$GL(n, \mathbb{Z})$, $Out(F_n)$ and everything in between: automorphism groups of RAAGs

KarenVogtmann

University of Warwick, Coventry, UK and Cornell University, Ithaca NY, USA
Email: kvogtmann@gmail.com

Abstract

A right-angled Artin group (RAAG) is a group given by a finite presentation in which the only relations are that some of the generators commute. Free groups and free abelian groups are the extreme examples of RAAGs. Their automorphism groups $GL(n, \mathbb{Z})$ and $Out(F_n)$ are complicated and fascinating groups which have been extensively studied. In these lectures I will explain how to use what we know about $GL(n, \mathbb{Z})$ and $Out(F_n)$ to study the structure of the (outer) automorphism group of a general RAAG. This will involve both inductive local-to-global methods and the construction of contractible spaces on which these automorphism groups act properly. For the automorphism group of a general RAAG the space we construct is a hybrid of the classical symmetric space on which $GL(n, \mathbb{Z})$ acts and Outer space with its action of $Out(F_n)$.

1 Introduction

In these lectures we will study the group of (outer) automorphism groups of a right-angled Artin group. Most of the material can be found in the papers [5, 6, 7, 8] which are all joint with Ruth Charney, some with additional authors. I will first go over some basic facts about right-angled Artin groups, then introduce inductive algebraic methods for studying these groups, then turn to more recent work on geometric methods. I will concentrate on describing and motivating the constructions but avoid proofs, however I will give explicit references to sources where the interested reader can find detailed proofs.

The Groups St. Andrews conference was run seamlessly by Colin Campbell, Edmund Robertson, Max Neuhoeffer, Colva Rodney-Dougal and Martin Quick, and I would like to thank them all warmly for inviting me to give these lectures.

2 Lecture 1

2.1 Definition of a RAAG

A right-angled Artin group, or RAAG for short, is a finitely-presented group whose relators (if any) are all simple commutators of generators. The extreme examples of RAAGs are free groups (with no relators), and free abelian groups (with all possible commutators of

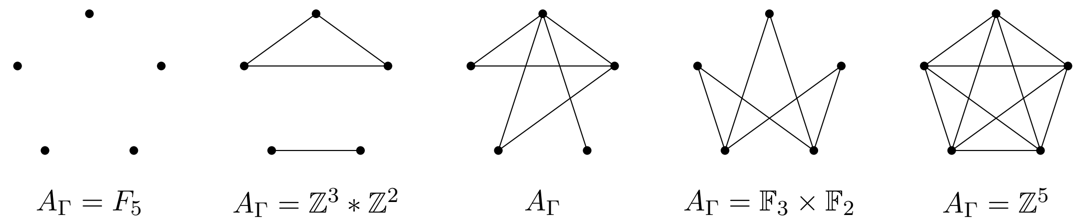

Figure 1. Graphs Γ and the associated RAAGs A_Γ

generators as relators). A RAAG is often specified by drawing a graph Γ with one vertex for each generator and one edge between two vertices if the corresponding generators commute (see Figure 1). Note that Γ is a simplicial complex, i.e. it has no loops or multiple edges.

If we start with any simplicial graph Γ the corresponding RAAG is denoted A_Γ. If Γ is disconnected with components $C_1, \ldots, C_k$, then A_Γ is the free product $A_{C_1} * \ldots * A_{C_k}$ and (just to maximize notational confusion) if Γ is a simplicial join $\Gamma = \Gamma_1 * \Gamma_2$, then A_Γ is the direct product $A_{\Gamma_1} \times A_{\Gamma_2}$.

2.2 Cell complexes with fundamental group A_Γ

Given a presentation $G = \langle X | R \rangle$ of a group there is a standard way of constructing a cell complex with fundamental group G, called the *presentation 2-complex*. This has one vertex, an edge for each generator $x \in X$, and a 2-cell for each relator $r \in R$. For a RAAG A_Γ, the 2-cells are all squares (see Figure 2).

The universal cover of the presentation 2-complex for A_Γ is not necessarily contractible, but it can be made contractible by attaching a few more cells. Recall that a *k-clique* in a graph is a complete subgraph with k vertices. If $\Delta \subset \Gamma$ is a k-clique, then the presentation 2-complex for A_Δ is a subcomplex of the presentation 2-complex for A_Γ, and is easily seen to be the 2-skeleton of a k-torus (constructed by gluing opposite sides of a k-dimensional cube). If we fill in this 2-skeleton with the entire k-torus for every k-clique in Γ, the result is called the *Salvetti complex* for A_Γ and is denoted S_Γ.

The Salvetti complex S_Γ is a cube complex which by construction satisfies Gromov's *link condition* and therefore supports a non-positively curved (*locally CAT(0)*) metric. In particular its universal cover is CAT(0) and therefore contractible. The Salvetti complex of a RAAG is in fact a particular kind of non-positively curved cube complex in which hyperplanes are well separated, called a *special cube complex* by Haglund and Wise [17]. We will say a little more about CAT(0) geometry and Gromov's link condition in Lecture 4, but for a thorough introduction to these concepts we refer to [3].

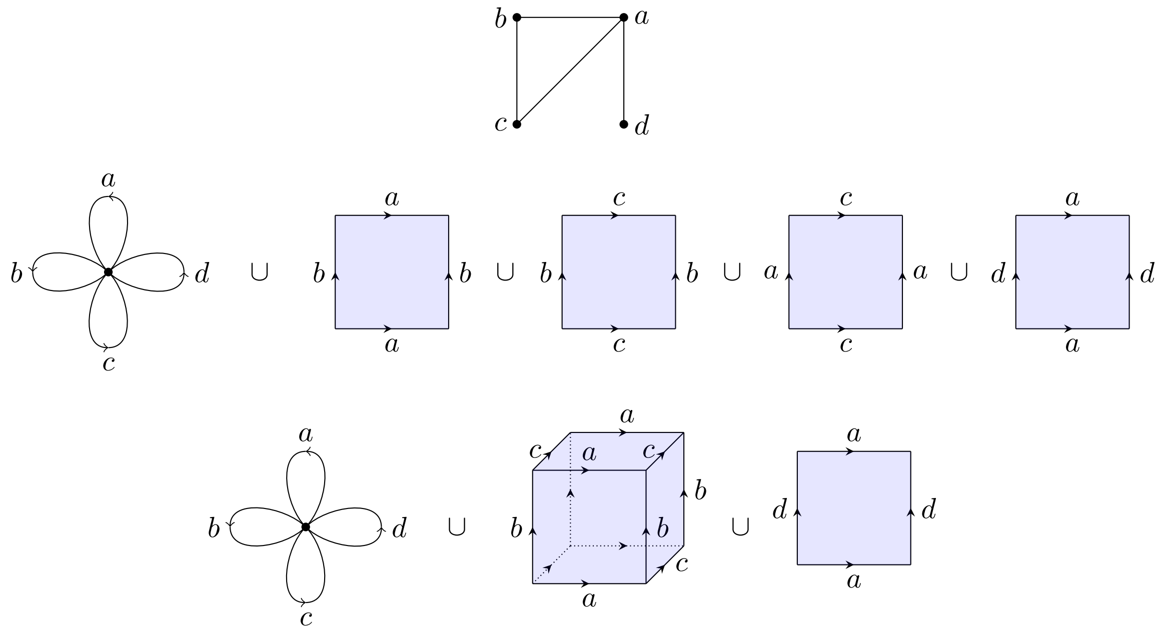

Figure 2. A graph, and kits for making its presentation 2-complex and its Salvetti complex

2.3 RAAGs and geometric group theory

RAAGs are important in geometric group theory for many reasons, including the fact that they have very interesting subgroups. They have been in the news lately because of Ian Agol's proof of Thurston's virtual fibering and virtual Haaken conjectures. A key step in those proofs is showing that the fundamental groups of closed hyperbolic 3-manifolds have finite-index subgroups which embed into RAAGs.

The extreme examples of RAAGs do not have such interesting subgroups. A subgroup of $\mathbb{Z}^n$ is a free abelian group of rank at most n. Things get slightly more interesting for F_n, where a subgroup is still a free group but can be of any rank, including infinity. Things got much more interesting when Stallings showed that the RAAG $F_2 \times F_2$ contains finitely generated subgroups which are not finitely presentable. In fact $F_2 \times F_2 \times \ldots \times F_2$ contains subgroups which are FP_{n-1} but not FP_n for all n, where FP_k is a k-dimensional algebraic finiteness property. This shows in particular that finitely-generated subgroups RAAGs are not necessarily RAAGs. Droms clarified the situation by characterizing exactly which RAAGs have the property that all of their finitely-generated subgroups are RAAGs: they are those for which the subgraph spanned by four vertices is never a square or a straight line [13]. Servatius, Droms and Servatius showed that if Γ is a pentagon, then A_Γ contains the fundamental group of a closed surface [26], and there is a great deal of recent work on surface subgroups of RAAGs by authors including S. Kim, T. Koberda, A. Duncan, I.

Kazachkov, M. Cassals-Ruiz, R. Weidman, I. Kapovich and A. Minasyan.

2.4 Automorphism groups of RAAGs

The emphasis of the present lectures is on automorphism groups of RAAGs. We will address the following three natural questions:

- How does the shape of Γ affect properties of $Out(A_\Gamma)$?
- $Aut(F_n), Out(F_n)$ and $GL(n,\mathbb{Z})$ share many basic properties. Which are in fact properties of $Out(A_\Gamma)$ for *any* Γ?
- How can we leverage information about $Out(F_n)$ and $GL(n,\mathbb{Z})$ to gain information about $Out(A_\Gamma)$?
- What techniques classically used to study $Out(F_n)$ and $GL(n,\mathbb{Z})$ can be adapted to $Out(A_\Gamma)$? We are especially interested in geometric techniques.

For the most part we will concentrate on the outer automorphism group $Out(A_\Gamma)$ instead of the full automorphism group $Aut(A_\Gamma)$. For $A_\Gamma = \mathbb{Z}^n$ there is no difference. For any Γ the abelianization map $A_\Gamma \to \mathbb{Z}^n$ induces a map on automorphism groups since the commutator subgroup is characteristic. But this map factors through $Out(A_\Gamma)$, and for $A_\Gamma = F_2$, the induced map on $Out(F_2)$ is an isomorphism, so *Out* seems the more natural comparison group. Further motivation is provided by the fact that we want to model automorphisms by maps on spaces with fundamental group A_Γ, and passing to $Out(A_\Gamma)$ means that we do not have to endow these spaces with basepoints and keep track of where the basepoint goes under the maps.

2.5 Generators for $Aut(A_\Gamma)$

For $A_\Gamma = \mathbb{Z}^n$ it is an easy consequence of the Euclidean algorithm that $Out(A_\Gamma) = GL(n,\mathbb{Z})$ is generated by the elementary matrices $A_{ij} = I_n + E_{ij}$ for $i \neq j \in \{1,\dots,n\}$ (where the only non-zero entry of E_{ij} is a 1 in the (i,j)-position) and the matrix

$$T = \begin{pmatrix} -1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ 0 & 0 & \dots & 1 \end{pmatrix}$$

If we let $GL(n,\mathbb{Z})$ act on $\mathbb{Z}^n$ on the right then A_{ij} sends $e_i \mapsto e_i + e_j$ and fixes e_k for $k \neq i$. This is called a *transvection*. In multiplicative notation for the free abelian group with generators $\{a_1,\dots,a_n\}$ these generators become $T\colon a_1 \mapsto a_1^{-1}$ and $A_{ij}\colon a_i \mapsto a_i a_j = a_j a_i$.

For $A_\Gamma = F_n$, the group $Out(A_\Gamma) = Out(F_n)$ is also generated by $T\colon a_1 \mapsto a_1^{-1}$ and by transvections, but right transvections $\rho_{ij}\colon a_i \mapsto a_i a_j$ are now different from left transvections $\lambda_{ij}\colon a_i \mapsto a_j a_i$ and the most natural presentation of $Out(F_n)$ (due to Gersten [14]) uses both. The fact that these generate $Out(F_n)$ was originally proved by Magnus [20], but the slickest proof is the one by Stallings using foldings of graphs [27].

For a general RAAG, not every transvection gives an automorphism: if a commutes with c but b does not, then the transvection $a \mapsto ab$ is not a homomorphism. This is the only thing that can go wrong, though: one just needs to check that everything that commutes with a also commutes with b; in this case we say the transvection $a \mapsto ab$ is Γ*-legal.* It is convenient to express this in terms of the defining graph Γ using the following standard terminology, which will be used throughout these lectures:

Definition 2.1 Let a be a vertex of Γ. The *link* of a is the full subgraph $lk(a)$ spanned by all vertices adjacent to a, and the *star* of a is the full subgraph $st(a)$ spanned by $lk(a)$ and a, i.e. $st(a)$ is the simplicial join $a * lk(a)$.

If Θ is a full subgraph of Γ, then the *link of* Θ is the intersection of the links of vertices in Θ

$$lk(\Theta) = \bigcap_{b\in\Theta} lk(b),$$

and the *star of* Θ is the simplicial join of Θ and $lk(\Theta)$

$$st(\Theta) = \Theta * lk(\Theta),$$

Twists and folds. Using the above notation, the condition for a transvection to be Γ-legal is: transitions $a \mapsto ab$ and $a \mapsto ba$ are Γ-legal if and only if

- $ab \neq ba$ and $lk(a) \subseteq lk(b)$, or
- $ab = ba$ and $st(a) \subseteq st(b)$.

This can be said more economically by the single condition $lk(a) \subseteq st(b)$, but it is often important to retain the distinction (commuting versus non-commuting is a critical difference here!) so we also introduce different terminology for the two types of transvections

Definition 2.2 If $ab = ba$, then a Γ-legal transvection $a \mapsto ab$ is a *twist.* If $ab \neq ba$ then a Γ-legal transvections $a \mapsto ab$ and $a \mapsto ba$ are called *(right and left) folds.*

The reason for this terminology will become clear when we discuss geometric models for these automorphisms.

Partial conjugations. Even if we can't transvect b onto a we can still try to conjugate a by b. If we do that, we must also conjugate everything which commutes with a, and everything that commutes with things that commute with a, etc. However, the vertices in $lk(b)$ don't know whether they've been conjugated by b or not, so if a and a' are separated by $lk(b)$, we could conjugate a by b but not a'. In other words, conjugating an entire component of $\Gamma - lk(b)$ by b gives an automorphism; this is called a (Γ-*legal*) *partial conjugation.*

Inversions and graph automorphisms. It is clear that a permutation of the generators will be an automorphism if and only if it extends to an automorphism of Γ, since Γ encodes the commuting relations. Since $Aut(A_\Gamma)$ does not contain all transvections, we can't assume that all of these permutations are products of transvections, so we include graph automorphisms in the generating set. Similarly, we add all inversions $a_i \mapsto a_i^{-1}$ instead of just $a_1 \mapsto a_1^{-1}$.

The types of automorphisms described in the last three paragraphs now do generate $Aut(A_\Gamma)$, by a theorem of Laurence and Servatius.

Theorem 2.3 *[19, 25] $Aut(A_\Gamma)$ is generated by graph automorphisms, inversions, and Γ-legal twists, folds and partial conjugations.*

In particular this shows that $Aut(A_\Gamma)$ (and therefore $Out(A_\Gamma)$) is finitely generated. Both $Aut(A_\Gamma)$ and $Out(A_\Gamma)$ are also finitely presented. This was proved in some special cases in [2], and an explicit finite presentation in all cases was given by Matt Day [11]. Day's proof closely follows McCool's proof that $Aut(F_n)$ is finitely presented using a "Peak reduction" algorithm (see [21]).

3 Lecture 2

In the last lecture we introduced right-angled Artin groups and their automorphisms. In this lecture we will show how to infer information about general $Out(A_\Gamma)$ from known facts about $Out(F_n)$ and $GL(n, \mathbb{Z})$.

We already mentioned that $Out(A_\Gamma)$ is finitely generated and finitely presented for all A_Γ. We claim that it also has higher-dimensional homological finiteness properties:

- $Out(A_\Gamma)$ has subgroups of finite index which are torsion-free, i.e. $Out(A_\Gamma)$ is *virtually torsion-free.*
- In fact, $Out(A_\Gamma)$ has lots of torsion-free finite index subgroups: for any given element $\phi \in Out(A_\Gamma)$ there is a torsion-free finite index which does not contain ϕ, i.e. $Out(A_\Gamma)$ is *residually finite.*
- The homology of any torsion-free finite index subgroup is finitely generated. In particular its homology vanishes above some point, i.e. $Out(A_\Gamma)$ has *finite virtual cohomological dimension.*

We will introduce our method for bootstrapping information about $Out(F_n)$ and $GL(n, \mathbb{Z})$ to $Out(A_\Gamma)$ by giving a proof that $Out(A_\Gamma)$ contains a torsion-free subgroup of finite index.

The proof that $GL(n, \mathbb{Z})$ is virtually torsion-free is quite easy, one just checks that the kernel of the natural map $GL(n, \mathbb{Z}) \to GL(n, \mathbb{Z}/3)$ has no torsion; the proof uses only the binomial theorem.

The proof that $Out(F_n)$ is virtually torsion-free relies on this calculation plus the non-trivial fact, due to Baumslag and Taylor [1] that the kernel of the natural map $Out(F_n) \to GL(n, \mathbb{Z})$ is torsion-free.

The kernel of the map $Out(A_\Gamma) \to GL(n, \mathbb{Z})$ for general A_Γ is also torsion-free; this is one consequence of a recent paper by Toinet [29]. We will avoid appealing to this, however, since our point is to illustrate the general bootstrapping method.

Since we are interested in a virtual notion, it suffices to pass to a subgroup of finite index. Let $Aut^0(A_\Gamma)$ denote the subgroup of $Aut(A_\Gamma)$ generated by inversions, transvections and partial conjugations (we are leaving out only the graph automorphisms), and let $Out^0(A_\Gamma)$ be the image of $Aut^0(A_\Gamma)$ in $Out(A_\Gamma)$.

Exercise 3.1 Show that $Aut^0(A_\Gamma)$ and $Out^0(A_\Gamma)$ are normal subgroups of $Aut(A_\Gamma)$ and $Out(A_\Gamma)$ respectively, and then that they have finite index.

To show that $Out^0(A_\Gamma)$ has a torsion-free subgroup of finite index the key idea we will exploit is that there are lots of subgroups in A_Γ which must be sent to conjugates of themselves by any automorphism. To describe these subgroups, we introduce some basic facts and some new terminology.

Lemma 3.2 *Let V be a set of vertices in Γ and Θ the full subgraph spanned by V. Then the subgroup generated by V is isomorphic to A_Θ.*

Such a subgroup is called a *special subgroup*. By convention, we set $A_\emptyset = 1$.

Now recall that a transvection $a \mapsto ab$ (or $a \mapsto ba$) is an automorphism of A_Γ if and only if $st(a) \subseteq lk(b)$. In this case we write $a \preceq b$. If $a \preceq b$ and $b \preceq a$ we say $a \sim b$; this defines an equivalence relation on the vertices of Γ. The notation is justified by the following observation.

Exercise 3.3 The set of equivalence classes of vertices of Γ is a partially ordered set, with partial order induced by $\preceq$.

A vertex is called *maximal* if its equivalence class is maximal in this partial order.

Let $[a]$ denote the full subgraph of Γ spanned by vertices equivalent to a.

Proposition 3.4 ([6], Proposition 3.2) *Let Γ be a connected graph and a a maximal vertex in Γ. Then any $\phi \in Out^0(A_\Gamma)$ is represented by some $f_a \in Aut^0(A_\Gamma)$ with*

$$f_a(A_{st(a)}) = A_{st(a)} \text{ and } f_a(A_{[a]}) = A_{[a]}.$$

Proof The proof is accomplished by checking that the statement is true for all of the generators of $Out^0(A_\Gamma)$. □

Proposition 3.5 ([7], Section 3) *For Γ connected and a maximal in Γ there are homomorphisms*

- **[Restriction]** $R_a \colon Out^0(A_\Gamma) \to Out^0(A_{st[a]})$
- **[Exclusion]** $E_a \colon Out^0(A_\Gamma) \to Out^0(A_{\Gamma - [a]})$
- **[Projection]** $P_a \colon Out^0(A_\Gamma) \to Out^0(A_{lk[a]})$

Sketch of proof. If f_a is the map representing $\phi \in Out^0(A_\Gamma)$ described in Proposition 3.4, then the restriction map sends ϕ to the class of the restriction of f_a to $A_{st[a]}$. This is well-defined because $A_{st[a]}$ is its own normalizer (see, e.g., [7], Proposition 2.2).

Exclusion is induced by the map $A_\Gamma \to A_{\Gamma-[a]}$ sending $v \mapsto 1$ if $v \in [a]$ and $v \mapsto v$ if $v \notin [a]$. This is well-defined because the normal subgroup generated by a maximal equivalence class $[a]$ is charactersitic, by Proposition 3.4.

Projection is the composition $P_a = E_a \circ R_a$. This is well-defined because $[a]$ is maximal in $st[a]$, so E_a is defined on the image of R_a. □

We can put all of the projection homomorphisms together to get a single homomorphism $P = \prod P_a$. The following theorem is the basic result which enables our bootstrapping technique.

Theorem 3.6 *[[6], Theorem 4.1 and [7], Section 3] Let Γ be connected, and set*

$$P = \prod_{[a]\, maximal} P_a \colon Out^0(A_\Gamma) \to \prod_{[a]\, maximal} Out^0(A_{lk[a]}).$$

Then $ker(P)$ is finitely generated and free abelian. The rank of $ker(P)$ is computable in terms of Γ.

Notice that $lk[a]$ is smaller than Γ. We would like to use this fact to do induction. There is, however a problem: all of the above results have the hypothesis that Γ is connected, but $lk[a]$ need not be connected, even if Γ is. There are two ways to get around this. First, since disconnected graphs give rise to free products of RAAGs we can sometimes take advantage of known results about free products. If these are not available, we can simply assume that all non-empty links are either connected or discrete (reducing us by induction to the general linear and free group cases).

We illustrate the first option by proving that $Out(A_\Gamma)$ has torsion-free subgroups of finite index. We take advantage of the following theorem of Guirardel and Levitt:

Theorem 3.7 ([15]) *Let $G = G_1 * G_2$ with G_i and $G_i/Z(G_i)$ torsion-free. If $Out(G_i)$ is virtually torsion-free for $i = 1, 2$ then so is $Out(G)$.*

With this in our repertoire we can now prove our theorem.

Theorem 3.8 ([6], Theorem 5.2) *$Out(A_\Gamma)$ has torsion-free subgroups of finite index, for any Γ.*

Proof If Γ is a disjoint union $\Gamma = \Gamma_1 \sqcup \Gamma_2$, then $A_\Gamma = A_{\Gamma_1} * A_{\Gamma_2}$, so by Theorem 3.7 it suffices to consider connected graphs Γ.

If Γ is a complete graph then $A_\Gamma = \mathbb{Z}^n$ and the theorem is true, as noted above. If Γ is not complete, then $lk[a]$ is non-empty for some maximal vertex a. For any vertex a of Γ

the maximal size of a clique in $lk[a]$ is strictly less than the maximal size of a clique in Γ. Therefore we can use this number, which we denote $m(\Gamma)$, to do induction.

If $m(\Gamma) = 1$ then Γ is discrete and $A_\Gamma = F_n$, in which case the theorem is true by the theorem of Baumslag and Taylor.

If $m(\Gamma) = 2$ then $lk[a]$ is discrete for all a, so we can use the map P defined in Theorem 3.6 to pull back a product of torsion-free finite index subgroups of $Out^0(A_{lk[a]})$ to obtain a torsion-free finite index subgroup of $Out(A_\Gamma)$.

Now induction on $m(\Gamma)$ together with Theorem 3.7 completes the proof. □

The groups A_Γ are residually finite; this follows from the fact that they are linear, which was proved by Davis and Januskiewicz [10]. Residual finiteness for $Aut(A_\Gamma)$ then follows from Baumslag's theorem that the automorphism group of any residually finite group is itself residually finite. Residual finiteness for $Out(A_\Gamma)$ is more subtle, but using the homomorphisms P, R and E above together with various inductive schemes we can also settle this question.

Theorem 3.9 ([7], Theorem 4.2) *For any RAAG A_Γ, $Out(A_\Gamma)$ is residually finite.*

This was also proved by Minasyan [22]. Both proofs rely on the result of Minasyan and Osin that if G_1 and G_2 are finitely generated groups with $Out(G_1)$ and $Out(G_2)$ residually finite, then $Out(G_1 * G_2)$ is residually finite [23].

Another result which can be proved using the maps P, R and E is

Theorem 3.10 ([6], Theorem 5.2) *$Out(A_\Gamma)$ has finite virtual cohomological dimension.*

Here again we rely on a result of Guirardel and Levitt about free products, namely

Theorem 3.11 ([15]) *Let $G = G_1 * G_2$ with G_i and $G_i/Z(G_i)$ torsion-free. If $Out(G_i)$ has finite virtual cohomological dimension for $i = 1, 2$ then so does $Out(G)$.*

If we do not have a suitable result in the wings for free products, we need to hypothesize that Γ is connected and the link of every non-maximal clique is either connected or discrete; such a graph Γ is called *homogeneous*. This is automatically true if Γ has no triangles (in which case the Salvetti complex is 2-dimensional, so we say A_Γ is two-dimensional). It is also true, e.g., if Γ is the 1-skeleton of a triangulated manifold. As an example, we can prove

Theorem 3.12 ([7], Theorem 5.5) *If Γ is homogeneous then every subgroup of $Out(A_\Gamma)$ is either virtually solvable or contains a free group of rank 2.*

We can also bound the maximum derived length of a solvable subgroup in terms of the shape of Γ (see [7], Section 6). The crudest such estimate is that this length is always less than or equal to the number of vertices in Γ.

4 Lecture 3

In the last lecture we studied $Out(A_\Gamma)$ via the projection map

$$P = \prod_{[a]\, maximal} P_a \colon Out^0(A_\Gamma) \to \prod_{[a]\, maximal} Out^0(A_{lk[a]})$$

and its free abelian kernel.

We remarked that we can use this to bound the virtual cohomological dimension of $Out(A_\Gamma)$. However, P is far from surjective, and $ker(P)$ is far from being maximal rank among abelian subgroups, so the upper and lower bounds this gives are not very good.

In this lecture we take a more geometric approach to the study of $Out(A_\Gamma)$ by attempting to realize $Out(A_\Gamma)$ as symmetries of an "outer space." As before the classical theory of $GL(n, \mathbb{Z})$ and $Out(F_n)$ provide guidance.

$GL(n, \mathbb{Z})$ acts on the symmetric space $SO(n)\backslash SL(n, \mathbb{R})$, and $Out(F_n)$ acts on Outer space. Useful features of these actions include:

- the spaces are contractible
- the actions are proper

These two properties imply that algebraic invariants of the groups can be computed by computing topological invariants of the quotient spaces; in particular the cohomology $H^*(\Gamma) \cong H^*(X/\Gamma)$. Further properties of the classical actions include

- the spaces are finite-dimensional

from which we can immediately conclude that the virtual cohomological dimension of the groups are finite, and

- there is a cocompact equivariant deformation retract

which implies that the group cohomology is finitely generated in all dimensions. Furthermore, the quotient of the retract by the action can be described combinatorially, making it possible to do explicit cohomology calculations, at least in small dimensions.

More sophisticated features include

- the spaces have *bordifications*, i.e. they can be enlarged to spaces with proper cocompact actions, whose cohomology at infinity is concentrated in one dimension.

By work of Bieri and Eckmann this implies that the groups are *virtual duality groups*, i.e. there is a dualizing module D and isomorphisms

$$H^*(G; A) \cong H_{d-*}(G, D \otimes A)$$

between cohomology with any coefficients A and homology with coefficients in $D \otimes A$.

Outer space for a general RAAG will be a hybrid species, combining features of both symmetric spaces and Outer space. So let us now review these spaces.

4.1 Symmetric space

The symmetric space $\mathbb{D}_n = SO(n)\backslash SL(n, \mathbb{R})$ has several useful alternate descriptions. A coset $SO(n)A$ gives a well-defined positive definite symmetric matrix $Q = A^tA$, identifying $\mathbb{D}_n$ with

- the space of positive definite quadratic forms Q on $\mathbb{R}^n$, modulo homothety.

If we fix the standard lattice $\mathbb{Z}^n \subset \mathbb{R}^n$, then any linear map $A\colon \mathbb{Z}^n \to \mathbb{R}^n$ defines a *marked lattice*, modulo homothety. If we also mod out by rotations, then $\mathbb{D}_n$ can also be described as

- the space of marked lattices $A\colon \mathbb{Z}^n \to \mathbb{R}^n$ modulo homothety and rotation.

Finally, the map $A\colon \mathbb{R}^n \to \mathbb{R}^n$ induces a map $\bar{A}\colon \mathbb{R}^n/\mathbb{Z}^n \to \mathbb{R}^n/A(\mathbb{Z}^n)$. We think of $\mathbb{R}^n/\mathbb{Z}^n$ as a standard torus T^n, $Y = \mathbb{R}^n/A(\mathbb{Z}^n)$ as a torus with a flat metric, and $\bar{A}$ as an isotopy class of homeomorphisms; then $\mathbb{D}_n$ is identified with

- the space of marked flat tori $\bar{A}\colon T^n \to Y$. modulo homothety.

In each case, the group $GL(n, \mathbb{Z})$ acts on the right. If $g \in GL(n, \mathbb{Z})$, then

- $SO(n)A \cdot g = SO(n)Ag$
- $Q \cdot g = g^tQg$
- $(A\colon \mathbb{Z}^n \to \mathbb{R}^n) \cdot g = Ag\colon \mathbb{Z}^n \to \mathbb{R}^n$
- $(\bar{A}\colon T^n \to Y) \cdot g = \overline{Ag}\colon T^n \to Y$.

Each of these descriptions of the symmetric space has its advantages. For example, the description as the space of positive definite quadratic forms makes it easy to see that $\mathbb{D}_n$ is contractible, since the set of positive definite quadratic forms a convex cone in the space of $n \times n$ matrices. The description that will be most relevant for us is the last, as a space of marked flat tori. Note that the action of $GL(n, \mathbb{Z})$ changes the marking, but does *not* change the flat metric.

4.2 Outer space

Outer space also has several useful descriptions. We can mimic the description of the symmetric space as the space of marked flat tori by defining Outer space as a space of *marked metric graphs.* To do this, we fix a *rose* R_n, i.e. a graph with one vertex and n directed edges, as a "model space" to play the role of the torus T^n. A metric on a graph X is simply an assignment of positive real lengths to its edges, making X into a metric space with the path metric. A *marking* is a homotopy equivalence $h\colon R_n \to X$. For technical reasons we don't allow our graphs to have univalent or bivalent vertices, and they must be finite. Marked graphs (X, h) and (X', h') are *equivalent* if there is an isometry or a homothety $f\colon X \to X'$ with $f \circ h$ homotopic to h'.

Definition 4.1 *Outer space* CV_n is the space of equivalence classes of marked metric graphs with fundamental group F_n

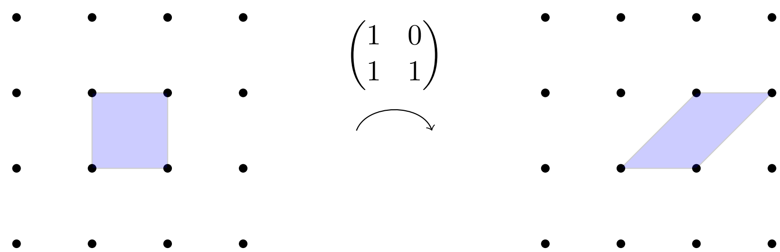

Figure 3. Action of A_{21} on Λ

$Out(F_n)$ acts on CV_n on the right by changing the marking, i.e. for $\phi \in Out(F_n)$ take a map $f\colon R_n \to R_n$ that induces ϕ on $\pi_1(R_n) \equiv F_n$ and set $(X,h)\cdot\phi = (X, h\circ f)$.

There is an obvious equivariant deformation retraction of CV_n onto the subspace consisting of marked metric graphs with no separating edges (simply shrink each separating edge to a point). This subspace, called *reduced Outer space* is sometimes more convenient to work with.

Given a marked metric graph X, the marking serves to identify F_n with the fundamental group of X. By looking at the universal cover $\widetilde{X}$ we thus obtain a simplicial tree with a free action of F_n. The fact that we don't allow X to be infinite or have univalent or bivalent vertices translates into the condition that the action is *minimal*, i.e. there are no F_n-invariant subtrees. Therefore an alternate description of CV_n is as the space of *free minimal actions of F_n on metric simplicial trees.* This is analogous to the description of the symmetric space as a space of lattices instead of as a space of flat tori. (There is a third definition of CV_n in terms of isotopy classes of spheres in a doubled handlebody which is extremely useful, but will not be relevant for these lectures.)

4.3 Lattices, tori and graphs in rank 2

To motivate our definition of outer space for a general RAAG, we first compare the symmetric space and Outer space in rank 2. In rank 2 the natural map $Out(F_2) \to GL(2,\mathbb{Z})$ induced by abelianization $F_2 \to \mathbb{Z}^2$ is an isomorphism and both (reduced) Outer space and the symmetric space can be identified with the hyperbolic plane. The spaces diverge dramatically in higher ranks, but the rank 2 picture gives us some insight into the general situation because we can look at the same space from two different points of view.

$GL(n,\mathbb{Z})$ is generated by elementary matrices, so consider the action of $A_{21} = \begin{pmatrix}1&0\\1&1\end{pmatrix}$ on the standard lattice $\Lambda = \mathbb{Z}^2 \subset \mathbb{R}^2$, marked by the identity. This sends $e_1 \mapsto e_1$ and $e_2 \mapsto e_1 + e_2$ (remember we are acting on the *right*). This action is illustrated in Figure 3

Note that the action of A_{21} does not change the lattice, it just changes the marking. Thus the orbit of $GL(n,\mathbb{Z})$ is a discrete subset of the space of all marked lattices. To get a

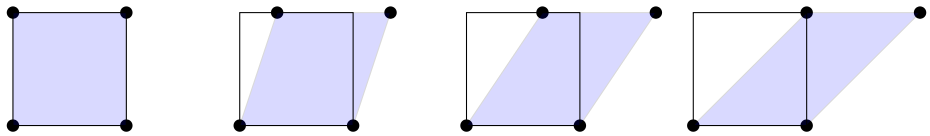

Figure 4. A path from Λ to $\Lambda \cdot A_{21}$.

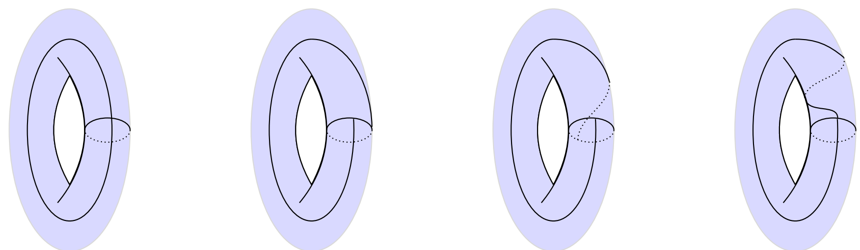

Figure 5. The same path, as a path of tori

path from Λ to $\Lambda \cdot A_{21}$ we must gradually shear the marked lattice, as in Figure 4. In the figure we have drawn the original fundamental domain for reference.

The path in the space of marked flat tori is obtained by identifying opposite sides of the fundamental domain for the lattices, as in Figure 5. In this figure, too, we have marked the original fundamental domain for reference.

This gives us a clue for what this path looks like if we think of it as a path in CV_2: if we puncture the torus its fundamental group is F_2 instead of $\mathbb{Z}^2$, and it deformation retracts onto the black graph. The path in CV_2 is then given by the graphs in Figure 6. Note that the total length of the graph is constant (equal to 2) in this picture. Under the deformation retraction of the punctured torus onto the graph, the action of A_{21} on $\mathbb{Z}^2$ becomes the action of $\rho_{21} \colon x_2 \mapsto x_2 x_1$ on F_2. We indicate this by coloring the loop representing x_2 in red.

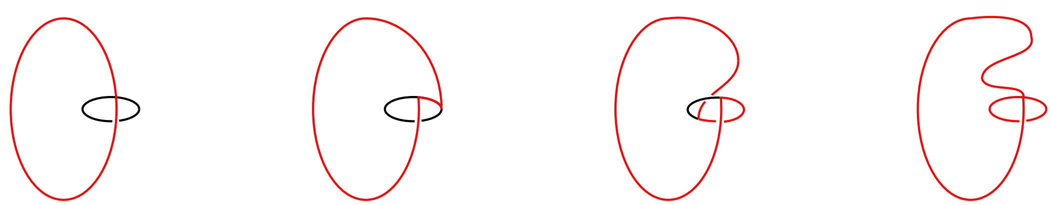

Figure 6. The same path, as a path in Outer space CV_2

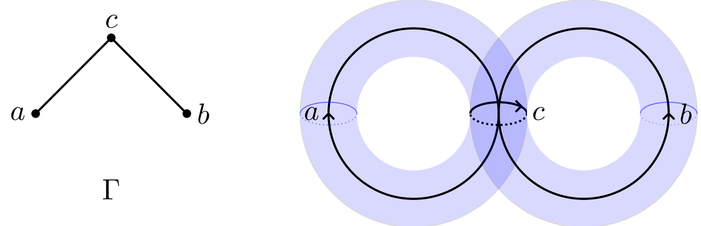

Figure 7. A simple graph Γ and its Salvetti complex S_Γ

4.4 A simple example

We have described CV_n as a space of marked metric graphs $R_n \to X$ and the symmetric space for $\mathbb{Z}^n$ as a space of marked flat tori $T^n \to X$. In each case we needed a model space and a homotopy equivalence to a metric space. For a general *RAAG* A_Γ we have a model space, namely the Salvetti complex S_Γ, so we would like to have an outer space of marked metric spaces $S_\Gamma \to X$. We now need to decide:

- What homeomorphism types X should we allow?
- What metric structures should we allow on these spaces X?

We begin by looking at a very simple RAAG, i.e.

$$\langle a, b, c | [a, b] = [b, c] = 1 \rangle.$$

This is the RAAG associated to the graph Γ with three vertices a, b and c and two edges, one from a to c and one from c to b. The Salvetti complex S_Γ is the union of two tori, glued along a common meridian curve labeled c (see Figure 7).

Generators for $Out(A_\Gamma)$ consist of the graph automorphism interchanging a and b, inversion in a, inversion in c, the twist $\tau_{ac} \colon a \mapsto ac = ca$, and the folds $\rho_{ab} \colon a \mapsto ab$ and $\lambda_{ab} \colon a \mapsto ba$.

The group A_Γ is the product of the cyclic group generated by c with the free group generated by a and b, and the Salvetti complex S_Γ is the product of the loop labeled c by the rose formed by the two longitudinal curves. To realize the twist τ_{ac} on S_Γ, we can perform a Dehn twist of the left-hand torus around a curve parallel to c but disjoint from c (see Figure 8).

To realize the transvection ρ_{ab} we *fold* the left-hand torus around the right-hand torus (see Figure 9). This is can also be described as first expanding the intersection circle into a cylinder, then collapsing a different cylinder (the bottom blue cylinder in the figure).

Thus to make a path from S_Γ to $S_\Gamma \cdot \tau_{ac}$ we need to gradually shear the metric on the left-hand torus, and to make a path from S_Γ to $S_\Gamma \cdot \rho_{ab}$ we need to pass through spaces X which are not homeomorphic to S_Γ. In our outer space for A_Γ we need to be able to vary both the homeomorphism type of spaces and the metrics on the spaces. But we want to

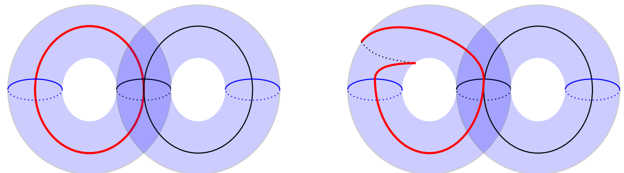

Figure 8. Realizing the twist τ_{ac} on S_Γ

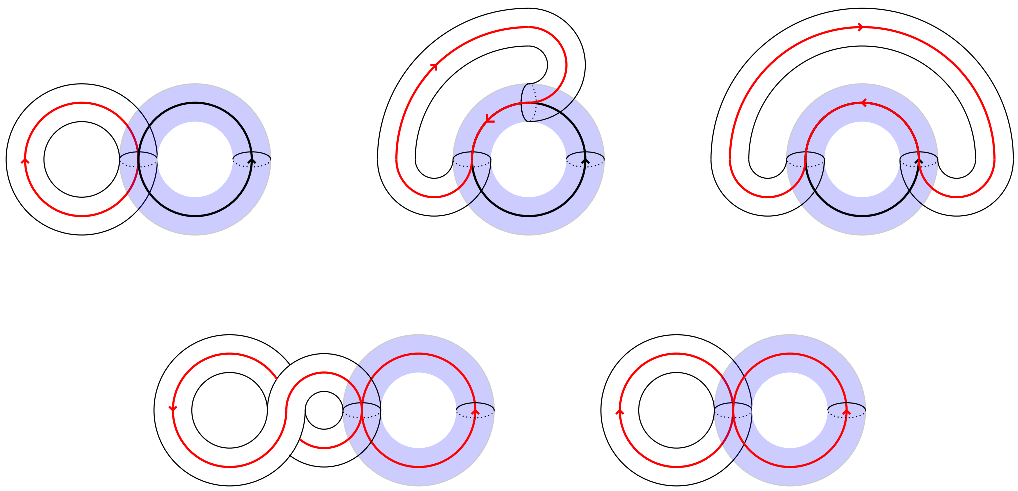

Figure 9. Realizing the fold ρ_{ab} on S_Γ

restrict both as much as possible so that we can control the topology and geometry of the space.

Note that both S_Γ and the intermediate complexes X are combinatorially cube complexes. With standard Euclidean metrics on the cubes they are non-positively curved cube complexes (i.e. their universal covers are CAT(0)), in fact they are *special* cube complexes, in the sense of Haglund and Wise [17]. However, we want to vary the (projective classes of) the metrics to allow shearing of the tori. This can be accomplished in the intermediate spaces X by giving all three cylinders flat right-angled metrics with the same circumference, then specifying the attaching maps to the two circles by shear parameters.

There are a priori two shear parameters for each cylinder, but shearing both ends of a cylinder by the same amount does not change the metric on X, so there are actually only three total parameters. These are not independent either, since shearing all three by the same amount simply twists the circle without affecting the metric; thus in the end we have only 2 independent shear parameters.

Shearing the top or bottom cylinder by an entire rotation changes the marking by a twist, and shearing in the opposite direction changes the marking by its inverse. Expanding and collapsing cylinders without shearing varies the space independently of the c direction, so may be thought of as moving around the space of metric graphs marked by the free subgroup generated by a and b, i.e. around reduced Outer space in rank 2. Since reduced Outer space in rank 2 is homeomorphic to $\mathbb{R}^2$, the entire moduli space of marked metric blowups is homeomorphic to the product $\mathbb{R}^2 \times \mathbb{R} \times \mathbb{R}$.

5 Lecture 4

In this lecture we show how to construct an outer space for any RAAG A_Γ; this will be a space of marked metric cube complexes. We then outline very briefly how to prove the space we have constructed is contractible and that the action is proper.

We recall the example we studied in the last lecture, where Γ has three vertices and two edges. To get a path from the standard Salvetti $id\colon S_\Gamma \to S_\Gamma$ to its image under the twist $\tau_{ac}\colon a \mapsto ac = ca$ we needed to shear the metric on the left-hand torus, while to get a path to its image under the fold $\rho_{ab}\colon a \mapsto ab$ we needed to expand a circle of the Salvetti into a cylinder, then collapse a different cylinder. The first operation involves changing the metrics on the spaces without changing their homeomorphism type, while the second operation can be described combinatorially in terms of "blowing up" and collapsing subcomplexes.

We begin our construction by determining which cube complexes we need, temporarily ignoring their metric structure. For $A_\Gamma = F_n$, this amounts to describing the (vertices of the) *spine* K_n instead of the full space CV_n, so we briefly recall that construction.

Outer space CV_n for a free group is the union of open simplices, one for each equivalence class of marked (combinatorial) graphs (X, h), where X is a finite graph with no bivalent vertices or separating edges (and hence no univalent vertices, either). The open simplex associated to (X, h) is obtained by assigning all possible positive real lengths to the edges of X, then either projectivising or (equivalently) normalizing so that the sum of the lengths is one. If we take small neighborhoods of these simplices we obtain an open cover of CV_n by contractible sets, such that the intersection of any two elements is either empty or is in the cover. The nerve of this cover is known as the *spine of Outer space*, and is an equivariant deformation retract of all of CV_n. The spine can be described combinatorially as the geometric realization of the partially ordered set of marked graphs, where the poset relation is given by forest collapse: $(X, h) > (X', h')$ if there is a forest $\Phi \subset X$ such that X' is obtained from X by collapsing each edge of Φ to a point, and h' is (homotopic to) the composition of h with this collapse. The full space CV_n can be recovered from the spine by putting the metric information back into the graphs.

Motivated by this, we will now construct a similar spine for any A_Γ.

5.1 The spine of outer space for A_Γ

For general A_Γ we need analogs of graphs, forests and forest collapses. The analog of a graph will be a particular type of non-positively curved cube complex adapted to Γ, which we call a Γ-complex. We first recall some standard background about cube complexes.

5.1.1 NPC cube complexes and hyperplanes

A *cube complex* is a CW complex X in which every cell is homeomorphic to a Euclidean cube (of some dimension) and the attaching maps identify faces with lower-dimensional cells by homeomorphisms.

If v is a vertex of a cube complex X there is an associated simplicial complex called the *link* of v and denoted $lk(v)$. This has one vertex for each half-edge terminating at v, and a set of half-edges spans a k-simplex if they belong to distinct edges of the same cube. Gromov gave a simple condition on links which guarantees that X can be given a metric of non-positive curvature. This says that if the 1-skeleton of a simplex appears in $lk(v)$, then the entire simplex must be in $lk(v)$. This is called the *flag condition* on links, and a cube complex X whose links satisfy the flag condition is said to be *NPC*.

Cubes in a cube complex are cut by *hyperplanes*. A hyperplane is dual to an equivalence class of edges, where the equivalence relation is generated by saying two edges are equivalent if they are parallel in some cube. If H is the hyperplane dual to $[e]$, then the intersection of H with a cube C is spanned by midpoints of edges of C which are in $[e]$; thus $H \cap C$ is either empty or is a codimension one linear subspace cutting C in half.

Example 5.1 The space X from the last lecture is an NPC cube complex. There are four hyperplanes. Three of them are circles midway up the three cylinders, and the other is a theta graph, with one edge running the length of each cylinder.

Example 5.2 S_Γ is an NPC cube complex with one k-cube for each k-clique in Γ. There is one hyperplane for each generator a of A_Γ (i.e. each vertex of Γ), and the associated hyperplane is isomorphic to the Salvetti complex $S_{lk[a]}$.

The *carrier* of a hyperplane H is the closure of the union of all cubes which intersect H. If the carrier of H is an embedded copy of $H \times [0,1]$, then collapsing each cube in the carrier to its intersection with H is called a *hyperplane collapse*, though maybe it should be called a carrier collapse. A hyperplane collapse is *trivial* if the resulting complex is still homeomorphic to X.

A set of hyperplanes $\{H_1, \dots, H_k\}$ in X is called a *hyperplane forest* if any cycle formed by edges dual to the H_i is null-homotopic. In this case each H_i determines a hyperplane collapse in which the images of the remaining H_j form a new hyperplane forest.

5.1.2 Marked Γ-complexes

Definition 5.3 A compact NPC cube complex X is called a Γ-*complex* if there is a hyperplane forest $\{H_1, \dots, H_k\}$ in X such that performing the associated hyperplane collapses (in any order) gives a cube complex isomorphic to S_Γ and

1. The hyperplane collapse associated to H_i is non-trivial for each i.
2. After each collapse, the image of any hyperplane in X is either a hyperplane or a subcomplex parallel to a hyperplane.

A *marking* of a Γ-complex X is a homotopy equivalence $h\colon S_\Gamma \to X$. Two marked Γ-complexes (X, h) and (X', h') are *equivalent* if there is an isomorphism of cube complexes $f : X \to X'$ with $h' \simeq f \circ h$.

The Salvetti complex S_Γ is of course an example of a Γ-complex. If we mark it with the identity map $id\colon S_\Gamma \to S_\Gamma$, the result is called the *standard Salvetti*.

The group $Out(A_\Gamma)$ acts on the set of marked Γ-complexes by changing the marking: any $\phi \in Out(A_\Gamma)$ can be induced by a homotopy equivalence $f\colon S_\Gamma \to S_\Gamma$, and we define $(X, h)\phi = (X, h \circ f)$.

The set of equivalence classes of marked Γ-complexes forms a partially ordered set under the relation of hyperplane collapse, and we define $\mathcal{M}_\Gamma$ to be the geometric realization of this poset.

5.1.3 The untwisted subgroup

We expect to have to shear the metric to find a path from the standard Salvetti to its image under a twist, as in the example from the last lecture. Since blowups, collapses and isometries don't do any shearing, we shouldn't even expect $\mathcal{M}_\Gamma$ to be connected, much less contractible. But it turns out that we can find a large contractible piece of $\mathcal{M}_\Gamma$ by ignoring the twists at first, i.e. we consider an orbit of the subgroup of $Out(A_\Gamma)$ generated by all other types of generators.

Definition 5.4 The *untwisted subgroup* $U(A_\Gamma)$ of $Out(A_\Gamma)$ is the subgroup generated by inversions, graph automorphisms, partial conjugations and folds.

The untwisted subgroup can be all of $Out(A_\Gamma)$ (e.g. if Γ is discrete or is without triangles and univalent vertices) or it can be finite (e.g. if Γ complete or is an n-gon with $n \geq 5$), and is generally somewhere in between. It is not usually normal (e.g. conjugating an inversion $v \mapsto v^{-1}$ by a twist $u \mapsto uv$ results in the square of the twist times the inversion).

Definition 5.5 Let $\mathcal{M}_\Gamma$ be the geometric realization of the poset of equivalence classes of marked Γ-complexes. Let $\sigma_0 = (S_\Gamma, id)$ be the standard Salvetti and $st(\sigma_0)$ the star of σ_0 in $\mathcal{M}_\Gamma$. The *spine of outer space* K_Γ is the orbit of $st(\sigma_0)$ under $U(A_\Gamma)$.

The following theorem is joint with Charney and Stambaugh.

Theorem 5.6 *[[8]] The spine K_Γ is contractible. The untwisted subgroup $U(A_\Gamma)$ acts cocompactly with finite stabilizers on K_Γ.*

The following subsections give a brief indication of the proof.

5.1.4 Blowups and Γ-Whitehead partitions

In order to prove Theorem 5.6 we need to understand exactly which complexes occur in the star of σ_0, i.e. which marked Γ-complexes collapse to (S_Γ, id). The opposite of a hyperplane collapse is called a *blowup*, and to explain how to find all blowups of σ_0 we first recall the situation for $A_\Gamma = F_n$.

If we can obtain a rose R_n by collapsing a single edge e of a graph G, then the other edges of G can be identified with the petals a_i of R_n. Each petal is an edge with two ends, a_i^+ and a_i^-. We can reconstruct G by saying which a_i^ϵ get attached to which end of e, i.e. by giving a partition of the set $\{a_1^+, a_1^-, \ldots, a_n^+, a_n^-\}$ into two subsets, called the *sides* of the partition. In fact, any marked graph in CV_n can be collapsed to a rose R_n by collapsing a maximal tree, and each edge in the maximal tree gives a partition of the ends $\{a_1^+, a_1^-, \ldots, a_n^+, a_n^-\}$ of the petals of R_n. A collection of partitions corresponds to a graph if and only if the partitions are pairwise compatible (P and Q are compatible if some side of P is disjoint from some side of Q).

Now let H be a hyperplane in a Γ-complex X, and suppose the corresponding hyperplane collapse is defined and gives the standard Salvetti (S_Γ, id). We will partition the edges in the 1-skeleton of S_Γ by looking at their pre-images in the 1-skeleton of $Y = X - (H \times (0,1))$. By condition (2) the image of H must be parallel to a hyperplane in S_Γ, say the hyperplane $S_{lk(v)}$ dual to v. If a_i is an edge in this image, then a_i has two pre-image loops in the 1-skeleton of Y, at the top and bottom of the hyperplane carrier $H \times [0,1]$. All other edges in S_Γ have one pre-image, and we will partition their ends according to whether they are attached at the top or bottom of the carrier. We cannot partition these arbitrarily, however, since there are constraints imposed by existence of the higher-dimensional cubes. For example, we cannot put the meridian of a torus at one vertex and the longitude at the other. Careful consideration of these constraints leads to the following definition.

Definition 5.7 Let m be a vertex of Γ, L_m the vertices in $lk(m)$ and $V_m^\pm$ the set of vertices in $\Gamma - lk(m)$ and their inverses. A partition P of $V_m^\pm$ is a Γ-*Whitehead partition* if

1. Each side of P has at least two elements.
2. Each side of P is a union of (vertices of) components of $\Gamma - lk(m)$ and their inverses *except*
   - P separates m from m^{-1} and
   - if $lk(a) \subseteq lk(m)$ then P may separate a from a^{-1}.

The vertex m is called a *maximal vertex for P*. Note that a Γ-Whitehead partition may have more than one maximal vertex but any two maximal vertices have the same link,

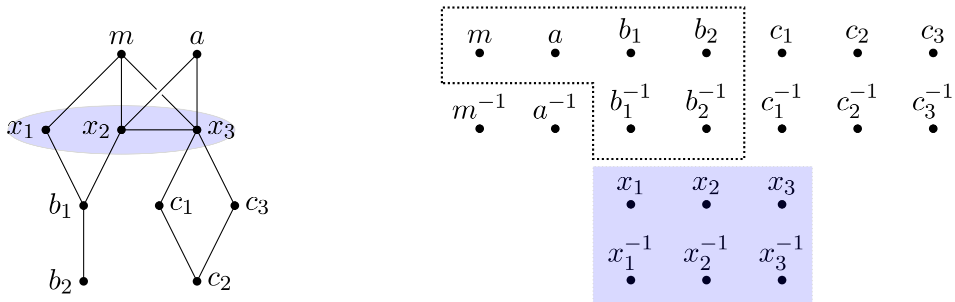

Figure 10. A graph Γ and a Γ-Whitehead partition

which we therefore call $lk(P)$.

An example of a Γ-Whitehead partition is shown in Figure 10.

The terminology "Γ-Whitehead" has a historical basis. Suppose the rose R_n is blown up by inserting a single edge e which partitions the half-edges of R_n. Then collapsing a different edge of the blowup gives a homotopy equivalence $R_n \to R_n$ which induces a *Whitehead automorphism* of F_n. If the newly collapsed edge was labeled with the generator a, this automorphism multiplies some generators by a (or a^{-1}) and conjugates some others by a (or a^{-1}).

Not every Whitehead automorphism of the free group on the generators of A_Γ induces an automorphism of A_Γ, but we can tell exactly which ones do. If P is a Γ-Whitehead partition as defined above, we can complete P to a partition $\widehat{P}$ of all of the generators of A_Γ and their inverses by putting $L^\pm = L_m \cup L_m^{-1}$ on one side of P (it doesn't matter which side). Then the induced Whitehead automorphism of the generators does give an automorphism of A_Γ.

We have seen that a two-vertex Γ-complex in $st(\sigma)$ gives a Γ-Whitehead partition. Conversely, given a Γ-Whitehead partition we can construct a two-vertex Γ-complex, which we call S^P. Here are instructions for its construction.

- Start with a copy of $S_{lk(P)} \times [0,1]$.
- For each a which is separated from a^{-1} by P (including each maximal vertex), glue in a copy of $S_{lk(a)} \times [0,1]$, attaching $S_{lk(a)} \times \{i\}$ by its inclusion into $S_{lk(P)} \times \{i\}$ for $i = 0,1$.
- For each remaining component C of $\Gamma - lk(P)$ attach a copy of $S_{lk(P)\cup C}$ via the inclusion of $lk(P)$, where components on opposite sides of P are attached at opposite ends of $S_{lk(P)} \times [0,1]$.

Collapsing the initial subcomplex $S_{lk(P)} \times [0,1]$ to its hyperplane $S_{lk(P)} \times \{\frac{1}{2}\}$ recovers the Salvetti complex S_Γ, so S^P is a Γ-complex in $st(\sigma)$.

Remark 5.8 For any vertex a of Γ the subcomplex $S_{st(a)}$ of S_Γ is a product $S_{lk(a)} \times S^1$. Another way of describing S^P is as the union of these subcomplexes, some of which have been subdivided by hyperplanes, modulo appropriate identifications. In particular, $S_{st(a)}$ embeds into the blowup S^P

5.1.5 Compatible partitions and iterated blowups

The Γ-complexes in $st(\sigma)$ which have exactly two vertices are those which can be obtained from σ by blowing up a single Γ-Whitehead partition. In order to obtain any Γ-complex in $st(\sigma)$ we may need to blow up several times. This is possible if we are given a collection of Γ-Whitehead partitions which are compatible, in the following sense.

Definition 5.9 Two Γ-Whitehead partitions P and Q are *compatible* if either

1. maximal elements of P and Q are distinct and commute, or
2. some side of P is disjoint from some side of Q

A precise recipe for constructing a Γ-complex from a collection of pairwise compatible Γ-Whitehead partitions is given in [8]. We omit the details here.

5.1.6 Contractibility of the spine

The proof that the spine K_Γ is contractible follows the general outline of the original proof that CV_n is contractible [9]. A vertex (S, h) of K_Γ is called a *Salvetti vertex* if S is homeomorphic to S_Γ. We define a total order on Salvetti vertices (S, h) by measuring the lengths of conjugacy classes of elements of A_Γ in S. More precisely, for each conjugacy class $w \in \pi_1(S_\Gamma) = A_\Gamma$ we record the length of the minimal loop in the 1-skeleton of S that represents $h(w)$, then list all these lengths in an infinite sequence. We then build K_Γ by gluing on stars of Salvetti vertices according to the lexicographical order of these sequences. We need to prove that a Salvetti vertex is determined by its length sequence, that there is a unique smallest Salvetti vertex, and that at each stage of the construction we are attaching the next star along a contractible subcomplex of its link. The proof of this last fact uses a variation of the classical Peak Reduction algorithm for free group automorphisms.

5.2 The full outer space

In order to get a contractible space on which all of $Out(A_\Gamma)$ acts properly, we need to add metric information to the marked Γ-complexes used to define K_Γ.

To explain the idea, we again recall the relation of the spine K_n to the full Outer space CV_n, from a slightly different point of view. The full space CV_n decomposes as a disjoint union of open simplices of various dimensions, where the simplex containing (X, h) is obtained by varying the (positive) lengths of the edges of X. If we allow an edge length to shrink to zero, we pass to a face of the simplex. Thus the closure of $\sigma(X, h)$ in CV_n is a simplex together with some of its faces, but some faces are missing: if we try to shrink a set of edges containing a loop to zero, we leave CV_n. If we formally add all of the missing faces to each $\sigma(X,h)$ we obtain a simplicial complex $\overline{CV}_n$ called the *simplicial closure of Outer space*. The spine K_n is a subcomplex of the barycentric subdivision $\overline{CV}'_n$, namely K_n is the subcomplex spanned by vertices of $\overline{CV}_n$ (i.e. faces of $\overline{CV}_n$) which are actually in CV_n.

We can verify that CV_n is homotopy equivalent to K_n using a suitable open cover of CV_n. For each vertex $v \in K_n$, let $U_v = st^o(v) \subset CV_n$ be the open star of v in $\overline{CV}'_n$. Each U_v is a contractible subset of CV_n and contains no other vertices of K_n. An intersection $U_{v_0} \cap \ldots \cap U_{v_k}$ is non-empty if and only if $v_0, \ldots, v_k$ are the vertices of a simplex of K_n, in which case the intersection is contractible. Thus the nerve of the cover $\{U_v\}$ is isomorphic to K_n and is homotopy equivalent to CV_n.

For general A_Γ we would like to do something similar, i.e. add metric information to marked Γ-complexes to produce a space of marked metric Γ-complexes and an open cover by sets $\{U_v\}$ corresponding to vertices $v \in K_\Gamma$. We want the nerve of this cover to be isomorphic to K_Γ and the cover to be by contractible sets with contractible intersections, so that the whole space is homotopy equivalent to K_Γ (and hence contractible).

5.2.1 Untwisted metrics

As in the free group case we can assign positive lengths to the edges of the (rectilinear) cubes of a Γ-complex X to obtain a set of metrics on X which forms an open simplex $\sigma(X,h)$ of marked metric Γ-complexes, one for each vertex $v = (X,h)$ of K_Γ. Let Σ_G denote the union of these open simplices, modulo the natural face relations. Formally adding missing faces to the simplices $\sigma(X,h)$ completes Σ_Γ to a simplicial complex $\overline{\Sigma}_\Gamma$. The spine K_Γ is a subcomplex of the barycentric subdivision $\overline{\Sigma}'_\Gamma$, and the space Σ_Γ is covered by open stars $st^o(v)$ in $\overline{\Sigma}'_\Gamma$ of vertices $v \in K_n$. The action of $U(A_\Gamma)$ on K_Γ extends naturally to a proper action on Σ_Γ and $\overline{\Sigma}_G$.

In the free group case this is all we needed to do, but for general A_Γ this is not enough...we also need to allow the metrics on some cubes to be sheared in order to get a space on which all of $Out(A_\Gamma)$ acts properly. This shearing is governed by the subgroup $T(A_\Gamma)$ generated by twists, so we next investigate this subgroup.

5.2.2 Twisted metrics

If we order the generators $\{a_1, \ldots, a_j\}$ of A_Γ we obtain a map $Out(A_\Gamma) \to GL(n,\mathbb{Z})$ induced by abelianization $A_\Gamma \to \mathbb{Z}^n$. This map sends the twist subgroup $T(A_\Gamma)$ injectively into $SL(n,\mathbb{Z})$. The image of $T(A_\Gamma)$ is generated by the matrices $A_{ij} = I_n + E_{ij}$ for i,j such that $st(a_i) \subseteq st(a_j)$. If the ordering of the generators is subordinate to the partial ordering on the vertices of Γ, the image of $T(A_\Gamma)$ is block upper triangular. The diagonal blocks correspond to equivalence classes of vertices and a non-zero upper block corresponds to an inclusion of stars.

Now let $T_{\mathbb{R}}(A_\Gamma) \subset SL(n,\mathbb{R})$ be the subgroup generated by matrices $A_{ij}(r) = I_n + E_{ij}(r)$ with r real, for i, j such that $st(a_i) \subseteq st(a_j)$. $T_{\mathbb{R}}(A_\Gamma)$ is contained in a parabolic subgroup of $SL(n,\mathbb{R})$ and $T(A_\Gamma)$ is a lattice in $T_{\mathbb{R}}(A_\Gamma)$. The quotient space

$$\mathbb{D}_\Gamma = (T_{\mathbb{R}}(A_\Gamma) \cap SO(n)) \backslash T_{\mathbb{R}}(A_\Gamma)$$

is a subspace of the symmetric space $\mathbb{D}_n = SO(n)\backslash SL(n,\mathbb{R})$. It is homeomorphic to a product of one symmetric space for each diagonal block and one copy of $\mathbb{R}$ for each pair $i > j$ with $st(a_i) \subsetneq st(a_j)$. In particular, it is a contractible subspace of $\mathbb{D}_n$.

Each point of $\mathbb{D}_\Gamma$ corresponds to a marked flat metric on T^n, but we will ignore the marking for a moment. If we regard the Salvetti complex S_Γ as a subcomplex of T^n, then the flat metric on T^n induces a metric on S_Γ, where the distance between two points is the length of the shortest path in S_Γ joining them. Note that the metric on each subtorus S_Δ corresponding to a clique $\Delta \subset \Gamma$ is flat. A metric on S_Γ induced in this way by a point in $\mathbb{D}_\Gamma$ is said to be Γ-*adapted*.

If X is a blowup of S_Γ then X contains a subcomplex X_Δ for each clique Δ which is a (possibly subdivided) copy of the cube complex S_Δ. A CAT(0) metric on X is Γ-*adapted* if the metric restricted to each X_Δ is equal to the flat metric on S_Δ obtained from some Γ-adapted metric on S_Γ.

Definition 5.10 A *marked metric* Γ-*complex* is a triple (X, h, d), where

1. X is a Γ-complex,
2. $h\colon S_\Gamma \to X$ is a homotopy equivalence, and
3. d is a Γ-adapted CAT(0) metric on X.

Two marked metric Γ-complexes (X, h, d) and (X', h', d') are *equivalent* if there is an isometry or a homothety $\iota : X \to X'$ with $\iota \circ h' \simeq h$.

5.2.3 Outer space for A_Γ

We now define outer space $\mathcal{O}_\Gamma$ to be the space of equivalence classes of marked metric Γ-complexes. The group $Out(A_\Gamma)$ acts on $\mathcal{O}_\Gamma$ on the right by changing the marking, i.e. given $\phi \in Out(A_\Gamma)$, choose a homotopy equivalence $f : S_\Gamma \to S_\Gamma$ inducing ϕ on $\pi_1(S_\Gamma)$; then $(X, h, d) \cdot \phi = (X, h \circ f)$.

Claim. *Outer space* $\mathcal{O}_\Gamma$ *is contractible, and* $Out(A_\Gamma)$ *acts properly.*

Caveat. I have refrained from calling this a Theorem since the details have not yet been posted on the arXiv.

Sketch of proof: We cover $\mathcal{O}_\Gamma$ by open sets U_v corresponding to vertices $v = (X, h)$ of K_Γ. Each U_v is homeomorphic to the product of $\mathbb{D}_\Gamma$ with the open star $st^o(v)$ of v in the barycentric subdivision $\overline{\Sigma}'_\Gamma$, and is hence contractible. The nerve of the cover $\{U_v\}$ is isomorphic to K_Γ and intersections $U_{v_1} \cap \ldots \cap U_{v_k}$ are either empty or contractible. Thus $\mathcal{O}_\Gamma$ is homotopy equivalent to K_Γ, which is contractible by Theorem 5.6. Finally, one must check the stabilizer of a point (X, h, d) in $\mathcal{O}_\Gamma$ under the action of $Out(A_\Gamma)$. The action of a twist moves (X, h, d) "up the $\mathbb{D}_\Gamma$-direction," and the stabilizer of (X, h, d) is isomorphic to the group of isometries of (X, d), which is finite.

5.3 Questions

1. The action of $Out(A_\Gamma)$ on O_Γ is not cocompact, since we're using *all* Γ-adapted metrics on the Γ-complexes X. Inside this space of metrics there should be an analog of Ash's *well-rounded retract* of $SO(n)\backslash SL(n, \mathbb{R})$, which is a cocompact deformation retract, equivariant with respect to the action of $SL(n, \mathbb{Z})$. Incorporating this idea should result in an outer space with a cocompact action.
2. Is the fixed point set of a finite subgroup of $Out(A_\Gamma)$ contractible (i.e. is $\mathcal{O}_\Gamma$ an $\underline{EG}$?)? Is it even non-empty, i.e. can every finite subgroup of $Out(A_\Gamma)$ be realized as isometries of a marked Γ-complex?
3. Is $Out(A_\Gamma)$ a virtual duality group? Is there a *bordification* of $\mathcal{O}_\Gamma$ which is a hybrid of the Borel-Serre bordification of the symmetric space $\mathbb{D}_n$ and the Bestvina-Feighn bordification of Outer space CV_n? If so, is bordified $\mathcal{O}_\Gamma$ highly connected at infinity?
4. The *metric theory* of symmetric spaces is classical and highly developed. There has also been a lot of activity recently on the metric theory of Outer space, using the asymmetric Lipschitz metric. Is there a good metric theory of $\mathcal{O}_\Gamma$? What are the geodesics? Can they be used to help classify elements of $Out(A_\Gamma)$?
5. Handel and Mosher recently proved that the 1-skeleton of the simplicial closure $\overline{CV}_n$ is a Gromov hyperbolic graph [16]. Is the 1-skeleton of $\overline{\mathcal{K}}_\Gamma$ Gromov hyperbolic? If so, is there an associated Gromov hyperbolic space on which all of $Out(A_\Gamma)$ acts?

References

[1] Gilbert Baumslag and Tekla Taylor, The centre of groups with one defining relator, *Math. Ann.* **175** 1968 315–319.

[2] Kai-Uwe Bux, Ruth Charney and Karen Vogtmann, Automorphism groups of RAAGs and partially symmetric automorphisms of free groups, *Groups Geom. Dyn.* **3** (2009), no. 4, 541–554.

[3] Martin Bridson and Andre Haefliger, *Metric spaces of non-positive curvature*, Springer-Verlag, New York, (1999).

[4] Ruth Charney, An introduction to right-angled Artin groups, *Geom. Dedicata* **125** (2007), 141–158.

[5] Ruth Charney, John Crisp and Karen Vogtmann, Automorphisms of two-dimensional right-angled Artin groups, *Geom. Top.* **11** (2007), 2227–2264

[6] Ruth Charney and Karen Vogtmann, Finiteness properties of automorphism groups of right-angled Artin groups, *Bull. Lond. Math. Soc.* **41** (2009), no. 1, 94–102.

[7] Ruth Charney and Karen 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.

[8] Ruth Charney, Nate Stambaugh and Karen Vogtmann, *Outer space for right-angled Artin groups, I*, arXiv:1212.4791

[9] Marc Culler and Karen Vogtmann, Moduli of graphs and automorphisms of free groups, *Invent. Math.* **84** (1986), no. 1, 91–119.

[10] Michael Davis and Tadeusz Januszkiewicz, Right-angled Artin groups are commensurable with right-angled Coxeter groups, *J. Pure Appl. Algebra* **153** (2000), no. 3, 229–235.

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

[12] Matthew B. Day, Full-featured peak reduction in right-angled Artin groups, *arXiv:1211.0078.*

[13] Carl Droms, Subgroups of graph groups, *J. Algebra* **110** (1987), no. 2, 519522.

[14] S. M. Gersten, A presentation for the special automorphism group of a free group, *J. Pure Appl. Algebra* **33** (1984), no. 3, 269–279.

[15] Vincent Guirardel and Gilbert Levitt, The outer space of a free product, *Proc. Lond. Math. Soc.* (3) **94** (2007), no. 3, 695–714.

[16] Michael Handel and Lee Mosher, The free splitting complex of a free group I: Hyperbolicity, arXiv:1111.1994.

[17] Frederic Haglund, Daniel T. Wise, Special cube complexes, *Geom. Funct. Anal.* **17** (2008), no. 5, 1551–1620.

[18] A. H. M. Hoare, Coinitial graphs and Whitehead automorphisms, *Can J. Math* **21** (1979), no. 1, 112–123.

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

[20] Wilhelm Magnus, Über n-dimensionale Gittertransformationen, *Acta Math.* **64** (1935), no. 1, 353–367.

[21] James McCool, Some finitely presented subgroups of the automorphism group of a free group, *J. Algebra* **35** (1975), 205–213.

[22] Ashot Minasyan, Hereditary conjugacy separability of right-angled Artin groups and its applications, *Groups Geom. Dyn.* **6** (2012), no. 2, 335–388.

[23] Ashot Minasyan and Denis Osin, Normal automorphisms of relatively hyperbolic groups, *Trans. Amer. Math. Soc.* **362** (2010), no. 11, 6079–6103.

[24] Daniel Quillen, *Higher Algebraic K-Theory, I*, Lect. Notes Math. **341** (1974), 85–147.

[25] Herman Servatius, Automorphisms of graph groups, *J. Algebra* **126** (1989), no. 1, 34–60.

[26] Herman Servatius, Carl Droms and Brigitte Servatius, Surface subgroups of graph groups *Proc. Amer. Math. Soc.* **106** (1989), no. 3, 573578.

[27] John R. Stallings, Topology of finite graphs, *Invent. Math.* **71** (1983), no. 3, 551–565.

[28] Nate Stambaugh, Toward an outer space for right-angled Artin groups, Ph.D. Dissertation, Brandeis University, August, 2011.

[29] Emmanuel Toinet, Conjugacy p-separability of right-angled Artin groups and applications, to appear in *Groups, Geometry and Dynamics.*

[30] Anna Vijayan, Compactifying the space of length functions of a right-angled Artin group, Ph.D Dissertation, Brandeis University, December, 2012.