Math. Proc. Camb. Phil. Soc. (2021), **170**, 523–547 © *Cambridge Philosophical Society* 2019
doi:10.1017/S0305004119000501
First published online 20 November 2019

Cube complexes and abelian subgroups of automorphism groups of RAAGs

BY BENJAMIN MILLARD

Department of Mathematics, University College London, 25 Gordon Street, London WC1H 0AY, U.K.
e-mail: benjamin.millard.17@ucl.ac.uk

AND KAREN VOGTMANN

Mathematics Institute, Zeeman Building, University of Warwick, Coventry CV4 7AL, U.K.
e-mail: kvogtmann@gmail.com

(*Received* 10 *April* 2019; *revised* 17 *October* 2019)

Abstract

We construct free abelian subgroups of the group $U(A_\Gamma)$ of untwisted outer automorphisms of a right-angled Artin group, thus giving lower bounds on the virtual cohomological dimension. The group $U(A_\Gamma)$ was studied in [5] by constructing a contractible cube complex on which it acts properly and cocompactly, giving an upper bound for the virtual cohomological dimension. The ranks of our free abelian subgroups are equal to the dimensions of *principal cubes* in this complex. These are often of maximal dimension, so that the upper and lower bounds agree. In many cases when the principal cubes are not of maximal dimension we show there is an invariant contractible subcomplex of strictly lower dimension.

2010 Mathematics Subject Classification: 20F65 (Primary); 20F28, 20F36 (Secondary)

1. *Introduction*

The class of right-angled Artin groups (commonly called RAAGs) contains the familiar examples of finitely generated free groups and free abelian groups. Though uncomplicated themselves, both examples have complex and interesting automorphism groups. In recent years these automorphism groups have been shown to share many properties, but also to differ in significant ways (see e.g. the survey articles [2, 16]). In this paper we study automorphism groups of general RAAGs, concentrating on the aspects they share with automorphism groups of free groups. These aspects are largely captured by the subgroup of untwisted automorphisms, as previously studied in [5]. Let us recall the definition.

A general RAAG is conveniently described by drawing a finite simplicial graph Γ. The RAAG is then the group A_Γ generated by the vertices of Γ, with defining relations that two generators commute if and only if the corresponding vertices are connected by an edge of Γ. By theorems of Laurence [13] and Servatius [14], the automorphism group of A_Γ is generated by inversions of the generators, graph automorphisms, admissible transvections (multiplying one generator by another) and admissible partial conjugations (conjugating

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

some subset of generators by another generator). Here transvections and partial conjugations are *admissible* if they respect the commutation relations. A transvection is called a *twist* if the generators involved commute. The subgroup of $Out(A_\Gamma)$ generated by twists injects into a parabolic subgroup of $SL(n, \mathbb{Z})$, where n is the number of vertices of Γ, and is well understood. The subgroup generated by all generators other than twists is the *untwisted subgroup* $U(A_\Gamma)$. This subgroup captures the part of $Out(A_\Gamma)$ most closely related to $Out(F_n)$. For example, if $A_\Gamma = F_n$ then $U(A_\Gamma) = Out(F_n)$, and $U(A_\Gamma)$ always contains the kernel of the map $Out(A_\Gamma) \to GL(n, \mathbb{Z})$ induced by abelianisation $A_\Gamma \to \mathbb{Z}^n$.

For free groups, the virtual cohomological dimension (VCD) of $Out(F_n)$ is equal to the maximal rank of a free abelian subgroup. The lower bound is established by exhibiting an explicit free abelian subgroup. For the upper bound, one considers the action of $Out(F_n)$ on a contractible space $\mathcal{O}_n$ known as *Outer space*. This action is proper, and $\mathcal{O}_n$ contains an equivariant deformation retract K_n known as the *spine* of Outer space, whose dimension is equal to the lower bound (see [**8**]).

For the subgroup $U(A_\Gamma)$ associated to a general RAAG, an analogous outer space $\mathcal{O}_\Gamma$ and spine K_Γ were defined in [**5**]. The dimension of K_Γ gives an obvious upper bound on the VCD of $U(A_\Gamma)$. Lower bounds were obtained in [**3**] by exhibiting free abelian subgroups ([**3**] actually exhibited free abelian subgroups in the entire group $Out(A_\Gamma)$, but these contain identifiable subgroups of $U(A_\Gamma)$). However, there was no clear relationship between the rank of these subgroups and the dimension of K_Γ, and there was often a large gap between the upper bound and lower bounds.

In this paper we address this problem. The spine K_Γ has the structure of a cube complex, and we produce free abelian subgroups in $U(A_\Gamma)$ of rank equal to the dimension of certain *principal* cubes in K_Γ. In the absence of a specific configuration in Γ we find principal cubes of dimension equal to the dimension of K_Γ, thus determining the exact VCD of $U(A_\Gamma)$.

The free abelian subgroups we produce are generated by a special type of automorphisms called Γ-*Whitehead automorphisms*. These generalise the generating set used by J.H.C. Whitehead in his work on automorphisms of free groups [**17**]. We show that for any graph Γ, our free abelian subgroups have the largest possible rank among those generated by Γ-Whitehead automorphisms, which we call the *principal rank* of $U(A_\Gamma)$.

Because $U(A_\Gamma)$ is analogous to $Out(F_n)$ it is tempting to conjecture that the VCD of $U(A_\Gamma)$ is equal to the principal rank. It is also tempting to conjecture that the principal rank is always equal to the dimension of K_Γ … but our results show that if the graph contains a specific configuration then the dimension of K_Γ is strictly larger than the principal rank. The first conjecture is still plausible, however, because at least in some cases when the dimension of K_Γ is too large we can show that K_Γ equivariantly deformation retracts onto a strictly lower-dimensional cube complex.

For $GL(n, \mathbb{Z})$, of course, the VCD is not equal to the rank of a free abelian subgroup, but rather is equal to the Hirsch rank of a certain (non-abelian) polycyclic subgroup. In light of the above conjecture, it is natural to ask whether $U(A_\Gamma)$ can contain a torsion-free, non-abelian solvable subgroup. For many graphs the answer is no. This was proved in [**6**] for graphs with no triangles, and more generally for graphs where the link of every vertex is either discrete or connected. If links are disconnected but not discrete, we do not know the answer.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

We remark that several authors have established upper and lower bounds on the VCD of the full group $Out(A_\Gamma)$. In particular bounds for graphs with no triangles were given in [**7**], the exact VCD for Γ a tree was established in [**3**] and other special cases were determined exactly in [**10**].

The paper is organised as follows. In Section 2 we review basic facts and notation about right-angled Artin groups and their automorphisms, and define the subgroup $U(A_\Gamma)$. In Section 3 we review the definitions and results from [**5**] that we will need in this paper. In Section 4 we construct free abelian subgroups of $U(A_\Gamma)$ using Γ-Whitehead automorphisms, and show that these subgroups have maximal possible rank among all such subgroups. Section 5 studies the dimension of K_Γ and gives a condition for this dimension to equal the principal rank. Section 6 works out some concrete examples. Finally, in Section 7 we show in certain cases how to find an invariant deformation retract of K_Γ of strictly lower dimension.

2. *Right-angled Artin groups and their automorphisms*

In this section we recall the basic definitions and notation for right-angled Artin groups and their automorphisms. For further details and proofs, we refer to [**5**] and the references therein.

Definition 2·1. Let Γ be a finite simplicial graph, i.e. a finite graph with no loops or multiple edges, with vertex set $V = \{v_1, \ldots, v_n\}$. The *right-angled Artin group* A_Γ is the group with one generator for every vertex of Γ and one commutator relation for each edge, i.e. A_Γ has the presentation

$$A_\Gamma = \langle v_1, \ldots, v_n | [v_i, v_j] = 1 \text{ whenever } v_i \text{ and } v_j \text{ are connected by an edge in } \Gamma \rangle.$$

It is shown in [**12**] that two words in the generators represent the same element of A_Γ if and only if they can be made identical by a process of switching adjacent commuting letters and cancelling where possible.

If Γ is a simplicial graph with vertex set V, recall that the *induced subgraph on* $U \subseteq V$ is the subgraph of Γ with vertex set U that contains all edges in Γ connecting any vertices in U.

Definition 2·2. Let v be a vertex of a simplicial graph Γ. The *link of* v, denoted $lk(v)$, is the induced subgraph on the set of vertices adjacent to v. The *star of* v, denoted $st(v)$, is the induced subgraph on the set of vertices in $lk(v)$ together with v itself.

We will need the fact, shown in [**13**], that the centraliser of a generator v is equal to the subgroup generated by the vertices in $st(v)$.

In the literature on right-angled Artin groups it is common to define a relation denoted $\leq$ on vertices of Γ by $v \leq w$ if $lk(v) \subseteq st(w)$. The notation is justified by defining an equivalence relation $v \sim w$ if $v \leq w$ and $w \leq v$; it is then easy to verify that this relation defines a partial order on equivalence classes $[v]$. A vertex is called *maximal* if its equivalence class is maximal in this partial ordering.

In fact for $v \neq w$ there are two mutually exclusive ways in which we can have $lk(v) \subseteq st(w)$: either $lk(v) \subseteq lk(w)$ or $st(v) \subseteq st(w)$. The distinction is important in this paper, so

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

when we need to make it we will use $v \leq_\star w$ to mean $st(v) \subseteq st(w)$ and $v \leq_\circ w$ to mean $lk(v) \subseteq lk(w)$ (Similarly, $v \geq_\star w$ means $st(v) \supseteq st(w)$ and $v \geq_\circ w$ means $lk(v) \supseteq lk(w)$.)

We also write $v \sim_\star w$ if $st(v) = st(w)$ and $v \sim_\circ w$ if $lk(v) = lk(w)$, and define $[v]_\star = \{w | w \sim_\star v\}$, $[v]_\circ = \{w | w \sim_\circ v\}$. Since either all elements of an equivalence class $[v]$ commute or none commute, at least one of $[v]_\star$ and $[v]_\circ$ is a singleton. If $[v]_\circ$ is not a singleton then $[v]$ is called a *non-abelian equivalence class*; otherwise $[v]$ is called an *abelian equivalence class* (in particular a singleton class is considered to be abelian).

Definition 2·3. A vertex v of Γ is *principal* if there is no w with $v <_\circ w$, i.e. with $lk(v)$ strictly contained in $lk(w)$.

All maximal vertices are principal, but there can be principal vertices which are not maximal. A simple example is a triangle with leaves at two of its vertices. The third vertex is principal but not maximal. Elements of non-singleton abelian equivalence classes are always principal:

LEMMA 2·4. *If $u \neq v$ but $u \sim_\star v$ then both u and v are principal vertices.*

Proof. If u is not principal there exists m with $u <_\circ m$, i.e. $lk(u) \subsetneq lk(m)$. Now $v \in lk(u) \subset lk(m)$, so $m \in lk(v) \subset st(v) = st(u)$. Since $m \neq u$ we must have $m \in lk(u)$, which is a contradiction.

2·1. *Automorphisms of RAAGs*

An invertible map $A_\Gamma \to A_\Gamma$ is an automorphism if and only if the images of commuting generators commute. In particular:

(i) the map sending a generator v to its inverse and fixing all other generators is an automorphism, called an *inversion*.

(ii) any automorphism of the defining graph Γ induces an automorphism of A_Γ, called a *graph automorphism*.

Inversions and graph automorphisms generate a finite subgroup of $Aut(A_\Gamma)$. We next describe two types of basic infinite-order automorphisms. Choose a vertex m and consider the components of $\Gamma - st(m)$.

(i) If there is a vertex u with $lk(u) \subseteq lk(m)$, then everything that commutes with u also commutes with m so the map ρ_{um} sending $u \mapsto um$ and fixing all other generators determines an automorphism, called a *right fold*. Since u and m do not commute, the map λ_{um} sending u to mu gives a distinct automorphism, called a *left fold*.

(ii) If C is a component of $\Gamma - st(m)$, then the map sending v to $m^{-1}vm$ for every $v \in C$ and fixing all other generators determines an infinite-order automorphism, called a *partial conjugation*. If $\Gamma - st(m)$ has only one component, this is an inner automorphism, since conjugating vertices of $st(m)$ by m has no effect.

By work of Laurence [13] and Servatius [14], the entire automorphism group $Aut(A_\Gamma)$ is generated by the above types of automorphisms together with *twists*, where:

(i) if $st(u) \subseteq st(v)$, the map τ_{uv} sending $u \mapsto uv = vu$ and fixing all other generators determines an automorphism called a *twist*.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

2·2. *The untwisted subgroup*

The natural map $Aut(A_\Gamma) \to GL(n, \mathbb{Z})$ induced by abelianisation $A_\Gamma \to \mathbb{Z}^n$ factors through the outer automorphism group $Out(A_\Gamma)$:

$$\begin{array}{ccc} Aut(A_\Gamma) & \longrightarrow & GL(n,\mathbb{Z}) \\ & \searrow \quad \nearrow & \\ & Out(A_\Gamma). & \end{array}$$

The subgroup $T(A_\Gamma) \subseteq Out(A_\Gamma)$ generated by twists injects into a parabolic subgroup of $GL(n, \mathbb{Z})$, and is well understood (see, e.g., [6]). In this paper we concentrate on the subgroup $U(A_\Gamma) \leq Out(A_\Gamma)$ generated by all other generators, i.e.

Definition 2·5. The *untwisted subgroup* $U(A_\Gamma)$ is the subgroup of $Out(A_\Gamma)$ generated by (the images of):

(i) inversions;
(ii) graph automorphisms;
(iii) (right and left) folds; and
(iv) partial conjugations.

The intersection $U(A_\Gamma) \cap T(A_\Gamma)$ is contained in the finite subgroup generated by graph automorphisms and inversions.

3. Γ-*Whitehead automorphisms, partitions and outer space for* $U(A_\Gamma)$.

The paper [5] studied $U(A_\Gamma)$ by constructing a contractible space $\mathcal{O}_\Gamma$ with a proper action of $U(A_\Gamma)$. In this section we review the definitions and results from [5] that we will need in this paper. Some of the terminology has been altered slightly, and we will point this out when it occurs. We refer to [5] for more details and all proofs.

3·1. Γ-*Whitehead automorphisms*

Whitehead studied $Aut(F_n)$ using a set of generators called Whitehead automorphisms. These were adapted in [5] to a give a set of elements of $Aut(A_\Gamma)$ called Γ-*Whitehead automorphisms*, whose images in $Out(A_\Gamma)$ along with graph automorphisms and inversions generate $U(A_\Gamma)$. These are infinite-order automorphisms which include folds and partial conjugations but also certain combinations of these.

For a free group with basis V, let $V^{\pm} = V \sqcup V^{-1}$ be the set of generators and their inverses. Suppose $P \subset V^{\pm}$ contains some element m but not m^{-1}. The *Whitehead automorphism* $\phi(P, m)$ is defined on the basis V by

$$\phi(P,m)(v) = \begin{cases} vm^{-1} & \text{if } v \in P,\ v^{-1} \in P^*,\ v \neq m^{\pm 1}, \\ mv & \text{if } v^{-1} \in P,\ v \in P^*,\ v \neq m^{\pm 1}, \\ mvm^{-1} & \text{if } v,\ v^{-1} \in P, \\ v & \text{otherwise (including } v = m^{\pm 1}). \end{cases}$$

The element m is called the *multiplier* of $\phi(P, m)$.

If V is the set of vertices of a simplicial graph Γ, then this formula defines an automorphism of A_Γ only for certain pairs (P, m). Specifically, for $m \in V^{\pm}$ consider the components

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

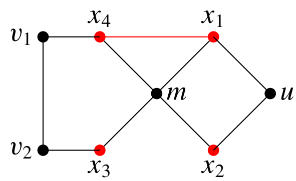

Fig. 1. Example 3·1.

C of $\Gamma - lk(m)$, where by $lk(m)$ we mean the link of the corresponding vertex $m^{\pm 1}$. A subset $U \subset V^{\pm}$ is *m-inseparable* if:

(i) C has only one vertex u, and $U = \{u\}$ or $U = \{u^{-1}\}$ (note this includes the case $u = m^{\pm 1}$); or
(ii) C contains more than one vertex and $U = C^{\pm}$, i.e. U is the union of all vertices in C and their inverses.

We denote by $\mathbf{I}(m)$ the collection of all m-inseparable subsets of $V^{\pm}$. Note that $\mathbf{I}(m) = \mathbf{I}(m^{-1})$, and if m and n have the same link then $\mathbf{I}(m) = \mathbf{I}(n)$.

Example 3·1. In the graph Γ in Figure 1 the link of the vertex m is the red subgraph, and the m-inseparable subsets are

$$\mathbf{I}(m) = \left\{\{m\}, \{m^{-1}\}, \{u\}, \{u^{-1}\}, \{v_1, v_1^{-1}, v_2, v_2^{-1}\}\right\}$$

Recall that a partition of a set into two subsets is *thick* if each side has at least two elements.

Definition 3·2. Let $m \in V^{\pm}$.

(i) A subset $P \subset V^{\pm}$ is called a *ΓW-subset based at m* if it is a union of elements of $\mathbf{I}(m)$ and contains m but not m^{-1}.
(ii) If P is a ΓW-subset based at m then $\phi(P, m)$ is a well-defined automorphism of A_Γ, called a *Γ-Whitehead automorphism.*
(iii) Let $P^* = V^{\pm} \setminus lk(m)^{\pm} \setminus P$. The three-part partition $\mathcal{P} = \{P|P^*|lk(m)^{\pm}\}$ of $V^{\pm}$ is called a *ΓW-partition based at m* if P (and therefore P^*) are ΓW-subsets and $\{P|P^*\}$ is a thick partition of $V^{\pm} \setminus lk(m)^{\pm}$. The subsets P and P^* are called the *sides* of $\mathcal{P}$.

Remark 3·3. For $A_\Gamma = F_n$ the above is the usual definition of a Whitehead automorphism. In [5], however, a Γ-Whitehead automorphism was defined as sending $m \mapsto m^{-1}$ instead of $m \mapsto m$. This makes the automorphism into an involution, and is useful for describing geometric aspects of $U(A_\Gamma)$. Since we are looking for free abelian subgroups we do not want involutions, so will use the more classical definition stated here.

In terms of the inseparable subsets $U \in \mathbf{I}(m)$, $\phi(P, m)$ is the composition of;

(i) right folds $v \mapsto vm^{-1}$ for $U = \{v\} \subset P$, $v \neq m^{\pm 1}$;
(ii) left folds $v \mapsto mv$ for $U = \{v^{-1}\} \subset P$, $v \neq m^{\pm 1}$; and
(iii) partial conjugations $v \mapsto mvm^{-1}$ for $U = C^{\pm} \subset P$ if C has at least two elements.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

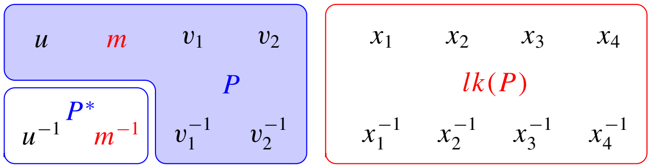

Fig. 2. Example of a ΓW-partition based at m for the graph in Figure 1.

Example 3·4. Continuing Example 3·1, we can take $P = \{m\} \cup \{u\} \cup \{v_1, v_1^{-1}, v_2, v_2^{-1}\}$ and $P^* = \{m^{-1}\} \cup \{u^{-1}\}$ to get a ΓW-partition

$$\mathcal{P} = \left\{\{m, u, v_1, v_1^{-1}, v_2, v_2^{-1}\} | \{m^{-1}, u^{-1}\} | \{x_1, x_1^{-1}, x_2, x_2^{-1}, x_3, x_3^{-1}, x_4, x_4^{-1}\}\right\}$$

based at m (see Figure 2). The Γ-Whitehead automorphism $\phi(P, m)$ sends $u \mapsto um^{-1}$, sends each $v_i \mapsto mv_im^{-1}$ and fixes m and the x_i. The Γ-Whitehead automorphism $\phi(P^*, m^{-1})$ sends $u \mapsto m^{-1}u$ and fixes all other generators.

LEMMA 3·5. *Let $\phi(P, m)$ be a Γ-Whitehead automorphism. Then*;

(i) $\phi(P, m)^{-1} = \phi(P \setminus \{m\} \cup \{m^{-1}\}, m^{-1})$;
(ii) *$\phi(P^*, m^{-1})$ is equal to $\phi(P, m)$ composed with conjugation by m, so the two are equal as outer automorphisms.*

Proof. Clear from the definitions.

For a ΓW-partition $\mathcal{P} = \{P | P^* | lk(m)^{\pm}\}$ based at m we define the outer automorphism $\varphi(\mathcal{P}, m)$ to be

$$\varphi(\mathcal{P}, m) = \begin{cases} \text{the image of } \phi(P, m) & \text{if } m \in P, \\ \text{the image of } \phi(P^*, m) & \text{if } m \in P^*. \end{cases}$$

We will call $\varphi(\mathcal{P}, m)$ an *outer Γ-Whitehead automorphism*. By Lemma 3·5, $\varphi(\mathcal{P}, m) = \varphi(\mathcal{P}, m^{-1})$, so we can think of the m in $\varphi(\mathcal{P}, m)$ as a vertex of Γ instead of an element of $V^{\pm}$.

Notation 3·6. We extend the relations $\leq, \sim, \leq_\circ, \sim_\circ, \leq_\star, \sim_\star$ etc. to elements of $V^{\pm}$ by saying a relation holds if and only if it holds for the corresponding vertices.

If P is a ΓW-subset based at m, let $max(P)$ be the elements $n \in P$ with $n \sim_\circ m$ and $n^{-1} \notin P$. Then P is also based at any $n \in max(P)$. Since all elements of $max(P)$ have the same link, we will write $\mathcal{P} = \{P | P^* | lk(P)\}$. There is a Γ-Whitehead automorphism $\phi(P, m)$ for each $m \in max(P)$.

Definition 3·7 ([5, definition 3·3]). Let $\mathcal{P}$ and $\mathcal{Q}$ be ΓW-partitions, with $\mathcal{P}$ based at m and $\mathcal{Q}$ based at n. Then $\mathcal{P}$ and $\mathcal{Q}$ are *compatible* if either:

(i) $P^\times \cap Q^\times = \emptyset$ for at least one choice of sides $P^\times \in \{P, P^*\}$ and $Q^\times \in \{Q, Q^*\}$; or
(ii) $[m, n] = 1$ but $st(m) \neq st(m)$.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

Remark 3·8. This is the definition of compatibility given in [5]. However the definition that is actually used in the proofs in that paper is weaker: condition (2) needs to be replaced by

(i) $[m, n] = 1$ but $m \neq n$.

We will call this *weak compatibility*. The proofs in this paper use the stronger notion of compatibility, but we show in Lemma 4·19 that this does not change the results of this paper.

If the bases m of $\mathcal{P}$ and n of $\mathcal{Q}$ do not commute, the following lemma constrains the relationships between sides of $\mathcal{P}$ and $\mathcal{Q}$.

LEMMA 3·9 ([5, Lemma 3·4]). *Suppose that $\mathcal{P} = \{P|P^*|lk(P)\}$ based at m and $\mathcal{Q} = \{Q|Q^*|lk(Q)\}$ based at n are compatible, m and n do not commute and $P \cap Q = \emptyset$. Then $P \cap lk(Q) = \emptyset$. In particular, $P \subseteq Q^*$ and $Q \subseteq P^*$.*

3·2. *Outer space $\mathcal{O}_\Gamma$ and its spine K_Γ*

In [5] an "outer space" $\mathcal{O}_\Gamma$ was defined on which $U(A_\Gamma)$ acts properly, and it was proved that $\mathcal{O}_\Gamma$ is contractible. The proof proceeds by retracting $\mathcal{O}_\Gamma$ equivariantly onto a *spine* K_Γ, which is the geometric realization of a partially ordered set (poset) of *marked* Γ*-complexes* (g, X) with $\pi_1(X) \cong A_\Gamma$.

The simplest example of a Γ-complex is the *Salvetti complex* S_Γ. This is the non-positively curved (i.e. locally CAT(0)) cube complex with a single 0-cell, one edge for each vertex of Γ, and one $(k+1)$-cube for each k-clique in Γ. A general Γ-complex X is a certain type of non-positively curved cube complex which can be collapsed along hyperplanes to produce the Salvetti complex. A *marking* is a homotopy equivalence $g: S_\Gamma \to X$ from a fixed standard Salvetti S_Γ whose fundamental group we identify with A_Γ, with the property that if $c: X \to S_\Gamma$ is a sequence of hyperplane collapses then the composition $c \circ g: S_\Gamma \to X \to S_\Gamma$ induces an element of $U(A_\Gamma)$ on the level of fundamental groups. The group $U(A_\Gamma)$ acts on vertices (g, X) of K_Γ by changing the marking.

Each Γ-complex X is constructed using a collection of pairwise-compatible ΓW-partitions (see [5] for the construction; we will not need to know the details). If we start with $X = S$ homeomorphic to S_Γ and fix a marking $g: S_\Gamma \to S$, the empty collection corresponds to the marked Salvetti (g, S), and the partially ordered set of all compatible collections of ΓW-partitions (ordered by inclusion) corresponds precisely to the star of (g, S) in K_Γ. In other words, each (ordered) compatible collection $(\mathcal{P}_1, \ldots, \mathcal{P}_k)$ corresponds to a k-simplex

$$\emptyset \subset \{\mathcal{P}_1\} \subset \{\mathcal{P}_1, \mathcal{P}_2\} \subset \cdots \subset \{\mathcal{P}_1, \ldots, \mathcal{P}_k\}$$

of the star; we abuse notation by writing

$$\emptyset \subset \mathcal{P}_1 \subset \mathcal{P}_1\mathcal{P}_2 \subset \cdots \subset \mathcal{P}_1\mathcal{P}_2\cdots\mathcal{P}_k.$$

The entire complex K_Γ is the orbit of a single such star, so the dimension of K_Γ is equal to the maximal size of a compatible collection of ΓW-partitions. (Lemma 4·19 shows that this size does not depend on whether one uses compatibility or weak compatibility.)

Since $Out(A_\Gamma)$ is known to have torsion-free subgroups of finite index, the fact that $U(A_\Gamma)$ acts properly on K_Γ gives:

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

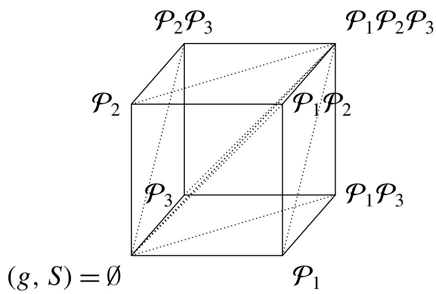

Fig. 3. The cube $c(\emptyset, \mathcal{P}_1\mathcal{P}_2\mathcal{P}_3)$ in the star of (g, S) in K_Γ.

THEOREM 3·10. *The* VCD *of $U(A_\Gamma)$ is less than or equal to the maximal size of a compatible collection of ΓW-partitions.*

3·3. *Cube complex structure of K_Γ*

Note that any ordering of $\{\mathcal{P}_1, \ldots, \mathcal{P}_k\}$ gives a k-simplex in the star of (g, S_Γ), and the union of all of these simplices forms a k-dimensional cube (see Figure 3). Thus K_Γ in fact has the structure of a cube complex, with one k-dimensional cube for each compatible collection $\Pi = \{\mathcal{P}_1, \ldots, \mathcal{P}_k\}$, which we will denote $c(\emptyset, \Pi)$. The faces of $c(\emptyset, \Pi)$ correspond to pairs $\Pi_1 \subset \Pi_2$ of subsets of Π; in particular the maximal faces of are of the form $c(\emptyset, \Pi \setminus \{\mathcal{P}\})$ and $c(\{\mathcal{P}\}, \Pi)$ for some $\mathcal{P} \in \Pi$.

4. *Free abelian subgroups of $U(A_\Gamma)$*

In this section we relate the dimension of K_Γ to abelian subgroups of $Out(A_\Gamma)$ by constructing abelian subgroups freely generated by outer Γ-Whitehead automorphisms associated to compatible collections of ΓW-partitions. We start by determining exactly when two of these commute.

4·1. *Commuting Γ-Whitehead automorphisms*

Definition 4·1. Let v be a vertex of Γ. A ΓW-partition $\mathcal{P}$ *splits* v if v and v^{-1} are in different sides of $\mathcal{P}$.

THEOREM 4·2. *Let $\phi(P, m)$ and $\phi(Q, n)$ be Γ-Whitehead automorphisms. If $[m, n] = 1$ then $\phi(P, m)$ commutes with $\phi(Q, n)$. If $[m, n] \neq 1$ let $\mathcal{P} = \{P|P^*|lk(P)\}$ and $\mathcal{Q} = \{Q|Q^*|lk(Q)\}$ be the associated ΓW-partitions. Then the outer automorphisms $\varphi(\mathcal{P}, m)$ and $\varphi(\mathcal{Q}, n)$ commute if and only if $\mathcal{P}$ and $\mathcal{Q}$ are compatible, $\mathcal{Q}$ does not split m and $\mathcal{P}$ does not split n.*

Proof. If m and n commute, the automorphisms clearly commute, so we only need to consider the case that m and n do not commute.

Suppose first that $\mathcal{P}$ and $\mathcal{Q}$ are compatible. Replacing (P, m) by (P^*, m^{-1}) and/or (Q, n) by (Q^*, n^{-1}) if necessary (which does not change $\varphi(\mathcal{P}, m)$ or $\varphi(\mathcal{Q}, n)$), then by the definition of compatibility we may assume that $P \cap Q = \emptyset$, $m \in P$ and $n \in Q$.

If both m^{-1} and n^{-1} are in $P^* \cap Q^*$, then $\phi = \phi(P, m)$ affects only elements of P and their inverses, and $\psi = \phi(Q, n)$ affects only elements of Q and their inverses. In particular

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

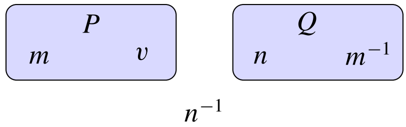

Fig. 4. A case in the proof of Theorem 4·2.

ϕ fixes n and ψ fixes m. If $x \in P$ and $x^{-1} \in Q$ then ϕ and ψ act on opposite sides of x. It follows that $\phi\psi(x) = \psi\phi(x)$ for all generators x.

If $n^{-1} \in P$ and $m^{-1} \in Q$, then $\phi\psi(m) = mnm$ while $\psi\phi(m) = nm$. Since these are not conjugate, $\phi\psi$ and $\psi\phi$ do not differ by an inner automorphism, i.e. they do not commute as outer automorphisms.

If $n^{-1} \in P^*$ but $m^{-1} \in Q$, then $\phi\psi(n) = n = \psi\phi(n)$ and $\phi\psi(m) = nm = \psi\phi(m)$ so we need a different argument to show that ϕ and ψ do not commute. Since P must have at least two elements, there is $v \in P$ with $v \neq m$ (see Figure 4). Since $P \subset Q^*$ by Lemma 3·9, v does not commute with m or n, so v, m and n generate a free group of rank three. Since $\phi\psi$ and $\psi\phi$ agree on two generators of this free group, they differ by an inner automorphism if and only if they are equal.

The effects of $\phi\psi$ and $\psi\phi$ on v are determined by the position of v^{-1}:

(i) if $v^{-1} \in P^* \cap Q^*$ then $\phi\psi(v) = vm^{-1}$ and $\psi\phi(v) = vm^{-1}n^{-1}$;
(ii) if $v^{-1} \in Q$ then $\phi\psi(v) = nvm^{-1}$ and $\psi\phi(v) = nvm^{-1}n^{-1}$;
(iii) if $v^{-1} \in P$ then $\phi\psi(v) = mvm^{-1}$ and $\psi\phi(v) = nmvm^{-1}n^{-1}$

Thus in all cases, $\phi\psi$ does not differ from $\psi\phi$ by an inner automorphism.

This argument applies also to the symmetric case $n^{-1} \in P$ but $m^{-1} \in Q^*$.

It remains to consider the possibility that $\mathcal{P}$ and $\mathcal{Q}$ are not compatible. In this case all four quadrants $P \cap Q$, $P \cap Q^*$, $P^* \cap Q$ and $P^* \cap Q^*$ are non-empty. Using Lemma 3·5 we may replace (P, m) by (P^*, m^{-1}) (which does not change $\varphi(\mathcal{P}, m)$) or by $(P \setminus \{m\} \cup \{m^{-1}\}, m^{-1})$ (which replaces $\varphi(\mathcal{P}, m)$ by its inverse), and similarly replace (Q, n) if necessary, to obtain one of the following configurations:

(i) if each quadrant contains one of $\{m, m^{-1}, n, n^{-1}\}$, then we may assume $m \in P \cap Q^*$, $n \in P \cap Q$, $m^{-1} \in P^* \cap Q$ and $n^{-1} \in P^* \cap Q^*$. Then $\phi\psi(n) = nm^{-1}$ and $\psi\phi(n) = nm^{-1}n^{-1}$ are not conjugate in A_Γ, so $\phi\psi$ and $\psi\phi$ do not differ by an inner automorphism.

(ii) if exactly two quadrants contain elements of $\{m, m^{-1}, n, n^{-1}\}$, then we may assume $m, n^{-1} \in P \cap Q^*$ and $n, m^{-1} \in P^* \cap Q$ so $\phi\psi(m) = nm$ and $\psi\phi(m) = mnm$, which are not conjugate in A_Γ.

(iii) if exactly 3 quadrants contain elements of $\{m, m^{-1}, n, n^{-1}\}$ then we may assume $m \in P \cap Q^*$, $n \in P^* \cap Q$, $m^{-1} \in P^* \cap Q^*$, and either $n^{-1} \in P^* \cap Q^*$ or $n^{-1} \in P \cap Q^*$. For either position of n^{-1} we have $\phi\psi(m) = \psi\phi(m)$ and $\phi\psi(n) = \psi\phi(n)$. Now $P \cap Q$ does not contain any element of $\{m, m^{-1}, n, n^{-1}\}$ but it cannot be empty, so let $v \in P \cap Q$. Note that v cannot commute with m or n, so m, n and v are the basis of a free subgroup of A_Γ. Therefore if $\phi\psi$ is conjugate to $\psi\phi$ we must have $\phi\psi(v) = \psi\phi(v)$. A calculation now shows that this is not the case for any position of v^{-1}.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

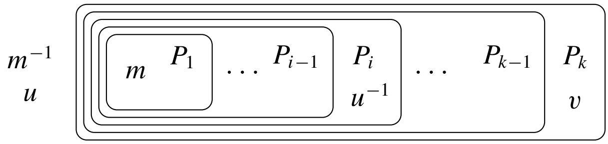

Fig. 5. Proof of Proposition 4·6.

COROLLARY 4·3. *If* Γ*-Whitehead automorphisms* $\phi(P, m)$ *and* $\phi(Q, n)$ *commute as outer automorphisms, then* $\phi(P, m)$ *acts on n either trivially or as conjugation by m.*

Proof. This is immediate from Theorem 4·2 and the definition of $\phi(P, m)$.

Let $m \in V^{\pm}$ and let $\mathcal{P} = (P, P^*, lk(P))$ be a ΓW-partition based at m. We define the *m-length of $\mathcal{P}$* to be the number of m-inseparable subsets in the side of $\mathcal{P}$ containing m.

LEMMA 4·4. *Let* $m \in V^{\pm}$ *and let* $\mathcal{P}$ *and* $\mathcal{Q}$ *be distinct* Γ*W-partitions based at m, with* m*-length*$(\mathcal{P}) = m$*-length*$(\mathcal{Q})$*. Then* $\mathcal{P}$ *and* $\mathcal{Q}$ *are incompatible.*

Proof. The sides of $\mathcal{P}$ and $\mathcal{Q}$ containing m are unions of elements of $\mathbf{I}(m)$. If they have the same m-length but are different, then all sides of $\mathcal{P}$ and $\mathcal{Q}$ must intersect non-trivially.

LEMMA 4·5. *Let* $m \in V^{\pm}$ *and let* $\mathcal{P}_1, \dots, \mathcal{P}_k$ *be pairwise-compatible* Γ*W-partitions based at m. Let* P_i *be the side of* $\mathcal{P}_i$ *that contains m. Then after reordering we may assume* $P_1 \subset P_2 \subset \cdots \subset P_k$.

Proof. For each $i \neq j$, $P_i \cap P_j$ contains m, so is not empty, and $P_i^* \cap P_j^*$ contains m^{-1}, so is not empty. Therefore, by compatibility, either $P_i \cap P_j^* = \emptyset$, which implies $P_i \subset P_j$, or $P_j \cap P_i^* = \emptyset$, which implies $P_j \subset P_i$. Therefore we can renumber the P_i in order of size to obtain $P_1 \subset \cdots \subset P_k$.

PROPOSITION 4·6. *Let* $m \in V^{\pm}$ *and suppose* $\mathcal{P}_1, \dots, \mathcal{P}_k$ *are pairwise compatible* Γ*W-partitions based at m. Then the subgroup of* $U(A_\Gamma)$ *generated by the* $\varphi(\mathcal{P}_i, m)$ *is free abelian of rank k.*

Proof. Let $P_1 \subset \cdots \subset P_k$ be the sides of the $\mathcal{P}_i$ that contain m as in Lemma 4·5 (see Figure 5), and let $\phi_i = \phi(P_i, m)$. Suppose $g = \phi_1^{n_1} \dots \phi_k^{n_k}$ is inner, and let $u \in P_k^*$, $u \neq m^{-1}$. Then

$$g(u) = \begin{cases} u & \text{if } u^{-1} \in P_k^*, \\ m^a u & \text{if } u^{-1} \in P_k, \end{cases}$$

where $a = \sum_{\ell=i}^{k} n_\ell$ if $u^{-1} \in P_i \cap P_{i-1}^*$. Since $g(u)$ is conjugate to u, we must have $a = 0$, i.e. $g(u) = u$ in all cases, so g is not just inner, but is actually the identity. Now let $v \in P_k \cap P_{k-1}^*$. Then

$$g(v) = \begin{cases} vm^{-n_k} & \text{if } v^{-1} \in P_k^*, \\ m^b m^{n_k} vm^{-n_k} & \text{if } v^{-1} \in P_k, \end{cases}$$

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

where $b = 0$ if $v^{-1} \in P_k \cap P^*_{k-1}$ and $b = \sum_{\ell=j}^{k-1} n_\ell$ if $v^{-1} \in P_j \cap P^*_{j-1}$ for some $j < k$. Since $g = id$, this implies $n_k = 0$ in all cases. Repeating this argument with $P_1 \subset \cdots \subset P_r$ for each $r < k$ gives $n_r = 0$ for all r.

PROPOSITION 4·7. *Let $\{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$ be a maximal compatible collection of ΓW-partitions based at m. Suppose $\mathcal{Q}$ is another ΓW-partition based at m. Then $\varphi(\mathcal{Q}, m)$ is in the subgroup G of $U(A_\Gamma)$ generated by the $\varphi(\mathcal{P}_i, m)$.*

Proof. Let P_i be the side of $\mathcal{P}_i$ containing m. By maximality of the collection together with Lemma 4·4 we know that $\mathbf{I}(m)$ has exactly $n+1$ elements $U_1, \ldots, U_{n+1}$ other than $\{m\}$ and $\{m^{-1}\}$ and (after setting $P_0 = \{m\}$ and possibly reordering) we have $P_i = P_{i-1} \cup U_i$. Define $P_{n+1} = P_n \cup U_{n+1}$ and set $V_i = U_i \cup \{m\}$. Then for all i with $1 \le i \le n+1$ we have $\phi(V_i, m) = \phi(P_i, m) \circ \phi(P_{i-1}, m)^{-1}$, so the corresponding outer automorphism is in G.

Each m-inseparable set in the side Q of $\mathcal{Q}$ containing m is one of the U_i, so we have $Q = \{m\} \cup U_{i_1} \cup \cdots \cup U_{i_k}$. Then

$$\phi(Q, m) = \phi(V_{i_1}, m) \circ \cdots \circ \phi(V_{i_k}, m),$$

so $\varphi(\mathcal{Q}, m)$ is in G.

We next show how Propositions 4·6 and 4·7 generalise to the situation where all partitions are based in the same abelian equivalence class.

LEMMA 4·8. *Let $\mathcal{P} = \{P|P^*|lk(P)\}$ be based at $v \in \Gamma$ and let $w \in \Gamma$ be a distinct vertex with $st(w) = st(v)$. Let P be the side of $\mathcal{P}$ containing v, set $P_{v,w} = P \setminus \{v\} \cup \{w\}$ and $\mathcal{P}_{v,w} = \left\{P_{v,w}|P^*_{v,w}|lk(w)^{\pm}\right\}$. Then*:

(i) *$\mathcal{P}$ and $\mathcal{P}_{v,w}$ are compatible*;

(ii) *if $\mathcal{R}$ is compatible with $\mathcal{P}$ then $\mathcal{R}$ is also compatible with $\mathcal{P}_{v,w}$*;

(iii) *if $\varphi(\mathcal{R}, s)$ commutes with $\varphi(\mathcal{P}, v)$ then $\varphi(\mathcal{R}, s)$ commutes with $\varphi(\mathcal{P}_{v,w}, w)$.*

Proof. For the first statement, notice that $P \cap P^* = \emptyset$ implies $P \cap (P_{v,w})^* = \emptyset$ since $w \in lk(v)$.

Now suppose $\mathcal{R}$ is based at s and is compatible with $\mathcal{P}$. If $[v, s] = 1$ and $st(v) \neq st(s)$, then $st(w) \neq st(s)$ so $\mathcal{R}$ is compatible with $\mathcal{P}_{v,w}$.

If $st(v) = st(s)$ or if $[v, s] \neq 1$ then by possibly renaming sides may assume $P \cap R = \emptyset$. The only element of $P_{v,m}$ which is not in P is w. If $st(s) = st(v) = st(w)$ then $w \in lk(R)$, and if $[s, v] \neq 1$ then $R \subset P^*$, which does not contain w. In either case $w \notin R$, so $P_{v,w} \cap R = \emptyset$ and $\mathcal{P}_{v,w}$ is compatible with $\mathcal{R}$.

For the third statement, by Theorem 4·2 it remains to check that if $[w, s] \neq 1$ then $\mathcal{R}$ does not split w and $\mathcal{P}_{v,w}$ does not split s. The first statement is clear since $w, w^{-1} \in lk(P)^{\pm}$, which does not intersect R. The second follows since $\mathcal{P}$ does not split s, and the only difference between $\mathcal{P}$ and $\mathcal{P}_{v,w}$ is the base w.

Remark 4·9. If $st(v) \subset st(w)$ and $P_{v,w} = P \setminus lk(w)^{\pm} \cup \{w\}$, then statements (*i*) and (*iii*) of Lemma 4·8 hold and statement (*ii*) holds unless $st(s) = st(w)$.

We say that $\mathcal{P}_{v,w}$ in Lemma 4·8 is obtained from $\mathcal{P}$ by *exchanging v for w*.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

COROLLARY 4·10. *Let Π be a maximal compatible collection of ΓW-partitions, and let $[v]$ be an abelian equivalence class of Γ. If $\mathcal{P} \in \Pi$ is based at $v \in [v]$, then Π contains every ΓW-partition that can be obtained from $\mathcal{P}$ by exchanging v for a different element $w \in [v]$.*

Definition 4·11. Let $\mathcal{P}$ be a ΓW-partition based at m. Define $\mathring{\mathcal{P}}$ to be the partition of $V^{\pm} \setminus st(m)^{\pm}$ obtained by intersecting each side of $\mathcal{P}$ with $V^{\pm} \setminus st(m)^{\pm}$.

LEMMA 4·12. *For a vertex m of Γ, let $\mathcal{P}_1, \ldots, \mathcal{P}_k$ be pairwise-compatible ΓW-partitions based at $m_i \in [m]_*$. Then for some ordering of the $\mathcal{P}_i$ and some choice of sides P_i we have $\mathring{P}_1 \subseteq \mathring{P}_2 \subseteq \cdots \subseteq \mathring{P}_k$.*

Proof. Let P_i be the side of $\mathcal{P}_i$ that contains m_i, and set $\mathring{P}_i = P_i \setminus \{m_i\}$. Fix $m \in [m]_*$ and for each i define $P_{i,m} = P_i \setminus \{m_i\} \cup \{m\} = \mathring{P}_i \cup \{m\}$. Then the $P_{i,m}$ are all compatible by Lemma 4·8, and by Lemma 4·5 we can renumber the $P_{i,m}$ in order of size to obtain $P_{1,m} \subseteq \cdots \subseteq P_{k,m}$. Removing m from each P_i now gives $\mathring{P}_1 \subseteq \mathring{P}_2 \subseteq \cdots \subseteq \mathring{P}_k$.

PROPOSITION 4·13. *Let $[m]$ be an abelian equivalence class and suppose $\Pi = \{\mathcal{P}_1, \ldots, \mathcal{P}_k\}$ is a compatible collection of distinct ΓW-partitions based at elements $m_i \in [m]$. Then the subgroup of $U(A_\Gamma)$ generated by the $\varphi(\mathcal{P}_i, m_i)$ is free abelian of rank k.*

Proof. Since $[m]$ is abelian the base m_i of each $\mathcal{P}_i$ is uniquely determined by $\mathcal{P}_i$, so we may partition Π into subsets Π_n with the same base $n \in [m]$. The subgroup generated by the $\varphi(\mathcal{P}_i, m_i) \in \Pi_n$ is free abelian by Proposition 4·6, and the intersection of any two of these is trivial since they use different multipliers. Therefore the subgroup generated by all of the $\varphi(\mathcal{P}_i, m_i)$ is the direct product of the subgroups A_n generated by the $\varphi(\mathcal{P}_i, m_i) \in \Pi_n$, so is free abelian of rank k.

PROPOSITION 4·14. *Let $\{\mathcal{P}_1, \ldots, \mathcal{P}_k\}$ be a maximal compatible collection of ΓW-partitions based at elements m_i of an abelian equivalence class $[m]$. Suppose Q is another ΓW-partition based at some $n \in [m]$. Then $\varphi(Q, n)$ is in the subgroup generated by the $\varphi(\mathcal{P}_i, m_i)$.*

Proof. Since Π is maximal, $n = m_i$ for some i by Lemma 4·10. Also, the partitions $\mathcal{P}_i$ based at n form a maximal collection of such partitions. So by Proposition 4·7 $\varphi(Q, n)$ is in the subgroup generated by the $\varphi(\mathcal{P}_i, n)$.

4·2. *Large abelian subgroups of $U(A_\Gamma)$*

Definition 4·15. For any subset $U \subset V$ of vertices of Γ, let $M(U)$ denote the largest possible size of a compatible collection of ΓW-partitions, each based at some $u \in U$.

Example 4·16. $M(V) = dim(K_\Gamma)$, by Theorem 3·10.

Example 4·17. $M(m) = |\mathbf{I}(m)| - 3$, since any ΓW-partition based at m gives a thick partition of $\mathbf{I}(m)$, and the largest compatible set of such partitions is obtained by adding one element of $\mathbf{I}(m)$ at a time.

Notation 4·18. Let Π be a compatible collection of ΓW-partitions, and $U \subset V$ a subset of vertices of Γ. Then:

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

(i) $\Pi_U = \{\mathcal{P} \in \Pi : \mathcal{P}$ is based at some $u \in U\}$ and
(ii) $\Pi^{\pm}$ is the set of ΓW-subsets of $V^{\pm}$ which are sides of elements of Π.

In this section we find a free abelian subgroup of $U(A_\Gamma)$ of rank $M(L)$, where L is the set of principal vertices of Γ, i.e. the set of vertices of Γ with maximal links. This subgroup will be generated by Γ-Whitehead automorphisms, and we will also show that every abelian subgroup freely generated by Γ-Whitehead automorphisms has rank at most $M(L)$. The following lemma shows that this bound is unchanged if we use the weaker notion of compatibility (see Remark 3·8).

LEMMA 4·19. *Let $U \subset V$ be any subset of vertices of Γ, and let $\mu(U)$ denote the largest possible size of a weakly compatible collection of ΓW-partitions, each based at some $u \in U$. Then $\mu(U) = M(U)$.*

Proof. Let Π be any collection of weakly compatible partitions of size $\mu(U)$. For each abelian equivalence class $[v]$ choose $m \in [v]$ such that $|\Pi_m|$ is largest. Remove all $\mathcal{P} \in \Pi_{[v]} - \Pi_m$ from Π, then add partitions $\mathcal{P}_{m,n}$ for each $\mathcal{P} \in \Pi_m$ and $n \in [v]$ with $n \neq m$. By Lemma 4·8 the resulting collection Π' is a (strongly) compatible collection, and since $|\Pi_m|$ was largest we have $|\Pi'| \geq |\Pi|$. Therefore, $\mu(U) \leq M(U)$. However, any compatible partitions are weakly compatible so $\mu(U) \geq M(U)$ giving equality.

In Lemma 4·20 to Proposition 4·22 we fix a compatible collection Π of ΓW-partitions. Recall that a partition *splits* a vertex v if v and v^{-1} are in different sides of the partition.

LEMMA 4·20. *Suppose $\mathcal{P} \in \Pi$ is based at m and $\mathcal{R} \in \Pi$ is based at $s \not\sim m$. If m and s do not commute and $\mathcal{R}$ splits some vertex in $[m]$, then $m <_\circ s$. In particular, if m is principal then all of $[m]^{\pm}$ is in the same side of $\mathcal{R}$.*

Proof. We are assuming $m \not\sim s$, so if $m \not<_\circ s$ there is some $v \in lk(m)$ which is not in $lk(s)$. This v is adjacent to every element of $[m]$ so all of $[m]$ is in the same component of $\Gamma - lk(s)$.

LEMMA 4·21. *Let m be a principal vertex of Γ, $\mathcal{P}_1, \ldots, \mathcal{P}_k \in \Pi_{[m]_\star}$ and let*

$$\emptyset = \mathring{P}_0 \subset \mathring{P}_1 \subseteq \cdots \subseteq \mathring{P}_k \subset \mathring{P}_{k+1} = V^{\pm} \setminus st(m)^{\pm},$$

where $\mathring{P}_1 \subseteq \ldots \subseteq \mathring{P}_k$ is the nest found in Lemma 4·12. Suppose $\mathcal{Q} \in \Pi \setminus \Pi_{[m]_\star}$ is based at n. If m does not commute with n, then there is a side Q of $\mathcal{Q}$ with $Q \subseteq \mathring{P}_i \cap \mathring{P}^*_{i-1}$ for some i with $1 \leq i \leq k+1$.

Proof. Since $\mathcal{Q}$ is compatible with each $\mathcal{P}_i$ and m does not commute with n, Lemma 3·9 implies that for each i there is some choice of side Q of $\mathcal{Q}$ so that either $Q \subset P_i$ or $Q \subset P_i^*$. Since the base m_i of $\mathcal{P}_i$ is principal, Q does not split m, by Lemma 4·20. Since $Q \subset P_i$ or $Q \subset P_i^*$, this means Q cannot contain either m_i or m_i^{-1}, so in fact either $Q \subset \mathring{P}_i$ or $Q \subset \mathring{P}_i^*$. We claim we can use the same side Q for all i. Replacing all P_i by P_i^* if necessary, we may assume $Q \subset \mathring{P}_i$ for at least one $i \leq k$ (this is because the $\mathring{P}_i^*$ also form a chain).

If $Q \subset \mathring{P}_1$ then $Q \subset \mathring{P}_j$ for all j and we are done. Otherwise, take the minimal i with $Q \subset \mathring{P}_i$. Since $Q \not\subset \mathring{P}_{i-1}$ we must have $Q^* \subset \mathring{P}_{i-1}$ or $Q^* \subset \mathring{P}^*_{i-1}$ or $Q \subset \mathring{P}^*_{i-1}$. If $Q^* \subset \mathring{P}_{i-1}$

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

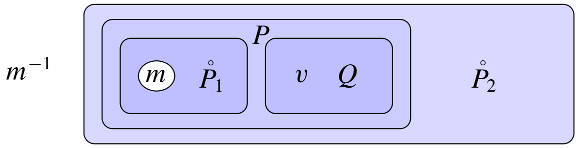

Fig. 6. Proposition 4·22.

then $Q \supset P^*_{i-1} \supset P^*_i$, contradicting $Q \subset P_i$. If $Q^* \subset \mathring{P}^*_{i-1}$ then $\mathring{P}_i \supset Q \supset P_{i-1}$ so Q splits m_{i-1}, contradicting $\mathcal{Q} \in \Pi - \Pi_{[m]_\star}$. So we must have $Q \subset P^*_{i-1}$, i.e. $Q \subset \mathring{P}_i \cap \mathring{P}^*_{i-1}$.

The strategy in several upcoming proofs will be to replace some $\mathcal{Q} \in \Pi$ by a "better" ΓW-partition $\mathcal{P}$ compatible with everything in Π except $\mathcal{Q}$, where the feature that makes $\mathcal{P}$ better will depend on the context. The following proposition gives us our main tool for doing this. The setup for this proposition is illustrated in Figure 6.

PROPOSITION 4·22. *Let m be a principal vertex of Γ, $\mathcal{P}_1 \in \Pi_m$ and $\mathcal{P}_2 \in \Pi_{[m]_\star}$, and choose sides P_1, P_2 with $\mathring{P}_1 \subset \mathring{P}_2$. Suppose $u \leq_\circ m$ is contained in $P_2 \cap P^*_1$. Let Q be a largest subset of $P_2 \cap P^*_1$ which is in $\Pi^\pm$ and is based at some $v \sim u$; if there are no such subsets, set $Q = \{u\}$. Let $\mathcal{P}$ be the ΓW-partition determined by $P = P_1 \cup Q$.*

If $\mathcal{R} \in \Pi - \Pi_{[m]_\star}$ is not compatible with $\mathcal{P}$, then some side R of $\mathcal{R}$ is contained in $\mathring{P}_2 \cap \mathring{P}^*_1$, contains Q and is based at some s with $s >_\circ u$.

Proof. Note that $\mathcal{P}$ is based at m. Since $\mathcal{R}$ is not compatible with $\mathcal{P}$ and $s \not\sim_\star m$, s and m do not commute.

Since s and m do not commute, then by Lemma 4·21 $\mathcal{R}$ has a side R in $\mathring{P}_1$, $\mathring{P}_2 \cap \mathring{P}^*_1$ or $\mathring{P}^*_2$. If either $R \subset \mathring{P}_1$ or $R \subset \mathring{P}^*_2$ then $\mathcal{R}$ is compatible with $\mathcal{P}$, so we must have $R \subseteq \mathring{P}_2 \cap \mathring{P}^*_1$. Since $\mathcal{R}$ is compatible with $\mathcal{Q}$ but not with $\mathcal{P}$ we must have $R \supset Q$.

Since Q was of maximal size, $s \notin [v] = [u]$. Thus either $v <_\circ s$ or there is some $x \in lk(u) \subseteq st(m)$ which is not in $lk(s)$. Such an x would be adjacent to both v and m so v and m would be in the same component of $\Gamma - lk(s)$, contradicting the fact that $\mathcal{R}$ separates m from v.

COROLLARY 4·23. *Let Π be a maximal collection of compatible ΓW-partitions and $[m]$ a non-abelian equivalence class of principal vertices of Γ. Then for any $m \in [m]$ the subset $\Pi_{[m]}$ can be replaced by a new set of partitions of the same size to obtain a compatible collection Π' with $\Pi'_{[m]} = \Pi'_m$.*

Proof. Fix $m \in [m] = [m]_\circ$ and suppose $\Pi_m = \{\mathcal{P}_1, \ldots \mathcal{P}_k\} \neq \emptyset$. Let $P_1 \subset \cdots \subset P_k$ be the sides of the $\mathcal{P}_i$ containing m.

Suppose $\mathcal{Q} \in \Pi_{[m]} \setminus \Pi_m$ is based at $n \sim m$. Since $[m]$ is nonabelian, n does not commute with m, so it must have a side Q contained in $P_i \cap P^*_{i-1}$ for some i. Take Q maximal with respect to inclusion among all such sides in $P_i \cap P_{i-1}$. Now take M maximal among all such sides properly contained in Q; if there is no such M, set $M = \{n\}$. By Proposition 4·22 (applied to $[m]_\star = \{m\}$), if some partition $\mathcal{R} \in \Pi \setminus \Pi_m$ is not compatible with the ΓW-partition $\mathcal{P}$ determined by $P_i \cup M$, then either it is equal to $\mathcal{Q}$ or it is based at some s with $n <_\circ s$. i.e. $lk(n) \subsetneq lk(s)$. But n is principal, so there is no such s. Since $\mathcal{Q}$ is the only partition in Π not compatible with $\mathcal{P}$, we may replace $\mathcal{Q}$ by $\mathcal{P}$ to obtain a new collection of the same size. We can continue this process until $\Pi_{[m]} = \Pi_m$.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

Definition 4·24. A ΓW-partition $\mathcal{P}$ based at m is *principal* if m is a principal vertex of Γ.

THEOREM 4·25. *Let L be the set of principal vertices of Γ. Then $U(A_\Gamma)$ contains a free abelian subgroup of rank $M(L)$.*

Proof. Let Π be a maximal compatible collection of principal ΓW-partitions, i.e. a collection of size $M(L)$.

By Corollary 4·23 we may assume $\Pi_{[m]} = \Pi_m$ for all nonabelian equivalence classes $[m]$. Using m as multiplier for each $\mathcal{P} \in \Pi_{[m]}$, the associated outer Γ-Whitehead automorphisms $\varphi(\mathcal{P}, m)$ pairwise commute.

If $\mathcal{P}$ and $\mathcal{Q}$ in Π are based at m and n with $[m, n] = 1$ then $\varphi(\mathcal{P}, m)$ and $\varphi(\mathcal{Q}, n)$ commute.

If $\mathcal{P}$ and $\mathcal{Q}$ in Π are based at m and n with $[m, n] \neq 1$ then Lemma 4·20 implies that $\mathcal{P}$ does not split n and $\mathcal{Q}$ does not split m, so $\varphi(\mathcal{P}, m)$ and $\varphi(\mathcal{Q}, n)$ commute by Theorem 4·2.

We now have a collection of pairwise-commuting infinite-order outer automorphisms $\varphi(\mathcal{P}_i, m_i)$ of size equal to $M(L)$, and we need to show they are independent. Choose sides P_i for $\mathcal{P}_i$ containing m_i, and set

$$\Phi = \phi(P_1, m_1)^{n_1} \cdots \phi(P_k, m_k)^{n_k}.$$

We must show that if Φ is inner then all $n_i = 0$.

Let $\{v_1, \ldots, v_\ell\}$ be the distinct m_i and define

$$\Phi_j = \prod_{m_i = v_j} \phi(P_i, m_i)^{n_i},$$

so $\Phi = \Phi_1 \ldots \Phi_\ell$. By Proposition 4·6 if any of the Φ_j are inner then the associated n_i are zero; in particular, if $\ell = 1$ we are done. So we may assume no Φ_i is trivial and $\ell > 1$.

If all v_i have the same star, then we are done by Proposition 4·13. Otherwise without loss of generality we may assume there is $x \in st(v_2)$ with $x \notin st(v_1)$.

Replacing $\phi(P_i, m_i)$ by $\phi(P_i^*, m_i^{-1})$ whenever $x \in P_i$ (which does not affect their images in $U(A_\Gamma)$) we may assume $\Phi(x) = xU$ for some word U in the m_i. Since Φ is conjugation by some element W, this implies $U = 1$, so W is in the centralizer of x, which is generated by $st(x)$. Since $v_1 \notin st(x)$, v_1 does not appear in any reduced expression for W.

Since Φ_1 is not trivial there is some vertex y with $\Phi_1(y) = v_1^a y v_1^b$, where a and b are not both zero. If we set $\Psi = \Phi_2 \cdots \Phi_\ell$ then $\Phi(y) = \Psi\Phi_1(y) = \Psi(v_1)^a \Psi(y)\Psi(v_1)^b$.

By Corollary 4·3, each $\phi(P_i, m_j)$ acts either trivially or as conjugation by m_i on each m_j. Thus $\Psi(v_1)$ is conjugate to v_1 by a word U in $v_2, \ldots, v_\ell$. So we have

$$\begin{aligned}\Phi(y) &= \Psi\Phi_1(y)\\ &= \Psi(v_1)^a \Psi(y)\Psi(v_1)^b\\ &= U^{-1}v_1^a U \Psi(y) U^{-1} v_1^b U.\end{aligned}$$

We also know that $\Phi(y) = W^{-1} y W$ for some W that does not contain the letter v_1. But v_1 does not commute with y so in order for the powers of v_1 in the expression for $\Phi(y)$ above to cancel it must be true that a reduced word representing $\Psi(y)$ does not contain y. In order for this to happen some $\phi(P_i, m_i)$ must have multiplier $m_i = y$. But if $y = m_i$ then $\Psi(y)$ is conjugate to y by Corollary 4·3 so the reduced word representing $\Psi(y)$ *does* contain y, giving a contradiction.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

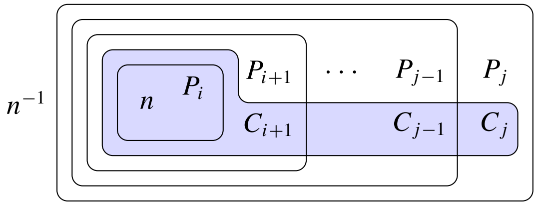

Fig. 7. Proof of Lemma 4·27.

Definition 4·26. Suppose $\varphi(\mathcal{P}_1, m_1), \ldots, \varphi(\mathcal{P}_k, m_k)$ generate a free abelian subgroup of $U(A_\Gamma)$, and let $\Pi = \{\mathcal{P}_1, \ldots, \mathcal{P}_k\}$. Suppose $[m]$ is abelian and $\Pi_{[m]} \neq \emptyset$. Then Π is $[m]$-*complete* if it contains every ΓW-partition $\mathcal{Q}$ such that:

(i) $\mathcal{Q}$ is based at some $n \in [m]$ and
(ii) $\varphi(\mathcal{Q}, n)$ commutes with all $\varphi(\mathcal{P}_i, m_i)$.

If Π is not $[m]$-complete, it can be completed by adding all possible $\mathcal{Q}$ satisfying the above conditions. The base n of any such $\mathcal{Q}$ is unique since $[m] = [m]_\star$, so $\varphi(\mathcal{Q}, n)$ is determined by $\mathcal{Q}$. All of these $\varphi(\mathcal{Q}, n)$ can be added to $\{\varphi(\mathcal{P}_i, m_i)\}$ to generate an abelian subgroup of possibly larger rank.

LEMMA 4·27. *Suppose $\varphi(\mathcal{P}_1, m_1), \ldots, \varphi(\mathcal{P}_\ell, m_\ell)$ generate a free abelian subgroup G and let $\Pi = \{\mathcal{P}_1, \ldots, \mathcal{P}_\ell\}$ be $[m]$-complete for some abelian $[m]$. Then Π contains a subcollection Π^c such that $\Pi^c_{[m]}$ is a compatible collection of ΓW-partitions and the $\varphi(\mathcal{P}_i, m_i)$ for $\mathcal{P}_i \in \Pi^c$ generate the same abelian subgroup G.*

Proof. Let Π_0 be a maximal compatible subcollection of $\Pi_{[m]}$. If $\mathcal{P} \in \Pi_0$ is based at $v \in [m]$ then by Lemma 4·8 $\mathcal{P}_{v,w}$ is also in Π_0 for every $w \in [m]$, since Π is $[m]$-complete.

Now consider $\mathcal{Q} \in \Pi_{[m]} \setminus \Pi_0$, based at some $n \in [m]$. By Lemma 4·5 we may choose sides P_i of the $\mathcal{P} \in \Pi_0$ based at n such that

$$\{n\} = P_0 \subset P_1 \subset \cdots \subset P_k \subset P_{k+1} = V^{\pm} \setminus lk(n)^{\pm} \setminus \{n^{-1}\}.$$

Take the largest $i \geq 0$ such that the side Q of $\mathcal{Q}$ containing n also contains P_i, and the smallest $j \leq k+1$ such that $Q \subset P_j$. For each ℓ with $i+1 \leq \ell \leq j$ let $C_\ell = P_\ell \cap Q$ (see Figure 7).

Then $P'_\ell = P_{\ell-1} \cup C_\ell$ is a ΓW-subset, and the ΓW-partition $\mathcal{P}'_\ell$ it determines is compatible with all $\mathcal{P} \in \Pi_0$. Furthermore, since $\varphi(\mathcal{Q}, n)$ commutes with all $\varphi(\mathcal{P}_j, m_j)$ it follows that $\varphi(\mathcal{P}'_\ell, n)$ does as well, so $\mathcal{P}'_\ell$ must be in Π_0 since Π is $[m]$-complete and Π_0 is maximal. Now

$$\varphi(\mathcal{Q}, n) = \left(\prod_{\ell=i+2}^{j} \varphi(\mathcal{P}'_\ell, n) \varphi(\mathcal{P}_{\ell-1}, n)^{-1} \right) \varphi(\mathcal{P}'_{i+1}, n)$$

so we may eliminate $\mathcal{Q}$ from $\Pi_{[m]}$ without affecting the abelian subgroup G. Continuing, we eliminate all partitions in $\Pi_{[m]}$ that are not in Π_0. Then $\Pi^c = \Pi \setminus \Pi_{[m]} \cup \Pi_0$ is the required collection.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

THEOREM 4·28. *Any free abelian subgroup of $U(A_\Gamma)$ generated by Γ-Whitehead automorphisms has rank at most $M(L)$.*

Proof. Suppose $\varphi(\mathcal{P}_1, m_1), \ldots, \varphi(\mathcal{P}_k, m_k)$ generate a free abelian subgroup G of rank $r > M(L)$, and let $\Pi = \{\mathcal{P}_1, \ldots, \mathcal{P}_k\}$. If $[m_i, m_j] \neq 1$ then $\mathcal{P}_i$ is compatible with $\mathcal{P}_j$ by Theorem 4·2. If $[m_i, m_j] = 1$ but $m_i \not\sim_\star m_j$ then $\mathcal{P}_i$ is compatible with $\mathcal{P}_j$ by the definition of compatibility. So the only incompatible pairs in Π live in the same $\Pi_{[m]}$ for some m with $[m] = [m]_\star$.

Fix such an m and add all necessary partitions to $\Pi_{[m]}$ so that Π is $[m]$-complete. The corresponding free abelian group contains G as a subgroup. By Lemma 4·27 there is a subcollection Π^c of Π such that $\Pi^c_{[m]}$ is a compatible collection and the corresponding group generated is the same, i.e. it still contains G as a subgroup. After repeating this for each equivalence class with $[m] = [m]_\star$ we may assume that Π is a compatible collection.

Now choose the $\varphi(\mathcal{P}_i, m_i)$ so that Π has the smallest possible number of non-principal partitions for a such a collection. By Corollary 4·23 we may assume $\Pi_{[m]} = \Pi_m$ for each principal nonabelian equivalence class $[m]$ and a choice of representative m. Let $\mathcal{Q} \in \Pi$ be a non-principal partition, based at a vertex u which has maximal link among the non-principal bases. Since u is non-principal, $u <_\circ m$ for some $m \in L$. Let $\Pi_{[m]} = \{\mathcal{P}_1, \ldots, \mathcal{P}_\ell\}$ $(\ell \geq 0)$ and choose sides P_i of $\mathcal{P}_i$ so that the $\mathring{P}_i$ are nested. By Lemma 4·21, there is a side Q of $\mathcal{Q}$ such that $Q \subset \mathring{P}_i \cap \mathring{P}^*_{i-1}$ for some $i \leq \ell + 1$. Maximise Q with respect to inclusion over all non-principal partitions in $\Pi_{[u]}$ with sides in $\mathring{P}_i \cap \mathring{P}^*_{i-1}$. By Proposition 4·22 if a partition $\mathcal{R}$ based at $s \not\sim m$ is incompatible with $P_{i-1} \cup Q$, then $\mathcal{R}$ has a side $R \subset \mathring{P}_i \cap \mathring{P}^*_{i-1}$ with $R \supset Q$ and $u <_\circ s$. Maximality of u tells us that $\mathcal{R}$ must in fact be a principal partition, so by replacing m with s and repeating the above arguments we will reach a point where no such incompatible $\mathcal{R}$ exists. At this point we claim that $P_{i-1} \cup Q = P_i$: if $P_{i-1} \cup Q$ was not in Π we could replace $\mathcal{Q}$ with the partition determined by $P_i \cup Q$ to arrive at a collection of the same size k but one fewer non-principal partition.

Since $\varphi(\mathcal{P}_i, m)$ commutes with $\varphi(\mathcal{Q}, n)$, $P_i = P_{i-1} \cup Q$ must contain both u and u^{-1}, i.e. $u^{-1} \in P_{i-1}$. This implies that $\mathcal{P}_{i-1}$ splits u, contradicting the commutativity conditions of Theorem 4·2.

COROLLARY 4·29. *Let G be an abelian subgroup of $U(A_\Gamma)$ of rank $M(L)$, freely generated by $\{\varphi(\mathcal{P}_i, m_i)\}$ with m_i principal. Suppose m_i is not maximal for some i, say $m_i <_\star w$. Then $\varphi(\mathcal{P}, w) \in G$, where $\mathcal{P}$ is the partition obtained from $\mathcal{P}_i$ by exchanging m_i for w, as defined in Remark 4·9.*

Proof. Let $\Pi = \{\mathcal{P}_i\}$, and let Π' be the $[w]$-completion of Π. Remark 4·9 shows that

$$\mathcal{P} = \mathcal{P}_i \setminus lk(w)^{\pm} \cup \{w\}$$

is contained in Π'. If G' is the corresponding free abelian subgroup then $G \leq G'$. Theorem 4·28 tells us that $rank(G) = rank(G')$ so $G = G'$, thus $\varphi(\mathcal{P}, w) \in G$.

5. *Virtual cohomological dimension*

By Theorem 3·10 we know $dim(K_\Gamma) = M(V)$ is an upper bound on the VCD of $U(A_\Gamma)$ and $M(L)$ is a lower bound by Theorem 4·25. In this section we give conditions under which $M(V) = M(L)$.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

LEMMA 5·1. *Let Γ be a connected graph and $u \not\sim v$ vertices with $d_\Gamma(u, v) \neq 2$ (where d_Γ is the length of a shortest path in Γ). Then any partition based at u is compatible with any partition based at v. In particular,*

$$M(u, v) = M(u) + M(v).$$

Proof. If $d_\Gamma(u, v) = 1$ then u and v commute. Since we are assuming $u \not\sim v$, the partitions are compatible. If $d_\Gamma(u, v) \geq 3$, let C_v denote the element of $\mathbf{I}(u)$ containing v (and v^{-1}) and C_u the element of $\mathbf{I}(v)$ containing u. Then C_v contains all elements of $\mathbf{I}(v)$ other than C_u and C_u contains all elements of $\mathbf{I}(u)$ other than C_v. This implies that any thick partition of $\mathbf{I}(v)$ separating v from v^{-1} is compatible with any thick partition of $\mathbf{I}(u)$ separating u from u^{-1}.

THEOREM 5·2. *Let Γ be a graph and $L \subseteq V$ its set of principal vertices. Suppose that every $u \in V \setminus L$ satisfies*

$$\textit{All principal } m \textit{ with } m >_\circ u \textit{ are in the same component of } \Gamma - lk(u). \tag{5·1}$$

Then $M(V) = M(L)$.

Proof. Let Π be a maximal pairwise-compatible collection of ΓW-partitions with $M(V)$ elements. We will produce a new collection of the same size in which all partitions are principal. By Corollary 4·23 we may assume $\Pi_{[m]} = \Pi_m$ for each principal nonabelian equivalence class and a choice of representative m.

Let u be maximal among all $u \in V \setminus L$ with $\Pi_u \neq \emptyset$. By Lemma 4·12, for any principal $m >_\circ u$ we have a nest

$$\mathring{P}_0 = \emptyset \subset \mathring{P}_1 \subseteq \cdots \subseteq \mathring{P}_k \subset \emptyset^* = \mathring{P}_{k+1},$$

where $\mathcal{P}_1, \ldots, \mathcal{P}_k$ ($k \geq 0$) are the elements of $\Pi_{[m]}$ and P_i is a side of $\mathcal{P}_i$. By Lemma 4·21, each $Q \in \Pi_{[u]}$ has a side $Q^\times \subset \mathring{P}_i \cap \mathring{P}^*_{i-1}$ for some i. For each m, let $i(m)$ be the smallest index such that $\mathring{P}_{i(m)}$ contains one of these $Q^\times$. Choose m such that $|P_{i(m)}|$ is minimal. Among the $Q \in \Pi_u$ with $Q^\times \subset \mathring{P}_{i(m)}$ choose one with $Q = Q^\times$ maximal.

The union $P_{i-1} \cup Q$ determines a ΓW-partition $\mathcal{P}$ based at m. If there is some $\mathcal{R} \in \Pi$ not compatible with $\mathcal{P}$ then by Proposition 4·22 $\mathcal{R}$ has a side $R \subset \mathring{P}_i$ containing Q and disjoint from $\mathring{P}_{i-1}$, and $\mathcal{R}$ is based at some $s >_\circ u$. By our choice of u this implies that s is principal, so either Q or Q^* is somewhere in the nest $\emptyset \subset \mathring{R}_1 \subset \cdots \subset \mathring{R}_\ell \subset \emptyset^*$ associated with s. Since $s >_\circ u$, Q does not contain s (if it did, it would split s and we would have $s \leq_\circ u$). Therefore Q is in the nest. Since $R = R_j$ for some j, we have $Q \subset R_j \subsetneq P_i$, contradicting minimality of $|P_i|$.

Now take a proper subset $M \subsetneq Q$ in $\Pi^\pm_u$ of maximal size. If there is no such M, take $M = \{u\}$. Proposition 4·22 applied to $\Pi \setminus Q$ shows that if $\mathcal{R}$ is not compatible with the ΓW-partition $\mathcal{P}'$ determined by $P_{i-1} \cup M$ then either $\mathcal{R} = Q$ or $\mathcal{R}$ has a side $R \subset \mathring{P}_i$ containing M and disjoint from $\mathring{P}_{i-1}$, and $\mathcal{R}$ is based at some $s >_\circ u$. By our choice of u, s is principal. But s and m are on different sides of Q, contradicting our hypothesis that all principal $v >_\circ u$ are in the same component of $\Gamma - lk(u)$. We can now replace Q by $\mathcal{P}'$ to get a new collection of the same size, with one fewer non-principal partition. Continuing, we can replace all non-principal partitions by principal partitions, showing $M(V) = M(L)$.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

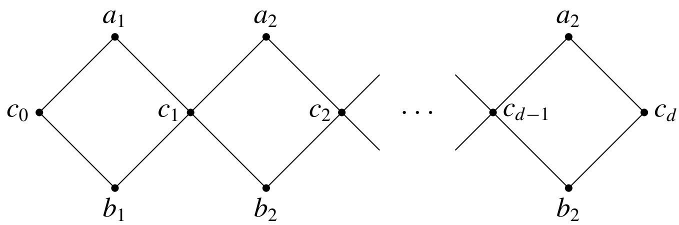

Fig. 8. String of diamonds.

The following is a special case of Theorem 5·2 which is often very easy to check.

COROLLARY 5·3. *If every non-principal equivalence class of vertices in* Γ *is* $<_\circ$ *at most one principal equivalence class, then* $M(V) = M(L)$.

6. *Examples*

In this section we give a few examples illustrating both the utility and the limits of Theorem 5·2.

Example 6·1. Let Γ be the graph with n vertices and no edges, i.e. $A_\Gamma = \mathbb{F}_n$.

Here there are no twists so $U(A_\Gamma) = Out(A_\Gamma)$. Since all vertices are maximal and equivalent, Corollary 4·23 implies $M(V) = M(m)$ for any choice of vertex m. Since $M(m) = 2n - 3$ (see Example 4·17), this gives (the correct) lower bound of $2n - 3$ for the VCD of $Out(F_n)$.

Example 6·2. Let Γ be a string of d diamonds, as shown in Figure 8.

Again there are no twists, so $U(A_\Gamma) = Out(A_\Gamma)$. The only non-principal vertices are c_0 and c_d and there are no ΓW-partitions based at either of these, so $M(V) = M(L)$. Let Π be a collection of size $M(V)$. We have $[a_i] = \{a_i, b_i\}$ for each $1 \leq i \leq d$ so by Corollary 4·23 we may assume $|\Pi_{\{a_i, b_i\}}| = |\Pi_{a_i}|$ for each i. We have $M(a_i) = 3$ if $2 \leq i \leq d - 1$, $M(a_d) = M(a_1) = 2$, $M(c_i) = 1$ if $2 \leq i \leq d - 2$ and $M(c_1) = M(c_{d-1}) = 2$. Therefore,

$$M(V) \leq 3(d - 2) + 4 + d - 3 + 4 = 4d - 1.$$

It is easy to find a collection of ΓW-subsets with $4d - 1$ elements (one is given explicitly in [**5**]), so in fact $M(V) = 4d - 1$ and $dim(K_\Gamma) = M(V) = \text{VCD}(Out(A_\Gamma))$.

Example 6·3. Let Γ be the graph in Figure 9.

Since Γ is a tree, the VCD of $Out(A_\Gamma)$ is equal to $e + 2\ell - 3 = 7 + 8 - 3 = 12$ [**3**]. The only twists are given by the leaf transvections. These form a normal free abelian subgroup of rank 4 (the number of leaves), with quotient $U(A_\Gamma)$, so it is natural to expect that the VCD of $U(A_\Gamma)$ is 8.

There are no ΓW-partitions based at any of the b_i since $\Gamma \setminus st(b_i)$ has only one component. Any partition based at v_1 is compatible with any partition based at a different vertex by Lemma 5·1, since $a_i, v_0 \in st(v_1)$ for each i. We have $M(v_1) = |\mathbf{I}(v_1)| - 3 = 5$. Now consider

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

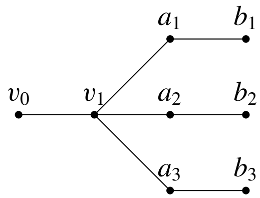

Fig. 9. Γ is a tree with $M(V) = 10$, but there is no free abelian subgroup of rank 10 generated by compatible collections of Whitehead automorphisms.

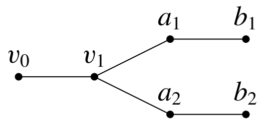

Fig. 10. For this tree $M(L) = 5$ but $\dim(K_\Gamma) = 6$.

partitions based at a_1, a_2 or a_3. Choose any one such partition $\mathcal{P}$, say based at a_1. Then for each a_i there are at most two choices of partition compatible with $\mathcal{P}$ since the side of $\mathcal{P}$ not containing a_i must be disjoint from the side of Q not containing a_1. Say a choice Q is based at m, then by repeating this argument on disjoint sides there is at most one choice of partition compatible with both $\mathcal{P}$ and Q, so $M(a_1, a_2, a_3) = 3$ and the largest possible number of ΓW-partitions based at principal vertices is $5 + 3 = 8$. Since $M(v_0) = 2$, we have $M(V) \leq 10$. In fact equality holds since the following list of five ΓW-subsets determines a compatible collection of distinct ΓW-partitions based in $\{a_1, a_2, a_3, v_0\}^{\pm}$:

$$\{a_1, v_0\}, \{v_0, a_1, a_1^{-1}, b_1, b_1^{-1}\}, \{a_2, v_0, a_1, a_1^{-1}, b_1, b_1^{-1}\}, \{v_0^{-1}, a_3, a_3^{-1}, b_3, b_3^{-1}\}, \{a_3^{-1}, v_0\}.$$

Thus $M(L) = 8 \leq \text{VCD}(U(A_\Gamma)) \leq dim K_\Gamma = 10$.

Example 6·4. A similar but slightly simpler example is when Γ is the tree in Figure 10.

A quick check yields $M(V) = M(v) + M(u, a_1, a_2)$ and $M(v) = 3$. Furthermore, arguing in the same fashion tells us that $M(u, a_1, a_2) \leq 3$, with a possible $\Pi_{\{u,a_1,a_2\}}$ being the ΓW-partitions determined by:

$$\{a_1, u\}, \{u, a_1, a_1^{-1}, b_1, b_1^{-1}\}, \{a_2^{-1}, u^{-1}\}.$$

Thus, $M(V) = 6$ so $\dim(K_\Gamma) = 6$ but we only find a subgroup $\mathbb{Z}^5 \leq U(A_\Gamma)$. In the following section it is shown that this particular Γ has $VCD(U(A_\Gamma)) = 5$.

7. *Reducing the dimension of K_Γ*

In this section we show that, in some cases with $M(V) > M(L)$, we can find an invariant contractible subcomplex of K_Γ of smaller dimension. We use the weak notion of compatibility throughout since that is what is actually used in [**5**] to define and prove contractibility of K_Γ.

Definition 7·1. A graph Γ is *barbed* if for all non-principal vertices u, $d_\Gamma(u, v) = 2$ implies $u <_\circ v$.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

LEMMA 7·2. *If Γ is barbed then every non-principal equivalence class is minimal and has only one element. Furthermore any ΓW-partition based at a non-principal element splits* only *that element.*

Proof. This is immediate.

All of the graphs in Section 6 are barbed. Examples 6·3 and 6·4 are examples of barbed graphs with $M(V) > M(L)$. We claim that if Γ is barbed and $M(V) > M(L)$, then K_Γ equivariantly deformation retracts to a smaller-dimensional complex. Specifically, every cube in K_Γ of dimension $M(V)$ has a free face, and the set of these free faces is invariant under the action of $U(A_\Gamma)$.

In Lemmas 7·3 to 7·7 we fix a collection Π of pairwise weakly compatible ΓW-partitions with $M(V) > M(L)$ elements. Recall that $\Pi^\pm$ denotes the collection of all sides of elements of Π.

LEMMA 7·3. *Let $Q \in \Pi^\pm$ be a non-principal ΓW subset, based at some $u \in Q$. If Q contains some $m \geq_\circ u$ other than u, then Q properly contains some $N \in \Pi^\pm$ with $u \in N$.*

Proof. Suppose the lemma is false, i.e. no $N \in \Pi^\pm$ is properly contained in Q and also contains u.

If there are no elements at all of $\Pi^\pm$ properly contained in Q, then the ΓW-partition determined by $\{u, m\}$ is (weakly) compatible with all elements of Π, contradicting maximality of Π.

Now take a largest $P \in \Pi^\pm$ properly contained in Q, based at some $n \in P$. If $n \not\geq_\circ u$ then there is some $v \notin lk(n)$ with $v \in lk(u) \subset lk(m)$, so u, v and m are all in the same component of $\Gamma - lk(n)$, so u, v, m and their inverses are all on the same side of P, i.e. all are outside P. If this is true for all largest P contained in Q we can add the partition determined by $\{u, m\}$ to Π, again contradicting maximality of Π.

If some largest P is based at a vertex $n \geq_\circ u$ then $P \cup \{u\}$ is a ΓW-subset and the corresponding ΓW-partition is (weakly) compatible with all elements of Π, once again contradicting maximality of Π.

Definition 7·4. A ΓW-partition $Q \in \Pi$ is *irreplaceable in* Π if Q is the only ΓW-partition compatible with all elements of $\Pi \setminus Q$.

Definition 7·5. A ΓW-partition $Q \in \Pi$ based at u is *sandwiched in* Π if there are principal $m \in Q$ and $n \in Q^*$ with $m, n >_\circ u$ such that both $Q_m = Q \setminus (\{m\} \cup lk(m)^\pm)$ and $Q_n^* = Q^* \setminus (\{n\} \cup lk(n)^\pm)$ are in $\Pi^\pm$ (See Figure 11).

LEMMA 7·6. *If a non-principal partition $Q \in \Pi$ is sandwiched in Π then Q is irreplaceable in Π.*

Proof. If Q_m and Q_n^* are both in $\Pi^\pm$, then any replacement for Q cannot have a side contained in Q_m or Q_n^* (by maximality of Π) and cannot split both m and n (since m and n are on different sides of Q so are not equivalent). Since Q is the only ΓW-partition that satisfies these conditions, Q is irreplaceable.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

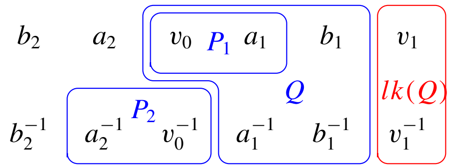

Fig. 11. Q determines a ΓW-partition for the tree in Figure 10 that is sandwiched, with $Q_{a_1^{-1}} = P_1$ and $Q^*_{a_2} = P_2$.

LEMMA 7·7. *Let Γ be a barbed graph and $Q \in \Pi^{\pm}$ innermost among non-principal sides, based at some $u \in Q$. If Q is not sandwiched in Π, then Q is replaceable by a principal partition.*

Proof. Since Γ is barbed, there are principal elements bigger than u on both sides of Q. By Lemma 7·3 there is a proper subset M of Q that is in $\Pi^{\pm}$ and contains u; take a largest such M. Since Q is an innermost non-principal subset, M must be principal based at some m, which must be $>_\circ u$ since M separates u from u^{-1}. Unless $M = Q_{m^{-1}} = Q \setminus \{m^{-1}\} \setminus lk(m)^{\pm}$, the set $Q^* \cup M \setminus lk(m)^{\pm}$ has at least two elements on each side so determines a principal ΓW-partition which can replace Q.

If $M = Q_{m^{-1}} = Q \setminus \{m^{-1}\} \setminus lk(m)^{\pm}$, we consider the other side Q^* of Q. By Lemma 7·3 there is also a ΓW-subset $N \subsetneq Q^*$ with $u^{-1} \in N$. Take a maximal such N. If N is based at u^{-1} then $N \cup Q_{m^{-1}}$ is principal, based at m, and can replace Q. So suppose N is based at $n \neq u^{-1}$. Since N splits u we must have $u \leq_\circ n$, and since Γ is barbed $[u] = \{u\}$ so in fact $n >_\circ u$ and n must be maximal. If $N = Q^*_{n^{-1}} = Q^* \setminus \{n^{-1}\} \setminus lk(n)^{\pm}$, then Q is sandwiched, contradicting our assumption. Therefore the set $Q \cup N \setminus lk(n)^{\pm}$ is a ΓW-subset and the corresponding principal ΓW-partition can replace Q.

THEOREM 7·8. *Let Γ be a barbed graph with $M(V) > M(L)$. Then the dimension of K_Γ is strictly larger than* VCD$(U(A_\Gamma))$.

Proof. Let Π be a maximal collection of weakly compatible ΓW-partitions with $M(V) > M(L)$ elements. Then Π determines a cube $c(\emptyset, \Pi)$ in K_Γ of dimension $M(V)$. We will find a free face of this cube, namely $c(\emptyset, \Pi \setminus Q)$ for some non-principal Q and use it to collapse the cube. We can do this equivariantly for all such cubes in all of K_Γ, thereby reducing the dimension of K_Γ by 1.

The cube $c(\emptyset, \Pi \setminus Q)$ is a free face of $c(\emptyset, \Pi)$ if and only if Q is irreplaceable. So we are looking for an irreplaceable Q in Π. Let $R \in \Pi^{\pm}$ be an innermost non-principal ΓW-subset. If the corresponding ΓW-partition $\mathcal{R}$ is sandwiched, then it is irreplaceable, by Lemma 7·6 so we may take $Q = \mathcal{R}$. If it is not sandwiched, then it can be replaced by a principal ΓW-partition $\mathcal{P}$, by Lemma 7·7, to form a new maximal collection Π'. This new collection has the same size, so must still contain a non-principal ΓW-partition.

Claim. If a non-principal $\mathcal{S} \in \Pi'$ based at v is sandwiched between S_m and S^*_n in Π', then it was already sandwiched in Π, so is irreplaceable in Π by Lemma 7·6.

Proof of claim. If $\mathcal{S}$ is sandwiched in Π' but not in Π then either S_m or S^*_n must be equal to the ΓW-subset we used in Lemma 7·7 to replace Q. In all cases this has a side of the form

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

$S = T \cup M$ where T is non-principal based at u and M is principal with $u \in M$. It follows that M is based at m (if $S = S_m$) and $u, v <_\circ m$ (or at n if $S = S_n^*$ and $u, v <_\circ n$) and that $u \neq v$. But then M splits both u and v, which cannot happen in a barbed graph.

Now let S be an innermost non-principal side in $\Pi'^{\pm}$. If $\mathcal{S}$ is sandwiched in Π' then by the claim it was already sandwiched in Π, so is irreplaceable in Π and we may take $\mathcal{Q} = \mathcal{S}$. If it is not sandwiched, we can replace it by a principal partition by Lemma 7·7. We continue replacing innermost non-principal sides until we encounter one that is sandwiched (which must exist since $M(V) > M(L)$) and hence irreplaceable.

As shown in [5], the star of a Salvetti S_Γ in K_Γ is the union of the cubes with S_Γ as a vertex, and these cubes are identified with weakly compatible collections of ΓW-partitions. The stabiliser S_Γ under the action of $U(A_\Gamma)$ is isomorphic to the subgroup generated by graph automorphisms and inversions. The effect of such an automorphism on the cubes in the star is to permute the labels of $V^\pm$. Since incidence relations are preserved, any such automorphism sends a ΓW-partition to the "same" ΓW-partition with the labels permuted. Since irreplaceable partitions are characterized by being sandwiched, such an automorphism sends sandwiched partitions to sandwiched partitions, and thus sends free faces to free faces. Thus collapsing these free faces is an equivariant operation, giving an equivariant deformation retraction of K_Γ.

Acknowledgements. We thank Benjamin Brück and Ric Wade for extremely useful comments on the first version of this paper. The second author was partially supported by a Royal Society Wolfson award.

REFERENCES

[1] C. BREGMAN. Automorphisms and homology of non-positively curved cube complexes. arXiv:1609.03602.

[2] M. R. BRIDSON and K. VOGTMANN. *Automorphisms of free groups, surface groups and free abeliangroups.* Problems on Mapping Class Groups and Related Topics, ed. by B. Farb, Proc. Symp. Pure Math. **74** (Amer. Math. Soc., Providence RI, 2006).

[3] K. BUX, R. CHARNEY and K. VOGTMANN. Automorphism groups of RAAGs and partially symmetric automorphisms of free groups. *Groups Geom. Dyn.* **3** (2009), no. 4, 541–554.

[4] R. CHARNEY. An Introduction to Right-angled Artin Groups. *Geom. Dedicata* **125** (2007), 141–158.

[5] R. CHARNEY, N. STAMBAUGH and K. VOGTMANN. Outer Space for Untwisted Automorphisms of right-angled Artin Groups. *Geom. Topol.* **21** (2017), no. 2, 1131–1178.

[6] R. CHARNEY and K. VOGTMANN. *Subgroups and Quotients of Automorphism Groups of RAAGs.* Low-dimensional and symplectic topology, 9–27, Proc. Sympos. Pure Math., **82** (Amer. Math. Soc., Providence, RI, 2011).

[7] J. CRISP, R. CHARNEY and K. VOGTMANN. Automorphism groups of two-dimensional right-angled Artin groups. *Geometry and Topology* **11** (2007), 2227–2264.

[8] M. CULLER and K. VOGTMANN. Moduli of graphs and automorphisms of free groups. *Invent. Math.* **84** (1986), no. 1, 91–119.

[9] M. B. DAY. Peak reduction and finite presentations for automorphism groups of right-angled Artin groups. *Geom. Topol.* **13** (2009), no. 2, 817–855.

[10] M. B. DAY and R. D. WADE. Relative automorphism groups of right-angled Artin groups. *J. Topology* **12** (2019), no. 3, 759–798.

[11] V. GUIRARDEL and A. SALE. Vastness properties of automorphism groups of RAAGs. *J Topol.* **11** (2018), no. 1, 30–64.

[12] S. HERMILLER and J.MEIER. Algorithms and geometry for graph products of groups. *J. Algebra* **171** (1995), no. 1, 230–257.

[13] M. R. LAURENCE. A generating set for the automorphism group of a graph group. *J. London Math. Soc.* (2) **52** (1995), no. 2, 318–334.

[14] H. SERVATIUS. Automorphisms of graph groups. *J. Algebra* **126** (1989), no. 1, 34–60.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press

[15] K. VOGTMANN. *GL(n,Z), Out(F_n) and everything in between: automorphism groups of RAAGs.* London Mathematical Society Lecture Note Series 422. Groups St Andrews (2013), 105–127 (Cambridge University Press, 2015).

[16] K. VOGTMANN. *The topology and geometry of automorphism groups of free groups.* Proceedings of the 2016 European Congress of Mathematicians 181—202. (EMS Publishing House, Zurich, 2018).

[17] J. H. C. WHITEHEAD. On certain sets of elements in a free group. *Proc. Lond. Math. Soc.* **2** (1936), 48–56.

https://doi.org/10.1017/S0305004119000501 Published online by Cambridge University Press