be the set of hyperplanes in $\tilde{X}$ that are *not* skewered by any element of A_Δ . By Lemma 5.5, this is the same set whether we are considering the action defined by h or by c_π .

Claim *If $H \in \mathcal{H}$, then exactly one half-space $H^+ \subset \tilde{X}$ contains an axis for every element of A_Δ .*

Proof Let $\tilde{\mathbb{K}}_\Delta$ be the lift of $\mathbb{K}_\Delta$ preserved by the action of A_Δ . The lift $\tilde{\mathbb{K}}_\Delta$ is convex, and hence is connected and contains an axis for every element of A_Δ .

Let H be a hyperplane in $\tilde{X}$. We claim that $H \in \mathcal{H}$ if and only if $H \cap \tilde{\mathbb{K}}_\Delta = \emptyset$, so that all of $\tilde{\mathbb{K}}_\Delta$ is on the same side of H , and that side contains an axis for every $g \in A_\Delta$.

If $H \cap \tilde{\mathbb{K}}_\Delta \neq \emptyset$, then $\tilde{\mathbb{K}}_\Delta$ contains an edge dual to H . Lemma 3.10 of [1] implies that every edge in $\tilde{\mathbb{K}}_\Delta$ is in some axis for some $g \in A_\Delta$, so H is skewered by an element of A_Δ , ie H is not in $\mathcal{H}$. Conversely, if H is skewered by an element $g \in A_\Delta$, then, by Lemma 5.5, every axis for g skewers H , so in particular some axis contained in $\tilde{\mathbb{K}}_\Delta$ skewers H , so H intersects $\tilde{\mathbb{K}}_\Delta$.

If both half-spaces determined by H contain axes for every element of A_Δ , then, by Lemma 5.5, H itself contains an axis for every element of A_Δ . Let e_H be an edge dual to H . Then e_H is contained in an axis for some element $w \in A_\Gamma$, since that is true of every edge in $\tilde{X}$, and $w \notin A_\Delta$ since w skewers H . By [1, Lemma 3.10], all of H is contained in the min set for w so w commutes with every element of A_Δ , ie A_Δ has nontrivial centralizer. This contradicts the assumption that $\text{lk}(\Delta) = \emptyset$. $\triangleleft$

We now define

$$\tilde{X}_\Delta = \bigcap_{H \in \mathcal{H}} H_\Delta,$$

where $H_\Delta = H^+ \setminus \kappa(H)$ is the largest subcomplex of $\tilde{X}$ contained in H^+ . This is independent of the blowup structure, and coincides with $\tilde{\mathbb{K}}_\Delta$ for the action defined by c_π . Therefore, the image $\mathbb{K}_\Delta$ of $\tilde{X}_\Delta$ in X is independent of the blowup structure. $\square$

Remark 5.6 In the terminology of [2], the set of hyperplanes $\mathcal{H}$ occurring in the proof of Proposition 5.4 are exactly those which are *inessential* for the action of A_Δ . Indeed, by [2, Proposition 3.2(ii)], the essential core $\text{Ess}(\tilde{X}, A_\Delta)$ for the action of A_Δ on $\tilde{X}$ consists of those hyperplanes skewed by some axis of A_Δ , ie the complement of $\mathcal{H}$. Thus, the proposition asserts the existence of a convex subcomplex of $\tilde{X}$ whose hyperplanes extend exactly to $\text{Ess}(\tilde{X}, A_\Delta)$, and on which A_Δ acts cocompactly.

Notation Since Proposition 5.4 shows that $\mathbb{K}_\Delta$ is independent of the blowup structure, we emphasize this by using the notation X_Δ instead of $\mathbb{K}_\Delta$ for the image of $\tilde{X}_\Delta$ in X .

Corollary 5.7 Let Δ be a U^0 -invariant subgraph with $\text{lk}(\Delta) = \emptyset$, and $G < U^0(A_\Gamma)$ a finite subgroup which is realized on a Γ -complex X with an untwisted marking $h: X \rightarrow \mathbb{S}_\Gamma$. Then the restriction of G is realized on $(X_\Delta, h|_{X_\Delta})$.

Proof An element of G is realized by an isometry $f: X \rightarrow X$. An isometry of X sends any blowup structure to a new blowup structure. If the isometry induces an element of $U^0(A_\Gamma)$, we have shown that X_Δ has not changed, so f must send X_Δ to itself. $\square$

Remark 5.8 We constructed X_Δ assuming that $\text{lk}(\Delta) = \emptyset$. If this is not the case, we can look instead at $\text{st}(\Delta)$, which always has empty link. We will see in Section 7 that the complex for $\text{st}(\Delta) = \Delta * \text{lk}(\Delta)$ breaks into a product $X_\Delta \times X_{\text{lk}(\Delta)}$, where X_Δ and $X_{\text{lk}(\Delta)}$ are Δ - and $\text{lk}(\Delta)$ -complexes, respectively.

6 Extendable Δ -complexes

Let Δ be a U^0 -invariant subgraph of Γ , $\rho: G \rightarrow U^0(A_\Gamma)$ a homomorphism from a finite group G and (X, h) a marked Γ -complex realizing ρ . In the last section we found a G -invariant Δ -complex (possibly subdivided) sitting inside X . In this section we consider the opposite problem: given a marked Δ -complex Y realizing the restriction $\rho_\Delta: G \rightarrow U^0(A_\Gamma)$, when can Y sit equivariantly inside a marked Γ -complex? If there is such a marked Γ -complex, we say Y is *extendable*. To determine when Y is extendable, we first define what it means for a Δ -partition to be extendable.

Definition 6.1 Let $\Delta \subset \Gamma$ be a U^0 -invariant subgraph and let $\mathfrak{Q}$ be a Δ -partition. We say that $\mathfrak{Q}$ is *extendable* if there exists a Γ -partition $\hat{\mathfrak{Q}}$ such that $\hat{\mathfrak{Q}} \cap \Delta^\pm = \mathfrak{Q}$.

Proposition 6.2 Let Δ be a U^0 -invariant subgraph of Γ and $\mathfrak{Q}$ a Δ -partition. Then $\mathfrak{Q}$ is extendable if and only if there is some base m of $\mathfrak{Q}$ such that

- (1) $\text{lk}_\Gamma(v) \subseteq \text{lk}_\Gamma(m)$ for every v split by $\mathfrak{Q}$, and
- (2) if v_1 and v_2 are in the same component of $\Gamma \setminus \text{st}_\Gamma(m)$, then $v_1^\pm, v_2^\pm$ are all in the same side of $\mathfrak{Q}$.

Proof The “only if” direction is immediate, since any extension of $\mathfrak{Q}$ is a Γ -partition. Note that any extension of $\mathfrak{Q}$ also splits m . In fact, it has to be based at m since Δ is U^0 -invariant, which implies there is no $v \in \Gamma \setminus \Delta$ with $\text{lk}(v) \supseteq \text{lk}(m)$.

For the converse, suppose $\mathfrak{Q} = (\mathrm{lk}_\Delta^\pm(m) \mid Q_1 \mid Q_2)$ satisfies conditions (1) and (2). We build a Γ -partition $(\mathrm{lk}_\Gamma^\pm(m) \mid \hat{Q}_1 \mid \hat{Q}_2)$ based at m as follows.

If $v \in \mathrm{lk}_\Delta(m)$, then $v \in \mathrm{lk}_\Gamma(m)$. If v (resp. v^{-1}) is in Q_i , put v (resp. v^{-1}) in $\hat{Q}_i$. This determines where to place all $v^{\pm 1}$ for $v \in \Delta$.

Now suppose $v \in \Gamma \setminus (\Delta \cup \mathrm{st}_\Gamma(m))$ and let C be the component of $\Gamma \setminus \mathrm{st}_\Gamma(m)$ containing v . If $C \cap \Delta = \emptyset$, put all vertices of C and their inverses in the same side of $\hat{\mathfrak{Q}}$ (either side will do). If $C \cap \Delta$ is nonempty, then some side Q_i of $\mathfrak{Q}$ contains an element $w \in \Delta$. We must have $w^{-1} \in Q_i$ as well, since otherwise $\mathrm{lk}(w) \subseteq \mathrm{lk}(m)$ by condition (1), which would imply that w was the only vertex of C . By condition (2), the side Q_i is independent of the choice of w . Put all vertices of C and their inverses into $\hat{Q}_i$. $\square$

Definition 6.3 A blowup $\mathbb{S}_\Delta^\Omega$ is *extendable* if every $\mathfrak{Q} \in \Omega$ is extendable. Note we are not assuming the extended partitions are compatible. A Δ -complex is *extendable* if it can be given the structure of an extendable blowup.

Remark 6.4 It is not hard to show that, if a Δ -complex is extendable with respect to one blowup structure, then it is extendable with respect to any blowup structure, but we will not need this fact.

6.1 U^0 -invariant subgraphs and extendability

In this subsection we give a condition that guarantees that a Δ -complex realizing ρ_Δ is extendable.

Definition 6.5 A G -action on a Γ -complex X is *reduced* if no orbit of hyperplanes is contained in any treelike set. A marked Γ -complex (X, h) realizing $\rho: G \rightarrow \mathrm{Out}(A_\Gamma)$ is *reduced* if the associated G -action on X is reduced.

If a marked Γ -complex (X, h) realizing ρ is not reduced, then some orbit $G.H$ is acyclic since it is contained in the treelike (acyclic) set associated to some blowup structure on X . We can collapse every hyperplane in $G.H$ to produce a new marked Γ -complex. The following lemma guarantees that the new Γ -complex still realizes ρ .

Lemma 6.6 Let X be an NPC cube complex and let $G \leq \mathrm{Isom}(X)$ be a subgroup. Suppose $\mathcal{S}$ is a collection of hyperplanes that is acyclic and G -invariant. If $G \rightarrow \mathrm{Out}(\pi_1(X))$ is injective, then the collapse map $c: X \rightarrow X//\mathcal{S}$ induces an injection $G \rightarrow \mathrm{Isom}(X//\mathcal{S})$.

Proof Let $\pi = \pi_1(X)$. Since $\mathcal{S}$ is acyclic, c is a homotopy equivalence. The fact that G preserves $\mathcal{S}$ means there is an induced map $\bar{c}: G \rightarrow \mathrm{Isom}(X//\mathcal{S})$. We obtain a commutative diagram

$$(1) \quad \begin{array}{ccc} G & \longrightarrow & \mathrm{Out}(\pi) \\ \downarrow \bar{c} & & \downarrow \cong \\ \mathrm{Isom}(X//\mathcal{S}) & \longrightarrow & \mathrm{Out}(\pi) \end{array}$$

By assumption, $G \leq \mathrm{Isom}(X)$ injects into $\mathrm{Out}(\pi)$; hence, G also injects under $\bar{c}$. $\square$

Let (X, h) be a marked Γ -complex realizing ρ . By Lemma 6.6, we may continue to collapse G -orbits of hyperplanes until we obtain a reduced marked Γ -complex realizing ρ . Note that the result of this process is not unique, but depends on the set of orbits we choose to collapse.

In Section 5 we produced a subcomplex X_Δ of X for any U^0 -invariant subgraph Δ with empty link. If (X, h) is reduced, then in any blowup structure the orbit of any hyperplane labeled by a partition $\mathcal{P}$ contains a hyperplane labeled by an element of V . Since the action of G preserves the subcomplex X_Δ , the same must be true for orbits in X_Δ , so the action of G on X_Δ is also reduced. As an example, note that, if (X, h) is reduced and the restriction of G to $U^0(A_\Delta)$ is trivial, then the subcomplex X_Δ must be equal to the Salvetti complex $\mathbb{S}_\Delta$. This follows since no nontrivial blowup of $\mathbb{S}_\Delta$ is reduced with respect to the trivial action. Thus, reduced realizations of G may be thought of as “minimal” Γ -complexes realizing G .

Let Δ be a U^0 -invariant subgraph of Γ . The next proposition states that being reduced is sufficient to guarantee extendability for any Δ -complex realizing the restriction ρ_Δ .

Proposition 6.7 *Let (X_Δ, h_Δ) be a marked Δ -complex that realizes ρ_Δ . If the action of G on X_Δ is reduced, then X_Δ is extendable.*

Proof Let V be the set of vertices in Δ . Choose a blowup structure $X_\Delta \cong \mathbb{S}_\Delta^\Omega$. Since the G -action is reduced, for every $\mathfrak{Q} \in \Omega$ there exists $g \in G$ such that $g.H_\mathfrak{Q} = H_w$ for some $w \in V$. By Proposition 6.2, in order to verify extendability we must show:

- (i) There is some base m of $\mathfrak{Q}$ such that $\text{lk}_\Gamma(v) \subseteq \text{lk}_\Gamma(m)$ for every v split by $\mathfrak{Q}$.
- (ii) If v_1 and v_2 are in the same component of $\Gamma \setminus \text{st}_\Gamma(m)$, then $v_1^\pm, v_2^\pm$ are all in the same side of $\mathfrak{Q}$.

Proof of (i) Let χ be a characteristic cycle for v in $\mathbb{S}_\Delta^\Omega$. The image $g.\chi$ is a path that crosses each hyperplane of $\mathbb{S}_\Delta^\Omega$ at most once, so the edge labels that are not in Ω spell a cyclic word that is the image of v under g ,

$$(2) \quad g(v) = x_1^{\epsilon_1} \cdots x_k^{\epsilon_k},$$

where all the x_i are distinct and $\epsilon_i = \pm 1$.

Claim 6.8 *If $x_i \in V$ labels an edge in $g.\chi$, then $\text{lk}_\Gamma(v) \subseteq \text{lk}_\Gamma(x_i)$. If $\mathfrak{Q} \in \Omega$ labels an edge in $g.\chi$, then $\mathfrak{Q}$ splits some x_i .*

Proof The double link $\text{dlk}(v) = \text{lk}(\text{lk}(v))$ is U^0 -invariant by Proposition 4.4(1). Since $v \in \text{dlk}(v)$, this implies that each x_i appearing in (2) must be in $\text{dlk}(v)$, ie $\text{lk}_\Gamma(v) \subseteq \text{lk}_\Gamma(x_i)$. Now suppose $\mathfrak{Q}$ labels an edge of $g.\chi$ for some $\mathfrak{Q} \in \Omega$. We claim that $\mathfrak{Q}$ splits at least one of the x_i . Indeed, after choosing a basepoint $*$ in X_Δ , we see that $g.\chi$ is freely homotopic to a concatenation of edge paths $\eta_1\eta_2 \cdots \eta_k$, where η_i is an edge path based at $*$ representing the element x_i . Each η_i is freely homotopic to a characteristic cycle for x_i ; hence, η_i crosses the hyperplane $H_\mathfrak{Q}$ an odd number of times if and only if $\mathfrak{Q}$ splits x_i . Since $\eta_1 \cdots \eta_k$ is freely homotopic to $g.\chi$ and $g.\chi$ crosses $H_\mathfrak{Q}$ exactly once, we must have that $\mathfrak{Q}$ splits some x_i . $\triangleleft$

The proof of (i) now follows directly from the following claim.

Claim 6.9 *For every $\mathfrak{Q} \in \Omega$, there exists $m \in \text{split}(\mathfrak{Q})$ such that $\text{lk}_\Gamma(v) \subseteq \text{lk}_\Gamma(m)$ for every $v \in \text{split}(\mathfrak{Q})$. Moreover, defining $\text{lk}_\Gamma(\mathfrak{Q}) = \text{lk}_\Gamma(m)$, the action of G preserves $\text{lk}_\Gamma(E)$ for every edge label $E \in V \cup \Omega$.*

Proof Define an increasing filtration $\emptyset = V_0 \subsetneq V_1 \subsetneq \cdots \subsetneq V_r = V$, where, for $i \geq 1$, $V_i \setminus V_{i-1}$ consists of all $v \in \Delta$ such that $\text{lk}_\Gamma(v)$ is maximal among elements of $V \setminus V_{i-1}$. Every partition E in Ω splits some element of V , so we can extend this to an increasing filtration $\emptyset = \mathcal{F}_0 \subsetneq \mathcal{F}_1 \subsetneq \cdots \subsetneq \mathcal{F}_r = V \cup \Omega$, by letting $E \in \mathcal{F}_i$ if H_E splits some generator of V_i . We will prove by induction that:

- (a) For every $\mathfrak{Q} \in \mathcal{F}_i$, there exists $m \in \text{split}(\mathfrak{Q}) \cap V_i$ such that $\text{lk}_\Gamma(v) \subseteq \text{lk}_\Gamma(m)$ for every $v \in \text{split}(\mathfrak{Q})$.
- (b) Define $\text{lk}_\Gamma(\mathfrak{Q}) = \text{lk}_\Gamma(m)$; then, for any $A, B \in \mathcal{F}_i$, if $g.H_A = H_B$, then $\text{lk}_\Gamma(A) = \text{lk}_\Gamma(B)$.

The base case $\mathcal{F}_0 = \emptyset$ holds vacuously. Suppose by induction that, for some $i \geq 1$, we have verified (a) and (b) for $\mathcal{F}_{i-1}$. Since the G -orbit of every $H_{\mathfrak{Q}}$ contains H_w for some $w \in V$, it follows from (b) that $\mathcal{F}_{i-1}$ is the union of all G -orbits of elements of V_{i-1} . Consider now $\mathfrak{Q} \in \mathcal{F}_i \setminus \mathcal{F}_{i-1}$. Then any generator in $\text{split}(\mathfrak{Q})$ lies in $V \setminus V_{i-1}$, and $\mathfrak{Q}$ splits some $m \in V_i \setminus V_{i-1}$. Let $g \in G$ be such that $g.H_{\mathfrak{Q}} = H_w$ for some $w \in V$. For any $v \in \text{split}(\mathfrak{Q})$, we know that $\text{lk}_\Gamma(v) \subseteq \text{lk}_\Gamma(w)$ by Claim 6.8. In particular, $\text{lk}_\Gamma(m) \subseteq \text{lk}_\Gamma(w)$ and therefore $w \in V_i$. On the other hand, since $\mathcal{F}_{i-1}$ is a union of G -orbits and does not contain $\mathfrak{Q}$, we know that $w \notin V_{i-1}$. Hence, $w \in V_i \setminus V_{i-1}$ and therefore $\text{lk}_\Gamma(w) = \text{lk}_\Gamma(m)$ as m is maximal among all elements of $V \setminus V_{i-1}$. It follows that $\text{lk}_\Gamma(v) \subseteq \text{lk}_\Gamma(w) = \text{lk}_\Gamma(m)$ for any $v \in \text{split}(\mathfrak{Q})$, which proves (a).

Let χ be a characteristic cycle for $v \in V_i \setminus V_{i-1}$. Claim 6.8 implies that any hyperplane crossed by $g.\chi$ has a Γ -link that contains $\text{lk}_\Gamma(v)$, and hence its label is in $\mathcal{F}_i$. Since $\mathcal{F}_{i-1}$ is a union of G -orbits, if some label appearing in $g.\chi$ is not in $\mathcal{F}_{i-1}$, it must be in $\mathcal{F}_i \setminus \mathcal{F}_{i-1}$, and therefore is equivalent to v . In particular, if $g.H_v = H_A$, then $\text{lk}_\Gamma(A) = \text{lk}_\Gamma(v)$. Since every label in $\mathcal{F}_i \setminus \mathcal{F}_{i-1}$ appears in such an orbit, we conclude that G preserves the Γ -link of each element of $\mathcal{F}_i$, which proves (b) and completes the inductive step. $\triangleleft$

Proof of (ii) Suppose $\mathfrak{Q} \in \Omega$ is based at m and $v_1^\pm, v_2^\pm$ lie on opposite sides of $\mathfrak{Q}$. Since X_Δ is a blowup of $\mathbb{S}_\Delta$, we know that v_1 and v_2 lie in different components of $\Delta \setminus \text{st}_\Delta(m)$. We must show that they lie in different components of $\Gamma \setminus \text{st}_\Gamma(m)$.

As shown in [1], the inverse image of the (unique) vertex of $\mathbb{S}_\Delta$ under the collapse map $\mathbb{S}_\Delta^\Omega \rightarrow \mathbb{S}_\Delta$ is a CAT(0) subcomplex of $\mathbb{S}_\Delta^\Omega$, consisting of cubes whose edges are all labeled by partitions. Denote this subcomplex by $\mathbb{C}^\Delta$. Choose a minimal length edge path α in $\mathbb{C}^\Delta$ between any characteristic cycle for v_1 and any characteristic cycle for v_2 . Then α crosses exactly those hyperplanes labeled by partitions containing $v_1^\pm$ and $v_2^\pm$ on opposite sides. In particular, it crosses the hyperplane labeled $\mathfrak{Q}$. For $i = 1, 2$, let χ_i be a characteristic cycle for v_i starting at either end of α , and consider the edge path $\gamma = \chi_1 \alpha \chi_2 \bar{\alpha}$. Under the collapse map to $\mathbb{S}_\Delta$, the loop γ represents the element $v_1 v_2$. Observe that, by the minimality of α , the hyperplanes crossed by α , χ_1 and χ_2 are pairwise disjoint. Given an element $g \in G$, we have

$$g_*(\gamma) = g_*(\chi_1)g_*(\alpha)g_*(\chi_2)g_*(\bar{\alpha}) = g_*(\chi_1)g_*(\alpha)g_*(\chi_2)\overline{g_*(\alpha)}.$$

Since α , χ_1 and χ_2 cross each hyperplane of X_Δ at most once, and cross pairwise distinct sets of hyperplanes, the same is true of $g_*(\alpha)$, $g_*(\chi_1)$ and $g_*(\chi_2)$. Therefore, the hyperplanes crossed by $g_*(\gamma)$ that are not labeled by partitions define a (cyclic) word in the generators that is the image of $v_1 v_2$ under the action of g . Thus, we may write $g_*(v_1 v_2) = w_1 u w_2 u^{-1}$, where

$$u = y_1^{\delta_1} \cdots y_r^{\delta_r}$$

is a word in pairwise distinct generators y_j and $\delta_j = \pm 1$.

Since X_Δ is reduced, there exists some $g \in G$ such that g maps the hyperplane labeled $\mathcal{P}$ to a hyperplane labeled by one of the y_j , and, by [Claim 6.9](#), m and y_j belong to the same Γ -equivalence class. It thus suffices to prove that v_1 and v_2 lie in different components of $\Gamma \setminus \text{st}(y_j)$ for each j .

For any vertex v in Γ , $\text{dlk}(v)$ is U^0 -invariant, up to conjugacy. Thus, the cyclically reduced form of the word $g_*(v_1)$, namely w_1 , must be a word in $\text{dlk}(v_1)$, and similarly w_2 must be a word in $\text{dlk}(v_2)$. Choosing the representative of g_* in $\text{Aut}(A_\Gamma)$ to be one that takes v_1 to w_1 , we then have $g_*(v_2) = u w_2 u^{-1}$, where u is a product of generators y_i such that $\text{st}(y_i)$ separates some element of $\text{dlk}(v_1)$ from some element of $\text{dlk}(v_2)$. But, in this case, $\text{st}(y_i)$ must also separate v_1 from v_2 , as required. $\square$

7 Γ -complexes for joins and disjoint unions

In this section we describe procedures for constructing Γ -complexes realizing G when Γ is either a join or a disjoint union of U^0 -invariant subgraphs realizing the restriction of G . We begin with the case where $\Gamma = \Gamma_1 * \Gamma_2$ is a join, which is straightforward.

Proposition 7.1 *If $\Gamma = \Gamma_1 * \Gamma_2$, then any Γ -complex is a product of a Γ_1 -complex and a Γ_2 -complex. Conversely, any product of a Γ_1 -complex and a Γ_2 -complex is a $\Gamma_1 * \Gamma_2$ -complex.*

Proof Suppose $\mathbb{S}_\Gamma^\Pi$ is a blowup structure on a Γ -complex X . Write

$$\Pi = \{\mathcal{P}_1, \dots, \mathcal{P}_k, \mathcal{Q}_1, \dots, \mathcal{Q}_l\},$$

where $\mathcal{P}_i$ is based at $x_i \in \Gamma_1$ and $\mathcal{Q}_j$ is based at $y_j \in \Gamma_2$. (Note that, if one base is in Γ_i , then all bases are in Γ_i .) Since $\Gamma = \Gamma_1 * \Gamma_2$, every x_i is adjacent to every y_j , intersecting each $\mathcal{P}_i$ with $V^\pm(\Gamma_1)$ gives a Γ_1 -partition $\mathcal{P}_i^1$, and intersecting each $\mathcal{Q}_j$ with $V^\pm(\Gamma_2)$ gives a Γ_2 -partition $\mathcal{Q}_j^2$. Thus, Π is a compatible collection of Γ -partitions if and only if $\Pi_1 = \{\mathcal{P}_1^1, \dots, \mathcal{P}_k^1\}$ is a compatible collection of Γ_1 -partitions and $\Pi_2 = \{\mathcal{Q}_1^2, \dots, \mathcal{Q}_l^2\}$ is a compatible collection of Γ_2 -partitions. We conclude that X is a product of the Γ_i complexes with blowup structures $\mathbb{S}_{\Gamma_1}^{\Pi_1}$ and $\mathbb{S}_{\Gamma_2}^{\Pi_2}$, respectively. $\square$

Proposition 7.2 *Suppose $\Gamma = \Gamma_1 * \Gamma_2$ and let $\rho: G \rightarrow U^0(A_\Gamma)$ be a homomorphism. Then Γ_1 and Γ_2 are U^0 -invariant and, if (X_i, h_i) for $i = 1, 2$ are marked Γ_i -complexes that realize the restrictions ρ_i of ρ to Γ_i , their product $X_\Gamma = X_1 \times X_2$, equipped with the product action of G and the product marking $h = h_1 \times h_2$, is a marked Γ -complex realizing ρ . Moreover, the action of G on X_Γ given by ρ is reduced if and only if the actions on X_1 and X_2 are.*




Figure 4: Building a Γ -complex for a disjoint union $\Gamma = \Gamma_1 \sqcup \Gamma_2 \sqcup \Gamma_3 \sqcup \Lambda$.

Proof Observe that $\Gamma_1 = \text{lk}(\Gamma_2)$ and $\Gamma_2 = \text{lk}(\Gamma_1)$; hence, by Proposition 4.4, both are proper U^0 -invariant subgraphs. For $i = 1, 2$, suppose X_i is a Γ_i -complex and that $h_i : X_i \rightarrow \mathbb{S}_{\Gamma_i}$ is a marking which realizes the restriction ρ_i . Define $X_\Gamma = X_1 \times X_2$ and let G act via the product action. Blowup structures on X_1 and X_2 give a blowup structure to $X_1 \times X_2$ by Proposition 7.1, and the product marking $h = h_1 \times h_2 : X_\Gamma \rightarrow \mathbb{S}_\Gamma$ realizes the action of G . The final statement of the lemma follows from the fact that hyperplanes of X_Γ are all of the form $H_1 \times X_2$ for H_1 a hyperplane of X_1 or $X_1 \times H_2$ for H_2 a hyperplane of X_2 , and that the action of G preserves the product decomposition. $\square$

We next consider the case when Γ is a disjoint union of (not necessarily connected) subgraphs. Given a Γ_i -complex X_i for each subgraph Γ_i that is not a singleton, we construct a Γ -complex X that contains each X_i as a subcomplex.

Definition 7.3 (Γ -amalgam) Suppose Γ is a disjoint union $\Gamma_1 \sqcup \cdots \sqcup \Gamma_k \sqcup \Lambda$, where Λ is discrete. Let Z_Λ be a graph satisfying:

- The rank of Z_Λ is equal to $|\Lambda|$.
- k vertices of Z_Λ are labeled by $\{1, \dots, k\}$.
- Any unlabeled vertex of Z_Λ has valence at least 3.

For each $i \in \{1, \dots, k\}$, let X_i be a Γ_i -complex. Form a new cube complex Y_Γ , called a Γ -amalgam, as follows. For each vertex v of Z_Λ , set $X_v = X_i$ if v is labeled by i , and set X_v to be a point otherwise. Now construct a complex by starting from the disjoint union $\bigsqcup X_v$ and attaching an edge from X_v to X_w whenever $\{v, w\}$ is an edge in Z_Λ . When $X_v = X_i$, we may attach the edge anywhere. The resulting complex can be given the structure of a cube complex: if the endpoint of one of the added edges lies at a point p in the interior of a cube $C \subseteq X_i$, we perform the cubical subdivision of C at p . Define the resulting cube complex to be Y_Γ . (See Figure 4.)

Observe that Y_Γ contains a subdivided copy X'_i of X_i as a subcomplex for each i . Moreover, collapsing each of these subcomplexes separately to a point defines a natural map $Y_\Gamma \rightarrow Z_\Lambda$, which is a bijection away from the X_i . Since each X_i is a Γ_i -complex, it does not have any separating hyperplanes. It follows that a hyperplane of Y_Γ is separating if and only if it comes from a separating edge of Z_Λ .




Figure 5: Determining the Γ -partition associated to a hyperplane H in X_Γ . All hyperplanes in the treelike set for X'_1 other than H have been collapsed in this figure (see proof of [Proposition 7.4](#)).

Proposition 7.4 Let $\Gamma = \Gamma_1 \sqcup \cdots \sqcup \Gamma_k \sqcup \Lambda$, where Λ is discrete, and let Y_Γ be a Γ -amalgam formed from Γ_i -complexes X_i and a graph Z_Λ . Let X_Γ be the complex formed from Y_Γ by collapsing all separating edges from Z_Γ . Then X_Γ is a Γ -complex, ie there exists a collection of Γ -partitions Π such that $X_\Gamma \cong \mathbb{S}_\Gamma^\Pi$.

Proof First note that the hyperplanes of X_Γ consist of the hyperplanes of each X'_i (which remain in X'_i) and the midpoints of nonseparating edges of Z_Λ . Choose a subdivided blowup structure on each $X'_i \subset Y_\Gamma$, corresponding to a compatible multiset of partitions Π_i . Then choose a maximal tree T_Λ in Z_Λ and label and orient the edges not in T_Λ by the elements of Λ . Let $\mathcal{T}$ be the union of all hyperplanes with labels in the Π_i and those dual to edges in T_Λ . Collapsing all hyperplanes in $\mathcal{T}$ gives the Salvetti $\mathbb{S}_\Gamma$.

To see that $\mathcal{T}$ is the treelike set for a blowup structure, we need to check that each $H \in \mathcal{T}$ determines a Γ -partition. Cut all edges of Z_Λ that are labeled by elements of Λ , labeling the initial half-edge λ^{-1} and the terminal half-edge λ (see [Figure 5](#)). Each hyperplane in $\mathcal{T}$ now determines an evident partition of the vertices of $\Gamma^\pm$. First consider the partition associated to an edge e of T_Λ . Since we chose T_Λ after collapsing each X'_i to a point, no element of Γ_i is split by this partition. The fact that Z_Λ has no separating edges implies that the partition associated to e separates some $\lambda \in \Lambda$ from its inverse, so this gives a Γ -partition based at λ .

A hyperplane in X_i partitions the vertices of $\Gamma_i^\pm$, and the only new vertices of Γ that might be split are in Λ , so have empty links; hence, the resulting partition is still a valid Γ -partition with its original base.

We must also check that the duplicate partitions in the Π_i give distinct Γ -partitions, and the singleton partitions in the Π_i give rise to legitimate Γ -partitions. Singleton and duplicate partitions in Π_i result from attaching an edge e of Z_Λ to the middle of a cube, which we then subdivide by duplicating all hyperplanes that intersect the cube. The edge e lies on a characteristic cycle for some $w \in \Lambda$.

Let H be one of the hyperplanes that has been duplicated to form new hyperplanes H' and H'' , so that now e terminates at a point between H' and H'' . If H was labeled by a partition $\mathcal{P}$ in Π_i , the new partitions $\mathcal{P}'$ and $\mathcal{P}''$ corresponding to H' and H'' agree on $\Delta^\pm$, but either w or w^{-1} (depending on the orientation of the cycle) lies in opposite sides of the extensions of $\mathcal{P}'$ and $\mathcal{P}''$ to Γ -partitions. If H was labeled by a vertex v , let H' be the duplicate hyperplane that did *not* get the label v , so H' corresponds

to a singleton partition $\mathcal{S}$. Then v (or v^{-1}) and w (or w^{-1}) lie on the same side of the extension of $\mathcal{S}$ to $\Gamma^\pm$, so the corresponding Γ -partition is not a singleton. $\square$

In passing from Y_Γ to X_Γ in the proof of [Proposition 7.4](#), we had to collapse the set of separating edges in Z_Γ . Since the collection of all separating hyperplanes in an NPC cube complex X is acyclic and invariant under $\text{Aut}(X)$, [Lemma 6.6](#) implies that collapsing them has no effect on realizing actions of finite subgroups of $\text{Out}(A_\Gamma)$ by isometries.

Proposition 7.5 *Let $\Gamma = \Gamma_1 \sqcup \cdots \sqcup \Gamma_k \sqcup \Lambda$, where Λ is discrete, and let $\rho: G \rightarrow U^0(A_\Gamma)$ be a homomorphism. If (X_i, h_i) are marked Γ_i -complexes realizing the restriction of ρ to Γ_i , then there exists a marked Γ -complex (X_Γ, h) realizing ρ such that each X_i is a subcomplex of X_Γ and $h|_{X_i} = h_i$.*


Proof Let $\bar{A}_\Gamma$ and $\bar{A}_{\Gamma_i}$ be the finite extensions of A_Γ and A_{Γ_i} , respectively, determined by G . We apply [\[9, Proposition 7.5\]](#) to $\Gamma = \Gamma_1 \sqcup \cdots \sqcup \Gamma_k \sqcup \Lambda$, with marked complexes (X_i, h_i) as input. The result is a marked cube complex X realizing the action of G on A_Γ . In the construction (see [\[9, Proposition 3.1 and Theorem 4.1\]](#)), the marked complex (X_Γ, h) is formed from a graph-of-groups decomposition of $\bar{A}_\Gamma$. The edge stabilizers are all finite, vertex stabilizers are the corresponding finite extensions of the $\bar{A}_{\Gamma_i}$, and the rank of the underlying graph is $|\Lambda|$. Each vertex is then “blown up” to a (possibly subdivided) copy of X_i equipped with the marking h_i , to which edges are attached (see [\[9, Remark 7.7\]](#)), though we only subdivide where necessary, ie where an added edge meets the interior of a cube. In particular, collapsing each of the X_i separately to points yields a graph of rank $|\Lambda|$. Thus, the hypotheses of [Proposition 7.4](#) are satisfied, and, by [Lemma 6.6](#), the resulting marked Γ -amalgam (X_Γ, h) is a marked Γ -complex which realizes ρ . $\square$

Remark 7.6 In [Proposition 6.7](#), the hypothesis that action of G on X_Δ is reduced is essential. In [Proposition 7.5](#), given a graph $\Gamma = \Gamma_1 \sqcup \cdots \sqcup \Gamma_k \sqcup \Lambda$ and a collection of Γ_i -complexes X_i , we constructed a Γ -amalgam, X_Γ , whose restriction to each Γ_i is a subdivision of the given complexes X_i . While we may assume that the original complexes X_i are reduced, the resulting subdivisions need not be, since a subdivided e_v edge has one segment labeled e_v and the rest labeled by partitions. If X_Γ is not reduced, we can collapse an acyclic set of hyperplanes in X_Γ so that the resulting complex X'_Γ is reduced, and hence its restriction X'_i to each Γ_i -subcomplex is also reduced. We claim that X'_i is still a subdivision of the original X_i . To see this, note that, since the orbit of every edge e_φ in X_i contains an e_v edge, after subdividing, the orbit of at least one segment of this edge will also contain an e_v edge. Thus, this segment will not be collapsed in the reduction process.

8 Realizing finite subgroups of $U^0(A_\Gamma)$

In this section we prove our main theorem.

Theorem 8.1 *Let Γ be a simplicial graph, G a finite group and $\rho: G \rightarrow U^0(A_\Gamma)$ a homomorphism. Then there is a Γ -complex X with an untwisted marking $h: X \rightarrow \mathbb{S}_\Gamma$ on which ρ is realized by isometries.*

Figure 6: Subgraphs of Δ referred to in the proof of [Theorem 8.1](#).

We are especially interested in the case that ρ is an inclusion, but the proof is inductive and the general case is used in the induction. The proof borrows a number of ideas from [\[8\]](#).

Proof For any U^0 -invariant subgraph Δ , we will say a marked Δ -complex (X_Δ, h_Δ) “realizes G ” as shorthand for “realizes the restriction $\rho_\Delta = r_\Delta \circ \rho: G \rightarrow U^0(A_\Delta)$ ”. Here the target of h_Δ is the subcomplex $\mathbb{S}_\Delta \subset \mathbb{S}_\Gamma$.

We proceed by constructing marked Δ -complexes realizing G for larger and larger U^0 -invariant subgraphs Δ of Γ , until we have one for all of Γ . Specifically, we define the *chain length* $\ell = \ell(\Delta)$ to be the length of a maximal chain of U^0 -invariant subgraphs $\emptyset = \Gamma_0 \subsetneq \Gamma_1 \subsetneq \cdots \subsetneq \Gamma_\ell \subsetneq \Delta$ and proceed by induction on $\ell(\Delta)$. At each stage we ensure that the G -action is reduced, so that the complexes we construct are extendable.

If $\ell(\Delta) = 0$, [Proposition 4.5](#) says that Δ is discrete, ie A_Δ is a free group. The classical realization theorem for free groups [\[5; 10; 15\]](#) says we can find a marked graph (X_Δ, h_Δ) on which G is realized by isometries. After collapsing G -invariant forests, we may assume X_Δ is reduced (in particular has no separating edges), so (X_Δ, h_Δ) is the desired marked Δ -complex.

Now suppose $\ell(\Delta) = i \geq 1$. Note that any proper U^0 -invariant subgraph of Δ has chain length strictly smaller than i , so we can construct marked complexes realizing G for any such subgraph by induction. Fix a maximal U^0 -invariant chain $\emptyset = \Gamma_0 \subsetneq \Gamma_1 \subsetneq \cdots \subsetneq \Gamma_i \subsetneq \Delta$ and let $\Theta = \Delta \setminus \Gamma_i$. The next claim follows from maximality of the chain.

Claim If $w, w' \in \Theta$, then $\text{lk}(w) \cap \Gamma_i = \text{lk}(w') \cap \Gamma_i$.

Proof By [Proposition 4.4\(5\)](#), the subgraph spanned by Γ_i and all vertices adjacent to Γ_i is invariant, so, by maximality of the chain, either every element of Θ is adjacent to Γ_i , or none of them are. In the latter case, the claim is vacuously true because $\text{lk}(w) \cap \Gamma_i = \emptyset$ for every $w \in \Theta$. In the former case, choose $w \in \Theta$ such that $\text{lk}(w) \cap \Gamma_i$ is maximal, let $W = \{w' \in \Theta \mid \text{lk}(w') \cap \Gamma_i = \text{lk}(w) \cap \Gamma_i\}$ and let Δ' be the subgraph spanned by Γ_i and W .

We will now show that Δ' is U^0 -invariant, so, by maximality of the chain, $\Delta' = \Delta$ and Θ is the graph spanned by W . By [Lemma 4.3](#), we need to prove:

- (1) If $x \in \Delta'$ and $\text{lk}(x) \subset \text{lk}(y)$, then $y \in \Delta'$.
- (2) If Δ' intersects more than one component of $\Gamma \setminus \text{st}(y)$, then $y \in \Delta'$.

We know that Δ is U^0 -invariant, so, under the hypotheses of either (1) or (2), y must be in Δ . If $y \in \Gamma_i$, we are done, so assume instead that $y \in \Theta$. We need to show that $y \in W$. If $x \in \Delta'$ and $\text{lk}(x) \subset \text{lk}(y)$, the invariance of Γ_i , together with the assumption that $y \notin \Gamma_i$, guarantees that $x \notin \Gamma_i$. That is, $x \in W$, so, by maximality, y is also in W . In case (2), the invariance of Γ_i guarantees that $\text{st}(y)$ does not separate Γ_i , so either it separates Γ_i from W , or it separates two elements of W from each other. In either case, we must have $\text{lk}(w) \subseteq \text{st}(y)$, so $\text{lk}(w) \cap \Gamma_i \subseteq \text{lk}(y) \cap \Gamma_i$; hence, $y \in W$. $\triangleleft$

Now set $\Phi = \text{lk}(\Theta) \cap \Gamma_i$, and let (X_i, h_i) be a reduced marked Γ_i -complex realizing G . Note that Φ is U^0 -invariant by [Proposition 4.4](#).

If $\Phi = \Gamma_i$, then the above claim implies that Δ is the join $\Gamma_i * \Theta$, and that $\Theta = \text{lk}(\Gamma_i) \cap \Delta$, so is U^0 -invariant. By induction, we can find a reduced marked Θ -complex (X_Θ, h_Θ) which realizes the action of G , so, by [Proposition 7.2](#), the product $(X_i \times X_\Theta, h_i \times h_\Theta)$ is a Δ -complex realizing G .

If $\Phi = \emptyset$, then Δ is the disjoint union of Γ_i and Θ , so, by [Proposition 7.5](#), we can build a Δ -complex realizing G using (X_i, h_i) and complexes for the components of Θ that are not singletons (these components are U^0 -invariant by [Proposition 4.4\(4\)](#)).

If Φ is a proper subgraph of Γ_i , let $\Xi = \text{lk}(\Phi) \cap \Gamma_i$ (see [Figure 6](#)). Now

$$\Delta = \Gamma_i \cup \text{st}_\Delta(\Phi) = \Gamma_i \cup (\Phi * (\Xi \sqcup \Theta))$$

and $\Gamma_i \cap \text{st}_\Delta(\Phi) = \Phi * \Xi$. Both Φ and Ξ are U^0 -invariant, so $\Phi * \Xi$ is also U^0 -invariant by [Proposition 4.4\(3\)](#). Since $\text{lk}_{\Gamma_i}(\Phi * \Xi) = \emptyset$, X_i contains a unique invariant (possibly subdivided) $(\Phi * \Xi)$ -complex $X_{\Phi * \Xi} = X_\Phi \times X_\Xi$ realizing G , by [Corollary 5.7](#). We may choose h_i so that it restricts to a marking $h_{\Phi * \Xi} = h_\Phi \times h_\Xi$ on $X_{\Phi * \Xi}$.

The subgraph $\text{st}_\Delta(\Phi)$ is also U^0 -invariant, and we build a $\text{st}_\Delta(\Phi)$ -complex $X_{\text{st}(\Phi)}$ realizing G as follows. We first build a complex $X_{\text{lk}(\Phi)}$ for $\text{lk}_\Delta(\Phi) = \Xi \sqcup \Theta$ using a copy of the complex X_Ξ we already found in X_i and complexes for the components of Θ that are not singletons, as we did in the case that $\Phi = \emptyset$ above. After reducing, [Proposition 6.7](#) ensures that $X_{\text{lk}(\Phi)}$ is extendable. By [Remark 7.6](#), the reduced complex still contains a subdivided copy of X_Ξ . We then take the product of $X_{\text{lk}(\Phi)}$ with a copy of the complex $X_\Phi \subset X_i$ to obtain a complex for $X_{\text{st}(\Phi)}$ realizing G with respect to the product marking.

If $\Xi \neq \emptyset$, then $\text{lk}_{\text{st}(\Phi)}(\Phi * \Xi) = \emptyset$, so the complex $X_{\text{st}(\Phi)}$ that we just built contains a unique (possibly subdivided) $(\Phi * \Xi)$ -complex realizing G . By construction, this is identical to the complex $X_{\Phi * \Xi}$ contained

in X_i , so we may glue X_i to $X_{\text{st}(\Phi)}$ by identifying these subcomplexes, thus forming a new complex X_Δ . The markings h_i and $h_{\text{st}(\Phi)}$ agree on $X_{\Phi * \Xi}$, so we obtain a marked complex (X_Δ, h_Δ) realizing G .

If $\Xi = \emptyset$, then $\text{lk}_{\text{st}(\Phi)}(\Phi) = \Theta$, so $X_{\text{st}(\Phi)} = X_\Phi \times X_\Theta$. The following claim will allow us to pick out a particular slice $X_\Phi \times \{p\} \subset X_\Phi \times X_\Theta$ to glue to the (unique) copy of X_Φ contained in X_i .

Claim *The restriction $\rho_\Theta: G \rightarrow U^0(A_\Theta)$ lifts to a homomorphism $G \rightarrow \text{Aut}(A_\Theta)$.*

Proof We use [3, Lemma 2.2], which says:

- (*) If Σ is a subgraph of Δ , then the normalizer in A_Δ of A_Σ is $A_{\text{st}_\Delta(\Sigma)} = A_\Sigma \times A_{\text{lk}_\Delta(\Sigma)}$.
- (**) If $\Sigma_1, \Sigma_2 \subset \Delta$, then $x A_{\Sigma_2} x^{-1} \leq A_{\Sigma_1}$ if and only if $\Sigma_2 \subset \Sigma_1$ and $x = x_1 x_2$ with $x_1 \in N_{A_\Delta}(A_{\Sigma_1})$ and $x_2 \in N_{A_\Delta}(A_{\Sigma_2})$.

Since all normalizers will be taken with respect to A_Δ , for the rest of this proof we omit A_Δ from the notation for normalizers.

We are assuming $\Xi = \emptyset$, so $\Phi = \text{lk}_\Delta(\Theta)$ and $\Theta = \text{lk}_\Delta(\Phi)$. Then (*) says $N(A_\Phi) = N(A_\Theta) = N(A_{\Phi * \Theta}) = A_{\Phi * \Theta}$. Furthermore, if $\text{lk}_\Delta(\Gamma_i) \neq \emptyset$, then $\Gamma_i = \Phi$ and $\Delta = \Phi * \Theta$. We have already taken care of this case, so we may assume that $\text{lk}_\Delta(\Gamma_i) = \emptyset$ and we have $N(A_{\Gamma_i}) = A_{\Gamma_i}$.

Let $g \in G$. Since Γ_i, Θ, Φ and $\Phi * \Theta$ are all U^0 -invariant, the corresponding special subgroups of A_Γ are all sent to conjugates of themselves by any representative of g in $\text{Aut}(A_\Gamma)$. Pick a representative $\hat{g}$ that sends A_{Γ_i} to itself, and suppose $\hat{g}(A_\Phi) = x A_\Phi x^{-1}$. Since $\hat{g}(A_\Phi) \leq \hat{g}(A_{\Gamma_i}) = A_{\Gamma_i}$, (**) says that $x = x_1 x_2$ with $x_1 \in N(A_{\Gamma_i}) = A_{\Gamma_i}$ and $x_2 \in N(A_\Phi)$. Then $x A_\Phi x^{-1} = x_1 x_2 A_\Phi x_2^{-1} x_1^{-1} = x_1 A_\Phi x_1^{-1}$, so, after composing $\hat{g}$ with conjugation by x_1 , we may assume $\hat{g}$ sends both A_{Γ_i} and A_Φ to themselves.

Now $\Phi \subset \Phi * \Theta$, so $\hat{g}(A_\Phi) \leq \hat{g}(A_{\Phi * \Theta}) = y A_{\Phi * \Theta} y^{-1}$ for some y . Since $\hat{g}(A_\Phi) = A_\Phi$, this gives $y^{-1} A_\Phi y \leq A_{\Phi * \Theta}$, so, by (**), $y \in N(A_{\Phi * \Theta}) N(A_\Phi) = A_{\Phi * \Theta}$, ie $\hat{g}(A_{\Phi * \Theta}) = y A_{\Phi * \Theta} y^{-1} = A_{\Phi * \Theta}$. Finally, $\Theta \subset \Phi * \Theta$, so $\hat{g}(\Theta) = z A_\Theta z^{-1} \leq \hat{g}(\Phi * \Theta) = \Phi * \Theta$ for some z , so (**) gives that $z \in N(A_{\Phi * \Theta}) N(A_\Theta) = A_{\Phi * \Theta} = N(A_\Theta)$, so $\hat{g}(A_\Theta) = A_\Theta$ as well.

Now let $g_1, g_2 \in G$ with $g_1 g_2 = g_3$, and find representatives $\hat{g}_1, \hat{g}_2$ and $\hat{g}_3$ as above. We know that $\hat{g}_1 \hat{g}_2 \hat{g}_3^{-1}$ is inner and preserves A_{Γ_i} and A_Θ , so the conjugating element lives in $N(A_{\Gamma_i}) \cap N(A_\Theta) = A_{\Gamma_i} \cap A_{\Phi * \Theta} = A_\Phi$. But conjugation by an element of A_Φ is trivial on A_Θ , ie the restriction of $\hat{g}_1 \hat{g}_2 \hat{g}_3^{-1}$ to A_Θ is the identity. Thus, $g \mapsto \hat{g}$ gives a lift of G to $\text{Aut}(A_\Theta)$. $\triangleleft$

By the claim, the action of G on X_Θ lifts to an action on $\tilde{X}_\Theta$. Since $\tilde{X}_\Theta$ is CAT(0) this action has a fixed point; projecting this fixed point back down gives a fixed point $p \in X_\Theta$. We now build X_Δ by gluing X_i to $X_{\Phi * \Theta} = X_\Phi \times X_\Theta$ along their common subspace $X_\Phi = X_\Phi \times \{p\}$. As above, the markings h_i and $h_{\Phi * \Theta}$ agree on the overlap, so give a marking $h_\Delta: X_\Delta \rightarrow \mathbb{S}_\Delta$ realizing G .

It remains to check that the complexes X_Δ that we have just built are actually Δ -complexes. We start by choosing a blowup structure $X_i \cong \mathbb{S}_{\Gamma_i}^{\Omega_i}$. Since the action of G on X_i is reduced, this blowup structure is extendable, and induces (possibly subdivided) extendable blowup structures on the subcomplexes X_Φ , X_Ξ and $X_\Phi \times X_\Xi$.

If $\Xi = \emptyset$, then finding a blowup structure is slightly easier, so we do that case first. In this case, Θ is U^0 -invariant, and we can choose an extendable blowup structure $\mathbb{S}_{\Theta}^{\Omega_\Theta}$ on X_Θ . Recall that we may have needed to subdivide X_Θ in order to make the fixed point a vertex; this means that Ω_Θ may contain trivial or duplicate partitions. The fixed point $p \in X_\Theta$ now lies in a *region*, ie a consistent choice of sides for each element of Ω_Θ . The structure on X_Θ together with the blowup structure on X_Φ now give a blowup structure on $X_\Phi \times X_\Theta$, by [Proposition 7.1](#). The partitions in Ω_i and Ω_Θ are all extendable, so in particular can be extended to Δ . We need to find extensions that form a compatible collection of Δ -partitions.

If $\mathcal{P} \in \Omega_i$, we need to decide where to put the vertices $v^\pm \in \Theta^\pm$ in our extension $\hat{\mathcal{P}}$. If $\mathcal{P}$ is based at $m \in \Phi$, they must all go into $\text{lk}(\hat{\mathcal{P}})$. If m is distance at least 2 from Θ , there is some $u \in \Phi$ with $u \notin \text{lk}(m)$ (since $\Xi = \emptyset$), so all vertices of Θ are in the same component of $\Delta \setminus \text{st}(m)$ as u , so all of $\Theta^\pm$ must go into the same side of $\mathcal{P}$ as u . (Note that the extendability of $\mathcal{P}$ guarantees that, since any two choices for u lie in the same component of $\Delta \setminus \text{st}(m)$, they must lie on the same side of $\mathcal{P}$, so there is no ambiguity here.) We also need to extend the partitions $\mathcal{Q} \in \Omega_\Theta$ to $\Delta^\pm$. All vertices in Φ must go into the link of each extension. Since Γ_i is U^0 -invariant, no vertex in Θ has a star that separates Γ_i , so vertices of $\Gamma_i \setminus \Phi$ and their inverses all have to go in the same side of $\mathcal{Q}$ for each $\mathcal{Q} \in \Omega_\Theta$. We put them all into the region determined by the fixed point p . It is now routine to check that all the extensions $\hat{\mathcal{Q}}$ and $\hat{\mathcal{P}}$ we have constructed are compatible. This verifies that X_Δ is a Δ -complex in the case $\Xi = \emptyset$. If the G -action on X_Δ is not reduced, we can reduce it to obtain an extendable Δ -complex.

If $\Xi \neq \emptyset$, we can find a blowup structure $X_{\text{lk}(\Phi)} \cong \mathbb{S}_{\text{lk}(\Phi)}^{\Omega_{\text{lk}(\Phi)}}$ that restricts to the given blowup structure on X_Ξ by [Proposition 7.4](#).

The procedure we used in the case $\Xi = \emptyset$ to extend $\mathcal{P} = (P_1 | P_2 | \text{lk}(\mathcal{P})) \in \Omega_i$ to $\Delta^\pm$ works again unless $\mathcal{P}$ is based at $m \in \Xi$. In this case, each $\mathcal{P} \cap \Xi^\pm$ is also the restriction of a partition $\mathcal{Q} = (Q_1 | Q_2 | \text{lk}(\mathcal{Q})) \in \Omega_{\text{st}(\Phi)}$. We form $\hat{\mathcal{P}}$ by adding $Q_i \cap \Theta^\pm$ to P_i for $i = 1, 2$.

We also need to extend partitions $\mathcal{Q} \in \Omega_{\text{st}(\Phi)}$ that are based at $m \in \Theta$ to Δ -partitions $\hat{\mathcal{Q}}$. The star $\text{st}(m)$ cannot disconnect Ξ or Γ_i since both are U^0 -invariant, so we add all of $\Gamma_i \setminus \Phi$ to the same side of $\hat{\mathcal{Q}}$ as Ξ . The extensions $\hat{\mathcal{P}}$ and $\hat{\mathcal{Q}}$ are now a compatible collection of Δ -partitions, giving X_Δ a blowup structure. Reducing X_Δ if necessary, this completes the induction and concludes the proof of the theorem. $\square$

Recall from [\[4\]](#) that K_Γ is a contractible subspace of $\mathbb{C}_\Gamma$ which is invariant under the action of $U^0(A_\Gamma)$. Points of K_Γ are Γ -complexes with untwisted markings, so, if ρ is an inclusion, [Theorem 8.1](#) gives the following statement.

Corollary 8.2 *The action of any finite subgroup of $U^0(A_\Gamma)$ on K_Γ has a fixed point.*

Finally, we obtain the following information about finite subgroups of $U^0(A_\Gamma)$.

Corollary 8.3 *The group $U^0(A_\Gamma)$ contains only finitely many conjugacy classes of finite subgroups.*

Proof By [Theorem 8.1](#), every finite subgroup of $U^0(A_\Gamma)$ is realized on a Γ -complex. Changing the marking produces a conjugate subgroup, so we may ignore the markings. There are only finitely many Γ -complexes since there are only finitely many partitions of $\Gamma^\pm$, so only finitely many compatible collections of Γ -partitions, each of which determines a unique blowup of $\mathbb{S}_\Gamma$. Finally, there are only finitely many cubical isometries of a finite cube complex. $\square$

References

- [1] **C Bregman, R Charney, K Vogtmann**, *Outer space for RAAGs*, *Duke Math. J.* 172 (2023) 1033–1108 [MR](#) [Zbl](#)
- [2] **P-E Caprace, M Sageev**, *Rank rigidity for CAT(0) cube complexes*, *Geom. Funct. Anal.* 21 (2011) 851–891 [MR](#) [Zbl](#)
- [3] **R Charney, J Crisp, K Vogtmann**, *Automorphisms of 2-dimensional right-angled Artin groups*, *Geom. Topol.* 11 (2007) 2227–2264 [MR](#) [Zbl](#)
- [4] **R Charney, N Stambaugh, K Vogtmann**, *Outer space for untwisted automorphisms of right-angled Artin groups*, *Geom. Topol.* 21 (2017) 1131–1178 [MR](#) [Zbl](#)
- [5] **M Culler**, *Finite groups of outer automorphisms of a free group*, from “Contributions to group theory”, *Contemp. Math.* 33, Amer. Math. Soc., Providence, RI (1984) 197–207 [MR](#) [Zbl](#)
- [6] **C W Curtis, I Reiner**, *Representation theory of finite groups and associative algebras*, *Pure Appl. Math.* 11, Interscience, New York (1962) [MR](#) [Zbl](#)
- [7] **F Haglund, D T Wise**, *Special cube complexes*, *Geom. Funct. Anal.* 17 (2008) 1551–1620 [MR](#) [Zbl](#)
- [8] **S Hensel, D Kielak**, *Nielsen realisation for untwisted automorphisms of right-angled Artin groups*, *Proc. Lond. Math. Soc.* 117 (2018) 901–950 [MR](#) [Zbl](#)
- [9] **S Hensel, D Kielak**, *Nielsen realization by gluing: limit groups and free products*, *Michigan Math. J.* 67 (2018) 199–223 [MR](#) [Zbl](#)
- [10] **D G Khramtsov**, *Finite groups of automorphisms of free groups*, *Mat. Zametki* 38 (1985) 386–392 [MR](#) [Zbl](#) In Russian; translated in *Math. Notes* 38 (1985) 721–724
- [11] **MR Laurence**, *A generating set for the automorphism group of a graph group*, *J. Lond. Math. Soc.* 52 (1995) 318–334 [MR](#) [Zbl](#)
- [12] **G Levitt, J-L Nicolas**, *On the maximum order of torsion elements in $GL(n, \mathbb{Z})$ and $\text{Aut}(F_n)$* , *J. Algebra* 208 (1998) 630–642 [MR](#) [Zbl](#)
- [13] **H Servatius**, *Automorphisms of graph groups*, *J. Algebra* 126 (1989) 34–60 [MR](#) [Zbl](#)
- [14] **J Smillie, K Vogtmann**, *Automorphisms of graphs, p -subgroups of $\text{Out}(F_n)$ and the Euler characteristic of $\text{Out}(F_n)$* , *J. Pure Appl. Algebra* 49 (1987) 187–200 [MR](#) [Zbl](#)
- [15] **B Zimmermann**, *Über Homöomorphismen n -dimensionaler Henkelkörper und endliche Erweiterungen von Schottky-Gruppen*, *Comment. Math. Helv.* 56 (1981) 474–486 [MR](#) [Zbl](#)

*Department of Mathematics, Tufts University
Medford, MA, United States*

*Department of Mathematics, Brandeis University
Waltham, MA, United States*

*Mathematics Institute, University of Warwick
Coventry, United Kingdom*

corey.bregman@tufts.edu, charney@brandeis.edu, k.vogtmann@warwick.ac.uk

Received: 15 October 2023 Revised: 14 February 2024

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Kathryn Hess
kathryn.hess@epfl.ch
École Polytechnique Fédérale de Lausanne

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Thomas Koberda	University of Virginia thomas.koberda@virginia.edu
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Markus Land	LMU München markus.land@math.lmu.de
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Norihiko Minami	Yamato University minami.norihiko@yamato-u.ac.jp
Octav Cornea	Université de Montréal cornea@dms.umontreal.ca	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Robert Oliver	Université Paris 13 bobol@math.univ-paris13.fr
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
David Futer	Temple University dfuter@temple.edu	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Vesna Stojanoska	Univ. of Illinois at Urbana-Champaign vesna@illinois.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Kristen Hendricks	Rutgers University kristen.hendricks@rutgers.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Daniel Isaksen	Wayne State University isaksen@math.wayne.edu	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2025 is US \$760/year for the electronic version, and \$1110/year (+\$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by [Mathematical Reviews](#), [Zentralblatt MATH](#), [Current Mathematical Publications](#) and the [Science Citation Index](#).

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

 **mathematical sciences publishers**

nonprofit scientific publishing

<https://msp.org/>

© 2025 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 25

Issue 4 (pages 1917–2526)

2025

The zero stability for the one-row colored $s\ell_3$ -Jones polynomial	1917
WATARU YUASA	
Quillen homology of spectral Lie algebras with application to mod p homology of labeled configuration spaces	1945
ADELA YIYU ZHANG	
Coarse Alexander duality for pairs and applications	1999
G CHRISTOPHER HRUSKA, EMILY STARK and HÙNG CÔNG TRẦN	
K -cwaist on complete foliated manifolds	2037
GUANGXIANG SU and XIANGSHENG WANG	
Line bundle twists for unitary bordism are ghosts	2053
THORSTEN HERTL	
The generalized Kauffman–Harary conjecture is true	2067
RHEA PALAK BAKSHI, HUIZHENG GUO, GABRIEL MONTOYA-VEGA, SUJOY MUKHERJEE and JÓZEF H PRZYTICKI	
Rigidity of elliptic genera for nonspin manifolds	2083
MICHAEL WIEMELER	
Shrinking without doing much at all	2099
MICHAEL FREEDMAN and MICHAEL STARBIRD	
Action of the Mazur pattern up to topological concordance	2115
ALEX MANCHESTER	
Kauffman bracket intertwiners and the volume conjecture	2143
ZHIHAO WANG	
Horizontal decompositions, II	2179
PAOLO LISCA and ANDREA PARMA	
On the nonorientable four-ball genus of torus knots	2209
FRASER BINNS, SUNGKYUNG KANG, JONATHAN SIMONE and PAULA TRUÖL	
Generalised Baumslag–Solitar groups and hierarchically hyperbolic groups	2253
JACK O BUTTON	
Geometric and arithmetic properties of Löbell polyhedra	2281
NIKOLAY BOGACHEV and SAMI DOUBA	
Formality of sphere bundles	2297
JIAWEI ZHOU	
A Quillen stability criterion for bounded cohomology	2317
CARLOS DE LA CRUZ MENGUAL and TOBIAS HARTNICK	
T -equivariant motives of flag varieties	2343
CAN YAYLALI	
Small Heegaard genus and $SU(2)$	2369
JOHN A BALDWIN and STEVEN SIVEK	
Harmonic measures and rigidity for surface group actions on the circle	2391
MASANORI ADACHI, YOSHIFUMI MATSUDA and HIRAKU NOZAWA	
Finite groups of untwisted outer automorphisms of RAAGs	2413
COREY BREGMAN, RUTH CHARNEY and KAREN VOGTMANN	
Computations on cobordism groups of projected immersions	2441
ANDRÁS CSÉPAI	
Rank-preserving additions for topological vector bundles, after a construction of Horrocks	2451
MORGAN P OPIE	
Power sum elements in the G_2 skein algebra	2477