Title
The Residual Finiteness of Triangle Artin Groups

Author
Meyer, Greyson Pesut

Publication Date
2025

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

THE RESIDUAL FINITENESS OF TRIANGLE ARTIN GROUPS

A dissertation submitted in partial satisfaction
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

MATHEMATICS

by

Greyson Meyer

March 2025

The Dissertation of Greyson Meyer
is approved:

Professor Martin Weissman, Chair

Assistant Professor Kasia Jankiewicz

Nic Brody, PhD

Peter Biehl
Vice Provost and Dean of Graduate Studies

Contents

Figures and Tables

List of Figures

List of Tables

Abstract

The Residual Finiteness of Triangle Artin Groups

by

Greyson Meyer

We prove that all triangle Artin groups of the form $A_{2,3,2n}$ where $n > 3$ are residually finite. To achieve this, we use the presentation for these groups previously employed by Wu and Ye to establish that each of them splits as a graph of groups. Building on techniques developed by Jankiewicz for other triangular subclasses of Artin groups, we adapt and extend these methods to show residual finiteness in this setting. Additionally, we developed a Python program to assist in specific computations for the case of $A_{2,3,8}$.

For Miranda Lima

Acknowledgments

I feel extremely lucky to have had the opportunity to work with Kasia Jankiewicz over the last four years. Every time I came to Kasia with a question or a mathematical obstacle, she always had an answer or a new perspective or a hint that would break the entire problem open. Despite being perhaps the busiest person I have ever known, she somehow always had time to meet with me each week. Asking her to be my advisor was one of the best decisions I have ever made. I truly could not have done this without her.

I would also like to thank Martin Weissman and Nic Brody for the countless nuggets of mathematical wisdom they have transferred to me throughout my time in the program. These three mathematicians made this experience fun and approachable. It is an honor to have this group read my dissertation.

And, of course, I would like to thank my friends and family for putting up with the highs and lows of discovery and failure that these last five years have been. I hope something here might make it all worth it in the end.

Part I.

Residual Finiteness of Triangle Artin Groups

Chapter 1.

Introduction

Given a group G, asking whether the word problem is solvable in G is a fundamental question to ask about its structure. Artin groups are a class of groups for which the solvability of the word problem is an open conjecture. There have been numerous efforts toward resolving this conjecture, many of which have been successful for particular subclasses of Artin groups. Squier proved algebraically that three triangle Artin groups split as graphs of free groups. These splittings are such that the residual finiteness of these Artin groups is apparent from the splittings. As will be proven later, residually finite groups have solvable word problem, making residual finiteness a property that group theorists desire Artin groups to have.

Jankiewicz desired to extend this result to other triangle Artin groups, and in doing so found that a similar procedure could be used to prove that certain classes of triangle Artin groups are also residually finite [[1] [2] [3]]. The work contained in this dissertation documents and continues this journey toward residual finiteness for the entire class of triangle Artin groups. The methods used do not depend on any structure inherent to triangle Artin groups, and thus present a hope for similar methods to be used to prove that other Artin groups, beyond the triangular class, are residually finite and have solvable word problem as well.

Section 2.1 will introduce Artin groups, along with known results about their structure. In order to understand and extend the work of Jankiewicz, we must introduce the fundamentals of Bass-Serre Theory, which will be explored in Section 2.2. With Bass-Serre Theory at our fingertips, we will then explore the concept of residual finiteness in more detail in Section 2.3. Finally, in Section 2.4, we have enough tools to understand how previous residual finiteness results were constructed in the context of triangle Artin groups.

Section 3 focuses on new results discovered by the author with the help of Jankiewicz. In Section 3.1 we discuss the structure of a particular class of triangle Artin groups, namely those of the form $A_{2,3,2n}$ for $n > 3$. In Section 3.2 we develop the methodology used in Section 3.3 to prove that all Artin groups $A_{2,3,2n}$ for $n > 4$ are residually finite. We end with Section 3.4 in which we prove, using the tools from Section 3.2, that the anomalous $A_{2,3,8}$ is residually finite as well.

Chapter 2.

Preliminaries

§ 2.1. Artin Groups

The star of this dissertation is a peculiar class of groups called Artin groups.

Definition 2.1.1 ([4])**.** An Artin group is a group that admits a presentation $\langle S|R\rangle$ where every relation in R takes the form $sts\ldots = tst\ldots$ and both sides of the equality are words of equal length. Each pair $\{s,t\} \subseteq S$ has at most one relation of this type in R.

Example 2.1.1. The Artin group $A_{2,3,7} = \langle a,b,c|ac = ca, aba = bab, bcbcbcb = cbcbcbc\rangle$

Since these relations have such a predictable shape, it is common to refer to an Artin group by its Coxeter diagram.

Definition 2.1.2 ([5])**.** The Coxeter diagram associated to an Artin group is a labeled graph constructed in the following manner:

- One vertex for each generator,
- An unlabeled edge between vertices s and t when $sts = tst$,
- A labeled edge between vertices s and t where $st \neq ts$ and $sts \neq tst$ labeled by the length of the relation involving the pair $\{s,t\}$ in R. If no such relation exists, the edge is labeled with ∞.

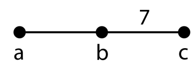

Figure 2.1.: Coxeter diagram for $A_{2,3,7}$

Due to the correspondence between the Artin relations and labelled graphs of this form, Artin groups are often referred to as $A(\Gamma)$ where Γ is the Coxeter diagram describing its relations. Coxeter diagrams originated with regards to the class of groups that shares its namesake.

Definition 2.1.3 ([4])**.** A Coxeter group is a group that admits a presentation

$$\langle s_1, s_2, ..., s_n | (s_i s_j)^{m_{i,j}} = 1, m_{i,i} = 1 \rangle$$

The connection between Artin groups and Coxeter groups via Coxeter diagrams is no coincidence. Artin groups were first introduced by Jacques Tits as a natural extension of Coxeter groups. If you remove the requirement that Coxeter groups are generated by involutions, then the resulting presentation is precisely an Artin presentation. Equivalently, every Artin group has a Coxeter group associated to it that can be realized by performing the quotient $A(\Gamma)/\langle\langle s_1^2, s_2^2, ..., s_n^2 \rangle\rangle$ where $A(\Gamma)$ has $\{s_1, s_2, ..., s_n\}$ as its generating set.

An important consequence of the association between Artin groups and Coxeter groups is the ability to now classify Artin groups based on properties that their associated Coxeter groups have. The first such subclass of Artin groups is called spherical Artin groups.

Definition 2.1.4 ([5])**.** An Artin group is called spherical if its associated Coxeter group is finite.

Unlike Artin groups which are always infinite (every generator has infinite order, for example), there exist Coxeter groups that are finite. Spherical Artin groups are the most well understood subclass of Artin groups. In modern terminology, spherical Artin groups are Garside groups [6]. A Garside group is the group of fractions of a Garside monoid. A Garside monoid is a monoid that is finitely generated, cancellative, its partial order with respect to divisibility gives us a lattice (gcd's and lcm's exist) and it contains a Garside element (an element whose left divisors are the same as its right divisors). Garside used this structure to prove a number of results about braid groups, and since then Artin group enthusiasts have applied similar approaches to spherical Artin groups. The Garside structure of spherical Artin groups has allowed group theorists to prove that spherical Artin groups have solvable word problem [6, 7], solvable conjugacy problem [8], and trivial torsion [9].

Unfortunately, non-spherical Artin groups, meaning Artin groups whose associated Coxeter groups are infinite, are not Garside groups. The barrier keeping these groups from being Garside groups is the lack of a Garside element. The Garside element arises from taking an element of maximal length from the Coxeter group W and pulling it back along the quotient map $q : A \to W$ in order to obtain an element of the Artin group A. When the Coxeter group is infinite, no such element of maximal length exists.

Another well-studied subclass of Artin groups are the Right Angled Artin Groups, or RAAGs for short. RAAGs are the Artin groups in which each pair of generators either commutes or there is no relation between them. They can also be thought of as Artin groups whose presentations are in correspondence with an unlabeled graph without length-1 loops or multiedges. Given such a graph Γ, we can construct a RAAG A_Γ by defining the generating set to be the vertices, and include a commuting relation between any two vertices that are joined by an edge. This graph Γ can actually tell us information about the RAAG just from its shape. Some results of this form are the following:

- A RAAG A_Γ is a direct product if and only if Γ can be partitioned into disjoint sets U_1 and U_2 where there is an edge between every $u_1 \in U_1$ and $u_2 \in U_2$ (aka Γ is a join). [4]
- A_Γ is a free product if and only if Γ is disconnected (this is immediately implied by the lack of relations between the genrators of different connected components).

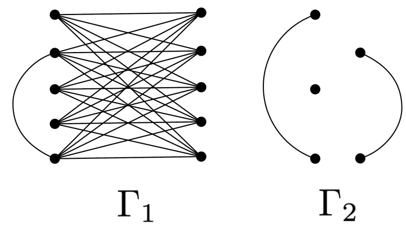

Figure 2.2.: $A(\Gamma_1) = \langle \sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_5 | \sigma_2\sigma_5 = \sigma_5\sigma_2 \rangle \times \langle \sigma_6, \sigma_7, \sigma_8, \sigma_9, \sigma_{10} \rangle$
& $A(\Gamma_2) = \langle \sigma_1, \sigma_3 | \sigma_1\sigma_3 = \sigma_3\sigma_1 \rangle * \langle \sigma_2 \rangle * \langle \sigma_4, \sigma_5 | \sigma_4\sigma_5 = \sigma_5\sigma_4 \rangle$

Another surprising result about RAAGs is the following:

Theorem 2.1.1. [4] *Every RAAG is isomorphic to a subgroup of finite index in some right-angled Coxeter group.*

Proof. In order to prove this theorem, we explicitly construct the embedding. Start with a RAAG A_Γ. Now construct a new graph $\tilde{\Gamma}$ so that each vertex $s_i \in V(\Gamma)$ now corresponds to two vertices $s_i, s_i' \in V(\tilde{\Gamma})$. Connect each s_i' to every other vertex except s_i. Now the embedding $\phi : A_\Gamma \to W_{\tilde{\Gamma}}$ comes from sending $s_i \mapsto s_i s_i'$. By multiplying on the right by s_i' we have $s_i s_i'$ being of infinite order since s_i and s_i' do not commute.

Observe that $\phi(s_i s_j s_i^{-1} s_j^{-1}) = s_i s_i' s_j s_j' s_i'^{-1} s_i^{-1} s_j'^{-1} s_j^{-1} = s_i s_i' s_i'^{-1} s_i^{-1} s_j s_j' s_j'^{-1} s_j^{-1} = s_i s_i^{-1} s_j s_j^{-1} = 1$ makes ϕ well-defined. Injectivity follows immediately from the fact that ϕ preserves the commutativity of A_Γ and $\phi(x)$ cannot contain any squares of generators of $W_{\tilde{\Gamma}}$ by the definition of ϕ.

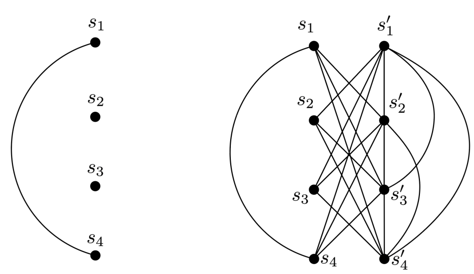

Figure 2.3.: Γ & $\tilde{\Gamma}$

Another fascinating class of Artin groups is the class of FC-type Artin groups. The definition of this class of groups comes explicitly from a property of the Coxeter diagram.

Definition 2.1.5 ([10]). An FC-type Artin group A_Γ is an Artin group where if T is a subset of the generating set S where every $t_i, t_j \in T$ has a finite Artin relation $R(t_i, t_j)$ in A_Γ, then the Artin subgroup $A_T \leq A_\Gamma$ is spherical.

It is easier to understand this definition by constructing an analogue $\tilde{\Gamma}$ of the Coxeter diagram for A_Γ. We again have a vertex in $\tilde{\Gamma}$ for every generator of A_Γ, but now we include an edge between two vertices if they have an Artin relation, including the case when the generators commute. We then label all edges with the length of the corresponding Artin relation. So $\tilde{\Gamma}$ can have edges labelled with 2 and 3, but no edges labelled with ∞. We can now rephrase Definition 2.1.5 by saying that A_Γ is FC-type if the Artin subgroup generated by every clique in $\tilde{\Gamma}$ is spherical. A clique in a graph is a subgraph containing vertices that are all pairwise adjacent to each other. Essentially, FC-type Artin groups are built from spherical Artin subgroups. Instead of requiring the associated Coxeter group to be finite, we require that specific Artin subgroups have a finite associated Coxeter group. FC-type Artin groups have a number of well-understood properties, chief of which is that they have solvable word problem [10]. Figure 2.4 shows an example of a graph whose corresponding Artin group is FC-type.

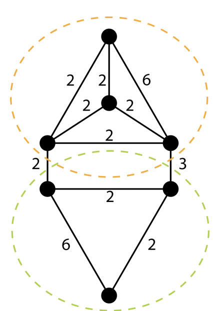

Figure 2.4.: Two maximal cliques in the diagram for an FC-type Artin group.

At the other end of the spectrum there is the class of 2-dimensional Artin groups.

Definition 2.1.6 ([2])**.** A_Γ is a 2-dimensional Artin group if no triple of its generators generates a spherical Artin group.

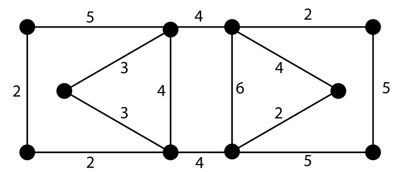

Figure 2.5.: Diagram for a 2-Dimensional Artin Group

Every 2-dimensional Artin group is built from explicitly non-spherical Artin subgroups. The Artin groups present in further sections will be primarily within the 2-dimensional class of Artin groups, specifically 2-dimensional Artin groups with 3 generators.

We have now briefly introduced a couple of the more obscure classes of Artin groups, but there are a number of other well known groups that exist under the Artin umbrella, namely: braid groups, free groups and free Abelian groups. Not only are Artin groups a fascinating class of groups in and of themselves, but they also count some of the most highly studied groups among their ranks.

But even though some classes of Artin groups have been well-studied, Artin groups as a class remain mysterious. Despite their pleasing and predictable group presentations, there are many simple questions that one can ask about Artin groups that still do not have definitive answers. Some of the big questions about Artin groups are the following:

- Are all Artin groups torsion-free?
- Which Artin groups have trivial center?
- Which Artin groups have solvable word & conjugacy problems?
- Isomorphism problem: When do two graphs Γ_1 & Γ_2 result in $A_{\Gamma_1} \cong A_{\Gamma_2}$?

As of yet these questions do not have universal answers, though there are answers known for certain classes of Artin groups. We will not be engaging with any of these questions directly for the remainder of the paper. Instead, we will be focusing on the concept of residual finiteness in the class of triangle Artin groups.

Definition 2.1.7 ([1])**.** A triangle Artin group is an Artin group that has 3 generators in its standard presentation.

Before we explore residual finiteness, we will take a brief detour into Bass-Serre Theory in order to motivate the context in which residual finiteness will later appear.

§ 2.2. Bass-Serre Theory

We begin this section with The Fundamental Theorem of Bass-Serre Theory, seeing as it beautifully encapsulates many of the concepts relevant for the later sections. We then will take time to explain what all of the terms mean and give a sketch of the proof.

Theorem 2.2.1 (The Fundamental Theorem of Bass-Serre Theory [11])**.** *Let* (Γ, G) *be a graph of groups and* $v \in \Gamma$ *a vertex. Then there exists a group* $H = \pi_1(\Gamma, G, v)$ *and a tree* T *such that* H *acts on* T *without inversions and* $T/H \cong (\Gamma, G)$.

In summation, The Fundamental Theorem of Bas-Serre Theory tells us that every graph of groups comes with a tree (called the Bass-Serre tree) on which its fundamental group acts, and the resulting quotient is the original graph of groups. Let us now break down each piece of the theorem.

Definition 2.2.1. A graph of groups, typically denoted (Γ, G), consists of an underlying graph Γ and a collection of groups G. The groups in G are assigned to the vertices and edges of Γ such that every edge group G_e injects into its adjacent vertex group(s).

Figure 2.6 shows two examples of graphs of groups, albeit without specifying the edge-group injections. We now have access to a mathematical object with the underlying structure of a graph, but whose every component contains a group-theoretic companion.

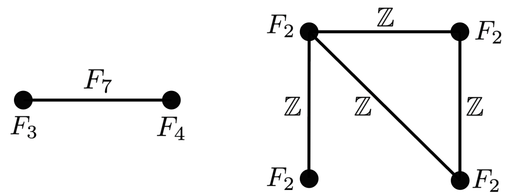

Figure 2.6.: Examples of graphs of groups.

Now that we understand the structure of graphs of groups, our next task is to understand how to calculate the fundamental group of such an object. A path through a graph of groups is a path through the underlying graph Γ, where at each vertex v in the path we take an element from the vertex group G_v in order to make a word of the form $g_0 e_0 g_1 e_1 \dots e_n g_n$. Here the g_i are elements from vertex groups and the e_i keep track of the which edges have been traversed in Γ. So a loop at a vertex v in (Γ, G) would be a word that begins and ends with a group element from the same vertex-group.

But fundamental groups are more than just groups built from loops in a space, they are homotopy classes of loops. Notice that thus far we have not utilized any information about the edge groups. Homotopies in graphs of groups are precisely where the edge groups are utilized. A homotopy in a graph of groups consists of the application of the following:

- The addition or removal of a subpath of the form $\sigma\sigma^{-1}$.
- For all $h \in G_e$, $e\iota_e(h) = \iota_{\bar{e}}(h)e$.

The first bullet point is standard. You can of course always add or remove cancellable subpaths. The second bullet point is the more interesting of the two. Here we have ι_e and $\iota_{\bar{e}}$ denoting the two injective group homomorphisms from G_e to its adjacent vertex group(s). Such a homotopy essentially allows the image of an element $h \in G_e$ in one vertex group to slide across the edge e to the other vertex group, changing from $\iota_e(h)$ to $\iota_{\bar{e}}(h)$ in the process.

We could attempt to calculate the fundamental group of a graph of groups in this way, by partitioning the loops into homotopy classes, but luckily there is a simpler way. Just like we can calculate the fundamental group of a graph by collapsing a maximal subtree and counting the remaining loops, graphs of groups have an analogous procedure. We can choose a maximal subtree $T \subseteq \Gamma$ and collapse it, but since there are algebraic objects involved with each vertex and edge being collapsed, we need to be a little more careful. We start by fixing a basepoint v and then choosing any edge e in T adjacent to v. We collapse e, identifying v', the other vertex adjacent to e, with v. In doing so, we now replace the vertex group G_v with $G_v *_{G_e} G_{v'}$, the amalgamated product of G_v and $G_{v'}$ along G_e.

Definition 2.2.2 ([11])**.** The amalgamated product of two groups $A = \langle S_A | R_A \rangle$ and $B = \langle S_B | R_B \rangle$ along a common subgroup C is a group denoted $A *_C B$ with the presentation $A *_C B = \langle S_A, S_B | R_A, R_B, \phi_A(c) = \phi_B(c) \forall c \in C \rangle$ where $\phi_A : C \to A$ and $\phi_B : C \to B$ are embeddings of C into A and B respectively.

Definition 2.2.3 ([11])**.** Let $A = \langle S_A | R_A \rangle$ and C be groups with two injective group homomorphisms $\phi_1 : C \to A$ and $\phi_2 : C \to A$. The HNN extension of A along C is the group denoted $A*_C$ with the presentation $\langle S_A, t | R_A, t\phi_1(c)t^{-1} = \phi_2(c) \forall c \in C \rangle$.

After collapsing T edge by edge in the manner previously described, we are then left with a graph consisting of one vertex v and k edges. The vertex group G_v is a potentially complicated amalgam of amalgamated products. The final step in computing $\pi_1(\Gamma, G)$ comes from the 2nd form of homotopy. Let e be one of the remaining edges.

$$e\iota_e(h) = \iota_{\bar{e}}(h)e \implies e\iota_e(h)e^{-1} = \iota_{\bar{e}}(h).$$

Notice that this looks identical to the new relation that arises in an HNN extension. To finish the calculation, we take the HNN extension of G_v along each G_e, introducing k new variables to G_v, corresponding to the remaining edges in Γ. At last we arrive at $\pi_1(\Gamma, G, v)$ as being an iterated HNN extension of an iterated amalgam of amalgamated products. While this is certainly complicated in abstract, the graphs of groups that we will be encountering in this paper are quite simple. In fact, their fundamental groups will simply be amalgamated products or HNN extensions most of the time.

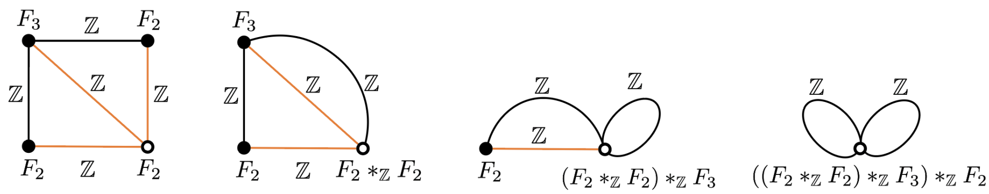

Figure 2.7.: Collapsing a maximal tree in a graph of groups.

Now that we understand the fundamental group that appears in the theorem, we move on to discuss the tree T on which this group will act. This tree is called the Bass-Serre tree of the graph of groups (Γ, G) and is constructed in the following manner. Its vertices are in bijection with $\bigcup_{v\in\Gamma} \pi_1((\Gamma, G))/G_v$. Similarly, its edges are in bijection with $\bigcup_{e\in\Gamma} \pi_1((\Gamma, G))/G_e$. Two vertices gG_{v_1} and gG_{v_2} in T are connected by the edge gG_e when v_1 and v_2 are connected by e in Γ.

The Bass-Serre tree is constructed in this way because we want the quotient of the action of $\pi_1((\Gamma, G))$ on T to result in (Γ, G). Consider a vertex $v \in \Gamma$. Then there is a collection of vertices in T whose vertex cosets are in bijection with $\pi_1((\Gamma, G))/G_v$. The action of $\pi_1((\Gamma, G))$ on T will permute this collection of vertices, allowing us to choose the vertex associated with $1G_v$ as the quotient representative for this orbit. This is true for each v, resulting in the vertex sets of (Γ, G) and $T/\pi_1((\Gamma, G))$ being in one-to-one correspondence. Similarly for the edges in T.

The Fundamental Theorem of Bass-Serre Theory tells us that for each graph of groups (Γ, G), we automatically get a tree T on which it acts. One fascinating thing about Bass-Serre theory is that it actually works the other way as well.

Theorem 2.2.2. [11] *Let H be a group acting on a tree T without edge inversions. Then T/H is a graph of groups where for each vertex $v \in T/H$, $H_v \cong Stab(v)$ and for each edge $e \in T/H$, $H_e \cong Stab(e)$.*

Moreover, in the case of T already being known to be the Bass-Serre tree of a graph of groups (Γ, G), the quotient $T/H \cong (\Gamma, G)$. When we discussed the intuition behind the construction of the Bass-Serre tree, we mentioned the collection $\pi_1((\Gamma, G))/G_v$ for each v. The coset $1G_v \cong G_v$ is one such element, and clearly $Stab(G_v) = G_v$. So if we choose G_v to be the orbit representative of $H \cdot 1G_v$, then the quotient T/H must be precisely (Γ, G).

In many ways, Bass-Serre Theory is the study of the structures that arise from groups acting on trees. While graphs of groups, amalgamated free products and Bass-Serre trees are the only objects that we will really be using from Bass-Serre theory in the rest of the paper, we will take time now to explore some of the consequences that arise from this beautiful area of study. We begin with the following elegant and powerful theorem.

Theorem 2.2.3. *If a group H acts freely on a tree, then H is a free group*

Proof. Let H act freely on a tree T. Then we can quotient T by H to get a graph of groups (Γ, G). By The Fundamental Theorem of Bass-Serre Theory, $H \cong \pi_1(\Gamma, G, v)$. Since the action of H on T is free, there are no vertex stabilizers. So T/H will have a trivial group for each G_v and G_e. So $H \cong \pi_1(\Gamma, G, v) \cong \pi_1(\Gamma, v)$ the fundamental group of a graph, which is of course a free group. ■

For the rest of the paper, when a group G is isomorphic to the fundamental group of a graph of groups, we will say that G splits as a graph of groups, or that the graph of groups is a splitting of G. A nice property of groups that split as a graph of groups is that they have normal forms for their elements [11]. A normal form is a standardized way of writing every element in your group. When a group's elements have a normal form, that normal form induces a solution to the word problem. Indeed, when you have a normal form, all you have to do to check whether two elements are equivalent in G is to convert them each into their normal form and then check whether or not the two normal forms are identical. In the context of graphs of groups, the existence of a normal form allows us to see clearly that the vertex groups G_v inject into $\pi_1((\Gamma, G))$.

We now give a brief synopsis of how Serre constructs the normal form of an element in an amalgamated product G. We start with a family of groups G_i with $i \in I$ and a group C that injects into each of them. The normal form of $g \in G$ looks like $g = f(c)f_{i_1}(s_1)...f_{i_n}(s_n)$ where $f : C \to G$ and $f_i : G_i \to G$ are the canonical injections and S_i is a collection of fixed coset representatives of $C\backslash G_i$. Basically the normal form is a word in the free product of the groups with all of the C-related material at the front. The right coset representative part is what allows us to slide our elements of C along the word to the front. Consider an element x in the word g. Then x is in some coset of $C\backslash G_i$, so $x = cy$ for some $c \in C$. Now let's say that the preceding element was z. Now $zx = zcy$ and zc is also an element of a coset and can be written $c'z'$ and so on.

Another way of viewing a group that splits as a graph of groups is by viewing it as the fundamental group of a graph of spaces. A graph of spaces is a geometric object constructed in a similar manner as a graph of groups, except we now assign a connected CW-complex, instead of a group, to every vertex and edge. Analogously to the construction of graphs of groups, we require edge spaces to embed into their adjacent vertex space(s). To go from a graph of spaces to a graph of groups, one must simply calculate the fundamental group of each vertex (resp. edge) space and assign that group to the vertex (resp. edge). To go from a graph of groups to a graph of spaces, for each vertex (resp. edge), simply assign a space whose fundamental group is the vertex (resp. edge) group, ensuring that the edge spaces still satisfy the embedding requirements.

Now that we have access to plentiful Bass-Serre theoretical tools, we can focus on the primary algebraic property that we desire to extract from triangle Artin groups, namely residual finiteness.

§ 2.3. Residual Finiteness

There are many equivalent definitions for a group being residually finite. We enumerate some of them in Definition 2.3.1.

Definition 2.3.1 ([2])**.** A group G is residually finite if any of the following hold:

1. For every nontrivial $g \in G$, there exists a finite group F and a surjective group homomorphism $\phi : G \to F$ such that $\phi(g) \neq 1_F$.

2. For every nontrivial $g \in G$, there exists a finite index normal subgroup $N_g \trianglelefteq G$ with $g \notin N_g$.

3. The intersection of all finite index normal subgroups of G is trivial.

4. The intersection of all finite index subgroups of G is trivial.

Proof. We prove that the items in the list above are all equivalent.

- “$1 \implies 2$": By assumption, $g \notin \ker\phi$. By the First Isomorphism Theorem, $G/\ker\phi \cong Im\phi$. The map ϕ being surjective makes $Im\phi \cong F$. So $[G : \ker\phi] = |F| < \infty$, making $\ker\phi$ a finite index normal subgroup not containing g.

- “$2 \implies 3$": The element g being nontrivial and $g \notin N_g$ for some finite index $N_g \trianglelefteq G$ forces $g \notin \bigcap_{N \trianglelefteq_{f.i} G} N$. So there can be no nontrivial elements in the intersection.

- “$3 \implies 4$": {finite index normal subgroups} $\subseteq$ {finite index subgroups}, so $\bigcap_{H <_{f.i} G} H < \bigcap_{N \trianglelefteq_{f.i} G} N = \{1\}$, forcing the intersection to be trivial.

- “$4 \implies 1$: Let $g \in G$ be nontrivial. Definition 4 being $\bigcap_{H <_{f.i} G} H = \{1\}$ forces the existence of a finite index subgroup $H_g < G$ with $g \notin H_g$. We can use such an index-n subgroup to define a map $\phi : G \to S_n$ by mapping each $h \in G$ to the permutation $G/H_g \to h(G/H_g)$. This map is onto, making $\ker\phi \trianglelefteq G$ of finite index. Furthermore, since $H_g = Stab(1H_g)$, $\ker\phi \leq H_g$. Therefore $\ker\phi$ is a finite index normal subgroup not containing g. So $G/\ker\phi$ is a finite group equipped with the standard projection map $p : G \to G/\ker\phi$. Since $g \notin \ker\phi$, this forces $p(g) \neq 1$.

Residual finiteness allows us to view a group in a locally finite fashion. A residually finite group is a group where every nontrivial element corresponds to a nontrivial element in a finite group, meaning that the group embeds in a product of finite groups. Residual finiteness can also be seen as testing for linearity since every linear group is residually finite [12]. Braid groups were proven to be linear independently by Krammer [13] and Bigelow [14], making them residually finite as well. An example of a class of groups that are not residually finite is the class of infinite simple groups. An infinite simple group is not residually finite since its lack of nontrivial normal subgroups prohibits the existence of an N_g defined in in Definition 2.3.1 (2). As mentioned in the preface, one main reason we care about residual finiteness is because a group being residually finite means that it has solvable word problem.

Theorem 2.3.1. *If a group G is residually finite then G has solvable word problem.*

Proof. Let $g \in G$. To decide whether $g = 1_G$, we begin by enumerating all of the finite index normal subgroups of G, beginning with the subgroups of lowest index. For each such N, define $q_N : G \to G/N$ to be the quotient map. By definition, if $w \neq 1_G$, there exists a finite index normal subgroup N_g such that $q_{N_g}(g) \neq 1_{N_g}$. We therefore compute $q_N(g)$ for all finite index normal $N \trianglelefteq G$. If we compute $q_{N_g}(g) \neq 1_{N_g}$ we have revealed that $g \neq 1_G$. If the process continues indefinitely, meaning that $q_N(g) = 1_N$ for all N, then $g = 1_G$.

A nice property of residual finiteness that we will exploit throughout the paper is the following.

Lemma 2.3.2. *Let $B < A$ be a finite index subgroup of A. If B is residually finite, then A is residually finite as well.*

Proof. Let $g \in A$ be nontrivial. If $g \in A - B$, then B is a finite index subgroup not containing g. If $g \in B$, then there exists a finite index subgroup $C \leq B$ not containing g. Since C is of finite index in B and B is of finite index in A, C is of finite index in A.

A simple example of how a class of groups can be found to be residually finite is as follows:

Property 2.3.2.1. Every finite rank free group F_k is residually finite.

Proof. Let $g \in F_k$. We can represent F_k as the fundamental group of the wedge of k circles $\bigvee_k S^1$. Draw the path corresponding to g in terms of oriented edges labeled according to the k loops, as shown in the middle of Figure 2.8. In order for this graph to be a cover of $\bigvee_k S^1$, we need each of the k edges to be entering and exiting every vertex in the path. To achieve this, we complete the path to a cover of $\bigvee_k S^1$ in the following way.

Let l_i be a circle in $\bigvee_k S^1$ and v be a vertex in our g-path. If there are no l_i edges adjacent to v attach an l_i loop to v. If there is only one l_i edge adjacent to v, and the edge is leaving v, follow the path along the l_i edges starting at v. The path will end at a vertex w. Attach an l_i edge from w to v. Similarly, if there is only one l_i edge adjacent to v, and it is entering v, follow the l_i edges backwards. This path will terminate at some vertex w. Add an l_i-edge from v to w.

This is a finite process since the g-path is finite and each step does not add any vertices. On completion of this process, we will have a finite graph X that covers $\bigvee_k S^1$. Therefore $\pi_1(X) < F_k$. Since X is a finite graph, it is therefore a finite cover of $\bigvee_k S^1$, making $[F_k : \pi_1(X)] < \infty$. The path g is not a loop in X by construction, making $\pi_1(X)$ a finite index subgroup of F_k that does not contain g. Therefore F_k is residually finite. ✠

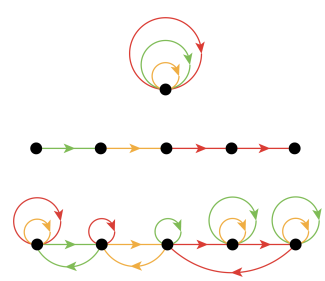

Figure 2.8.: Constructing a finite index subgroup of $\langle x, y, z\rangle$ that does not contain the element yxz^2.

In this paper we will be concerned with groups more complex than free groups, namely amalgamated free products. It is unknown what properties are necessary and sufficient to guarantee that a given amalgam is residually finite. However, there are cases in which properties of the component groups in the amalgam do guarantee that the amalgam is residually finite.

Theorem 2.3.3. *If A, B, C are finite groups with $C < A$ and $C < B$, then $A *_C B$ is residually finite.*

Proof. Consider the finite group $A \times B$. There is a natural surjective group homomorphism $\phi : A *_C B \longrightarrow A \times B$ defined by mapping $g = ca_1b_1...a_nb_n$ in its normal form to $g \mapsto (ca_1...a_n, cb_1...b_n)$. Let $K = \ker\phi$. Since $A \times B$ is finite, $[A *_C B : K] < \infty$. Since $\phi|_A : A \longrightarrow A \times B$ and $\phi|_B : B \longrightarrow A \times B$ are both injective by the construction of ϕ, this forces $A \cap K = B \cap K = \{1\}$. Since $K < A *_C B$, K acts on the Bass-Serre tree T of $A *_C B$. Recall that every vertex $v \in T$ corresponds to a coset gA (resp. gB) for some $g \in A *_C B$. So the stabilizer of such a vertex would be gAg^{-1} (resp. gBg^{-1}). Let $gag^{-1} \in gAg^{-1}$. Then $\phi(gag^{-1}) = \phi(g)\phi(a)\phi(g)^{-1} = 1$ forces $\phi(a) = 1$, which makes $a \in K$. But $A \cap K = \{1\}$, so this forces $a = 1$ and $gag^{-1} = 1$. So $gAg^{-1} \cap K = \{1\}$ for all g. An identical procedure shows that $gBg^{-1} \cap K = \{1\}$ as well. In the context of Bass-Serre theory, this tells us that K does not contain any vertex-stabilizers of T. Therefore the action of K on T is free. By Theorem 2.2.3, this makes K a finite rank free group. By Property 2.3.2.1 this makes K residually finite. The subgroup K is therefore finite index and residually finite, making $A *_C B$ residually finite by Lemma 2.3.2.

There are also examples of amalgamated products that are not residually finite, though these are of a much more exotic variety than examples like the one above. Bhattacharjee constructed the first example of a non-residually finite amalgam [15]. This amalgam surprisingly happens to be the amalgam of two free groups along a common subgroup of finite index. Other examples include lattices in the automorphism group of a product of two trees that split as twisted doubles of free groups along a finite index subgroup [16, 17]. We will discuss twisted doubles shortly.

§ 2.4. Residual Finiteness of Artin Groups

Now that we have a good amount of background knowledge built up, we will use this section to summarize the findings & methodologies that led to the new results presented in the next section. The journey begins with Craig C. Squier's paper [18] in which he proves that the triangle Artin groups $A_{2,4,4}, A_{3,3,3}$ and $A_{2,3,6}$ all split as graphs of finite rank free groups. Squier begins his proofs in [18] by mapping onto a dihedral Artin group, an Artin group with two generators, like $B = \langle a, b | abab = baba \rangle$.

This is precisely the dihedral Artin group that he maps

$$A_{2,4,4} = \langle a, b, c | ac = ca, abab = baba, bcbc = cbcb \rangle$$

onto via the map $\phi : A_{2,4,4} \longrightarrow B$ defined by $a \mapsto a, b \mapsto b, c \mapsto 1$. He designed this map so as to make checking that ϕ is a group homomorphism a simple exercise. Also, the kernel can be immediately seen to be $K = \langle\langle c \rangle\rangle$. He uses this K to create a specific group presentation for $A_{2,4,4}$ which reveals a semidirect product structure.

To create the new presentation for $A_{2,4,4}$, he begins by defining a new presentation for B, writing $B = \langle a, \pi | a\pi^2 a^{-1} = \pi^2 \rangle$ by identifying $\pi = ab$. Rewriting B in this way gives B a familiar structure, that of an HNN extension. Indeed, this new presentation can be written as $B \cong \langle \pi \rangle *_{\langle \pi^2 \rangle}$ with both injections being the inclusion map. How does splitting B tell us anything about $A_{2,4,4}$? The answer comes from the following very useful lemma.

Lemma 2.4.1. *1. Let $p : G \longrightarrow A *_C B$ be a surjective group homomorphism. Then $G \cong p^{-1}(A) *_{p^{-1}(C)} p^{-1}(B)$.*

2. Let $p : G \longrightarrow A *_C$ be a surjective group homomorphism.

Then $G \cong p^{-1}(A) *_{p^{-1}(C)}$.

Proof. Let $H = A *_C B$ be an amalgamated product. Then H acts on its Bass-Serre tree T, and the quotient of this action is a graph with 2 vertices v_1, v_2, and one edge e. Then there are vertices $\tilde{v_1}$ and $\tilde{v_2}$ that are stabilized by A and B respectively. The edge between $\tilde{v_1}$ and $\tilde{v_2}$, $\tilde{e}$, is stabilized by C. The homomorphism $p : G \to H$ induces a group action of G on T by $\tilde{v} \mapsto p(g)\tilde{v}$ where $g \in G$ and $\tilde{v}$ is an arbitrary vertex in T. Therefore $\tilde{v_1}$, $\tilde{e}$ and $\tilde{v_2}$ are stabilized by $p^{-1}(A)$, $p^{-1}(B)$ and $p^{-1}(C)$, respectively, under this action. So T/G is the graph of groups Γ with 2 vertices and one edge. The vertex groups of Γ are $p^{-1}(A)$ and $p^{-1}(B)$, and the edge group is $p^{-1}(C)$, making $p^{-1}(H) \cong G \cong p^{-1}(A) *_{p^{-1}(C)} p^{-1}(B)$. Similarly for $p : G \to A*_C$. ▓

We can now apply Lemma 2.4.1 to the map $\phi : A_{2,4,4} \to B$. Since B splits as an HNN extension and ϕ is surjective, we get that $A_{2,4,4}$ splits as an HNN extension as well. Using the kernel K and an infinitely generated presentation for $A_{2,4,4}$ that he uses to elucidate the semidirect product structure, he proves that $A_{2,4,4} \cong \langle c, ba\rangle *_{\langle baba,c,bcb^{-1}\rangle}$, which is an explicit description of the splitting of $A_{2,4,4}$ as an HNN extension. Here, a is the new variable included in the HNN extension and the two injective maps are defined by $\phi_1(baba) = \phi_2(baba) = baba$, $\phi_1(c) = \phi_2(c) = c$, $\phi_1(bcb^{-1}) = bcb^{-1}$ and $\phi_2(bcb^{-1}) = babacbcb^{-1}c^{-1}a^{-1}b^{-1}a^{-1}b^{-1}$.

Squier then moves on to perform a similar proof with the group

$$A_{3,3,3} = \langle a, b, c | aca = cac, aba = bab, bcb = cbc\rangle$$

In this case he chooses the dihedral Artin group $A = \langle a, b | aba = bab\rangle$. Again, we construct a map $\phi : A_{3,3,3} \to A$ by mapping $a \mapsto a, b \mapsto b, c \mapsto b^{-1}ab$. Verifying that ϕ is well-defined is as easy as checking that

$$\phi(cac) = b^{-1}abab^{-1}ab = b^{-1}babb^{-1}ab = a^2b = ab^{-1}bab = ab^{-1}aba = \phi(aca)$$

and

$$\phi(bcb) = bb^{-1}abb = ab^2 = b^{-1}bab^2 = b^{-1}abab = b^{-1}abbb^{-1}ab = \phi(cbc)$$

Once again, we focus on $K = \ker\phi = \langle\langle bcb^{-1}a^{-1}\rangle\rangle$. Squier denotes this new generating element $x = bcb^{-1}a^{-1}$ and conjugates x by every element of $A_{3,3,3}$ to obtain another infinitely generated presentation for $A_{3,3,3}$. The generating set is $\{a, b, x_n, y_n, z_n, u_n, v_n, w_n | \forall n \geq 0, n \in \mathbb{Z}\}$

where $x_0 = x$. There are numerous relations that all come from conjugating each of the generators by a and b respectively. Such conjugations result in the recurrence relations with which we define the generators for each n. This $\{x_n, y_n, z_n, u_n, v_n, w_n\}$ is a generating set for K by construction, and even makes K a free group freely generated by $\{x_n, y_n.z_n, u_n, v_n, w_n\}$. This presentation allows him to prove that T is the semidirect product of K with A.

He then splits A as an amalgamated product by proving that

$$A \cong \langle \pi = ab \rangle *_{\langle \pi^3 = \Delta^2 \rangle} \langle \Delta = aba \rangle$$

which can easily be seen by the fact that $\pi^3 = (ab)^3 = ababab = (aba)(bab) = (aba)(aba) = \Delta^2$ in A. Since A splits as an amalgamated product, we can again apply Lemma 2.4.1 to the map ϕ to obtain that $A_{3,3,3}$ indeed splits as an amalgamated product.

He then goes on to find the explicit free groups involved in the amalgam. Lemma 2.4.1 tells us that $A_{3,3,3} \cong \langle \pi, K \rangle *_{\langle \Delta^2, K \rangle} \langle \Delta, K \rangle$. It remains to show that the groups in the amalgam are free groups. The generators of K were constructed in such a way that conjugating x_0, y_0 and z_0 by Δ allows one to generate all of the rest of the generators of K, thereby making $\langle \Delta, K \rangle \cong \langle \Delta, x_0, y_0, z_0 \rangle \cong F_4$. Similarly, conjugating generators of K by Δ^2 reveals that $\langle \Delta^2, K \rangle \cong \langle \Delta^2, x_0, y_0, z_0, u_0, v_0, w_0 \rangle \cong F_7$. Unfortunately, such a simple process cannot be applied to $\langle \pi, K \rangle$. For this component of the amalgam, we have to use to a theorem of Stallings.

Theorem 2.4.2 ([19]). *Let G be a torsion-free group and $H < G$ a finite index subgroup. If H is a free group, then so is G.*

In our case we know that both A and K are torsion-free. Squier proves that $A_{3,3,3}$ is the semidirect product of A and K, and since both of the components of the semidirect product are torsion-free, this forces $A_{3,3,3}$ to be torsion-free as well. Therefore $\langle \pi, K \rangle < A_{3,3,3}$ forces $\langle \pi, K \rangle$ to also be torsion-free. We have already proven that $\langle \Delta^2, K \rangle \cong \langle \pi^3, K \rangle$ is a free group, and is clearly an index-3 subgroup of $\langle \pi, K \rangle$, so Theorem 2.4.2 tells us that $\langle \pi, K \rangle$ is a free group. But what is its rank? To calculate its rank, Squier uses the following lemma.

Lemma 2.4.3. [20] *Let F be a free group of rank r and H be a subgroup of index k in F. Then H is a free group of rank $rk - k + 1$.*

In our case we do not know the value of r, but we do know that $\langle \pi^3, K \rangle$ is an index-3 subgroup of $\langle \pi, K \rangle$ and that $\langle \pi^3, K \rangle$ is a free group of rank 7. Therefore $7 = 3r - 3 + 1 = 3r - 2$, so $r = 3$, making $\langle \pi, K \rangle \cong F_3$. So $A_{3,3,3}$ splits as $F_3 *_{F_7} F_4$.

He then uses nearly the exact same process to split

$$A_{2,3,6} = \langle a, b, c | ac = ca, aba = bab, bcbcbc = cbcbcb \rangle$$

We start by mapping $\phi : A_{2,3,6} \to A = \langle a, b | aba = bab \rangle$ where $a \mapsto a$, $b \mapsto b$ and $c \mapsto a^3$. Checking that ϕ is well-defined is again a simple exercise:

$$\phi(ac) = aa^3 = a^3a = \phi(ca)$$

and

$$\phi(bcbcbc) = ba^3ba^3ba^3 = ba^2(aba)a(aba)a^2 = ba^2(bab)a(bab)a^2 = ba(aba)bab(aba)a =$$
$$ba(bab)bab(bab)a = (bab)ab(bab)b(aba) = (aba)ab(aba)b(bab) = aba(aba)(bab)bab =$$
$$aba(bab)(aba)bab = a(bab)(aba)(bab)ab = a(aba)(bab)(aba)ab = a^2(bab)a(bab)a^2b =$$
$$a^2(aba)a(aba)a^2b = a^3ba^3ba^3b = \phi(cbcbcb)$$

The exact same reasoning, albeit with a slightly different generating set for K results in

$$A_{2,3,6} \cong \langle \pi, K \rangle *_{\langle \pi^3, K \rangle} \langle \Delta, K \rangle \cong F_3 *_{F_7} F_4$$

Squier's splittings of these three Artin groups are interesting in their own right, but it was a consequence of these results that sparked the inspiration for the new results presented in the next section. Notice that $A_{3,3,3}$ and $A_{2,3,6}$ both split as amalgamated products where the amalgamating subgroup is of finite index in each component. This allows us to construct the following short exact sequence:

$$1 \to F_7 \to A_{2,3,6} \to \mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z} \to 1$$

$\mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$ admits a standard projection map $\phi : \mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z} \to \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ given by $\phi(a_1b_1a_2b_2...a_nb_n) = (a_1a_2...a_n, b_1b_2...b_n)$. Denote the surjective map in the short exact sequence by $p : A_{2,3,6} \to \mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$. Then $\ker\phi < \mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$ and $p^{-1}(\ker\phi) < A_{2,3,6}$. The group F_7 injects into $p^{-1}(\ker\phi)$ since $p(F_7) = 1 \in \ker\phi$, so we get a new short exact sequence of groups

$$1 \to F_7 \to p^{-1}(\ker\phi) \to \ker\phi \to 1$$

We now analyze $\ker\phi$. Denote the generator of $\mathbb{Z}/3\mathbb{Z}$ by a and the generator of $\mathbb{Z}/2\mathbb{Z}$ by b. Then $\ker\phi \cong \langle\langle aba^2b, a^2bab\rangle\rangle$ since we need every element, pre-conjugation, to contain 3 as and 2 bs, and all such words of this type are conjugates of these two generating elements. To show that $\ker\phi$ is a free group, it suffices to prove that the action of $\ker\phi$ on the Bass-Serre tree T of $\mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$ is free.

Recall that the vertices of T correspond to $((\langle a\rangle * \langle b\rangle)/\langle a\rangle) \cup ((\langle a\rangle * \langle b\rangle)/\langle b\rangle)$. So each vertex-stabilizer comes in the form $g\langle a\rangle g^{-1}$ or $g\langle b\rangle g^{-1}$ for some $g \in \langle a\rangle * \langle b\rangle$. Every element ga^kg^{-1} (resp. gb^kg^{-1}) is mapped to $a^k \in \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ (resp. b^k) by ϕ, which is nontrivial in the codomain precisely when a^k (resp. b^k) is nontrivial in the domain. Therefore $g\langle a\rangle g^{-1}$ and $g\langle b\rangle g^{-1}$ must intersect $\ker\phi$ trivially. This proves that the action of $\ker\phi$ on T is free, making $\ker\phi$ a free group by Theorem 2.2.3. So by the First Isomorphism Theorem,
$(\mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z})/\ker\phi \cong \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, making $[\mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z} : \ker\phi] = 6$.

Why did we expend effort towards proving that $\ker\phi$ is a finite index free subgroup of $\mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$? Well now if we inspect the short exact sequence above, we have a free group as the rightmost component, which automatically makes the short exact sequence a split exact sequence. This is due to the fact that defining a group homomorphism “backwards" along p is now as simple as choosing a preimage element for each element in $\ker\phi$. Such a map is well-defined because free groups have no relations. Since the short exact sequence is split exact, this tells us that $p^{-1}(\ker\phi) \cong F_7 \times \ker\phi$, the direct product of free groups.

Lemma 2.4.4. *Let G and H be residually finite groups. Then $G \times H$ is residually finite.*

Proof. Let $(g, h) \in G \times H$ be nontrivial. Consider the natural projection maps $p_G : G \times H \longrightarrow G$ and $p_H : G \times H \longrightarrow H$. Assume $g \neq 1_G$. Then $p_G(g, h) = g$ is a nontrivial element of G. The group G is residually finite, so there exists a finite group F_g and surjective homomorphism $\phi_g : G \longrightarrow F_g$ such that $\phi_g(g) \neq 1_{F_g}$. So $f = \phi_g \circ p_G$ is a surjective group homomorphism to a finite group such that $f(g, h)$ is nontrivial. Assume $g = 1_G$. Then $h \neq 1_H$ and $f' = \phi_h \circ p_H$ is the surjective homomorphism to a finite group such that $f'(g, h)$ is nontrivial, since H is also residually finite. ■

So $p^{-1}(\ker \phi) \cong F_7 \times \ker \phi$ is the direct product of residually finite groups, making $p^{-1}(\ker \phi)$ residually finite. Furthermore, since p is a surjection, we get that $[A_{2,3,6} : p^{-1}(\ker \phi)] = [\mathbb{Z}/3\mathbb{Z} * \mathbb{Z}/2\mathbb{Z} : \ker \phi] = 6$. So $p^{-1}(\ker \phi)$ is a finite index residually finite subgroup of $A_{2,3,6}$, making $A_{2,3,6}$ residually finite by Lemma 2.3.2. The same holds for $A_{3,3,3}$.

The subclass $\{A_{2,4,4}, A_{3,3,3}, A_{2,3,6}\}$ is a special subclass of triangle Artin groups. These are all of the Euclidean triangle Artin groups.

Definition 2.4.1.

- A spherical triangle Artin group $A_{M,N,P}$ is a triangle Artin group satisfying $\frac{1}{M} + \frac{1}{N} + \frac{1}{P} > 1$.
- A Euclidean triangle Artin group $A_{M,N,P}$ is a triangle Artin group satisfying $\frac{1}{M} + \frac{1}{N} + \frac{1}{P} = 1$.
- A hyperbolic triangle Artin group $A_{M,N,P}$ is a triangle Artin group satisfying $\frac{1}{M} + \frac{1}{N} + \frac{1}{P} < 1$.

We refer to such triangle Artin groups as "Euclidean" due to the fact that the triangle with angles $\frac{2\pi}{M}$, $\frac{2\pi}{N}$ and $\frac{2\pi}{N}$ is a triangle in Euclidean space. The Coxeter group associated to a Euclidean Artin group can therefore tile a Euclidean plane by taking said triangle as the fundamental domain and assigning each Coxeter generator to be the reflection across the corresponding edge of the triangle.

Definition 2.4.1 allows us to fully classify every triangle Artin group based on the type of triangular tiling that results from such a process. Squier proved that Euclidean triangle Artin

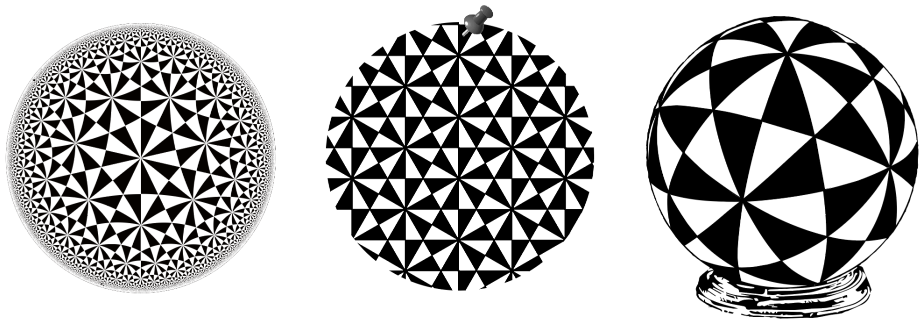

$(\frac{\pi}{2}, \frac{\pi}{3}, \frac{\pi}{5})$-triangle tiling $(\frac{\pi}{2}, \frac{\pi}{3}, \frac{\pi}{6})$-triangle tiling $(\frac{\pi}{2}, \frac{\pi}{3}, \frac{\pi}{7})$-triangle tiling

Figure 2.9.: Triangular tilings of $\mathbb{H}^2$, $\mathbb{R}^2$ & S^2.

groups split as graphs of groups. It is natural to wonder whether the same holds for spherical triangle Artin groups and hyperbolic triangle Artin groups.

We begin with spherical triangle Artin groups. Jankiewicz proved the following:

Theorem 2.4.5 ([1]). *If A is a spherical irreducible Artin group that splits as a graph of free groups, then A is dihedral or $\mathbb{Z}$.*

Proof. Every dihedral Artin group has the standard presentation

$$A_M = \langle a, b | (a, b)_M = (b, a)_M \rangle$$

When $M = 2m$, set $x = ab$ to get that $A_M \cong \langle a, x | ax^m a^{-1} = x^m \rangle \cong \langle x \rangle *_{\langle x^m \rangle} \cong \mathbb{Z}*_{\mathbb{Z}}$, which is a splitting as a graph of free groups. When $M = 2m + 1$, we set $x = ab$ and $y = (a, b)_M$ to get that $A_M = \langle x, y | x^m = y^2 \rangle \cong \langle x \rangle *_{\langle x^m \rangle \cong \langle y^2 \rangle} \langle y \rangle \cong \mathbb{Z} *_{\mathbb{Z}} \mathbb{Z}$, which is again a splitting as a graph of free groups.

Conversely, assume that an irreducible spherical Artin group A splits as a nontrivial graph of free groups. Then we can consider this graph of free groups as the fundamental group of a graph of spaces X, in which every space is a graph. The geometric realization of X is therefore a collection of graphs (corresponding to the vertex-spaces) with 2-cells used to join these graphs. These 2-cells form cylinders between the loops identified by the edge-space embeddings. Let $\tilde{X}$

be the universal cover of X. Then $\tilde{X}$ is a 2-dimensional CW complex, meaning that it has the cellular chain complex

$$\dots \xrightarrow{\partial_3} C_2(\tilde{X}) \xrightarrow{\partial_2} C_1(\tilde{X}) \xrightarrow{\partial_1} C_0(\tilde{X}) \to 0 \tag{2.4.1}$$

We would like to calculate $H_2(\tilde{X}) \cong \frac{\ker \partial_2}{Im\partial_3}$. Because $\tilde{X}$ is 2-dimensional, we know that $C_k(\tilde{X}) = 0$ for $k \geq 3$ and thus $Im(\partial_3) = 0$. By the construction of $\tilde{X}$, every 2-cell has the boundary $\ell_1 t \ell_2 \bar{t}$ where ℓ_1 and ℓ_2 are nontrivial loops in the vertex-graphs in the graph of spaces that were identified via the 2-cells. The loops ℓ_1 and ℓ_2 unfold into open paths in the universal cover $\tilde{X}$. We use t and $\bar{t}$ to denote preimages of the edge used to identify the basepoints in the graph of spaces. Therefore $\ell_1 t \ell_2 \bar{t}$ is a nontrivial loop in the 1-skeleton of $\tilde{X}$. Since this 2-cell was arbitrary, this proves that $\ker(\partial_2) = 0$ and thus $H_2(\tilde{X}) = 0$.

By Hurewicz Theorem [21], $\tilde{X}$ being a universal cover and thus simply connected, forces the Hurewicz homomorphism $h_* : \pi_2(\tilde{X}) \to H_2(\tilde{X})$ to be an isomorphism. Therefore $\pi_2(\tilde{X}) = 0$ and we can iterate this process to get that $\pi_n(\tilde{X}) = 0$ for all $n \geq 0$.

Consider the Serre fibration $\tilde{X} \to X$. Every fiber F is discrete, making $\pi_n(F) = 0$ for all $n > 0$. Putting this information into the long exact sequence induced by the fibration forces $\pi_n(X) \cong \pi_n(\tilde{X}) = 0$ for all $n > 1$. This makes X aspherical and a $K(A, 1)$ space by definition. Therefore, by $H^n(X, \mathbb{Z}) = H^n(A, \mathbb{Z})$ for all $n > 0$. Since X is a 2-dimensional CW-complex, $H^n(X) = 0$ when $n > \dim X = 2$. Therefore the cohomological dimension of A is at most 2. But spherical Artin groups are known to have cohomological dimension equal to the number of generators in their standard presentation [22]. Therefore A is forced to have at most 2 generators, which makes A either dihedral or $\mathbb{Z}$.

Now that we have successfully split every spherical Artin group that can be split as a graph of free groups, it remains to split the hyperbolic triangle Artin groups. Jankiewicz successfully split a significant portion of the hyperbolic triangle Artin groups by working with a different presentation for Artin groups called the Brady-McCammond presentation.

Definition 2.4.2 ([23])**.** For $M, N, P \geq 3$,
$A_{M,N,P} \cong \langle a, b, c, x, y, z | x = ab, y = bc, z = ca, r_M(a, b, x), r_N(b, c, y), r_P(c, a, z)\rangle$ where:

- $r_M(a, b, x) \implies x^m = bx^{m-1}a$ when $M = 2m$ and
- $r_M(a, b, x) \implies x^m a = bx^m$ for $M = 2m + 1$

and similarly for $r_N(b, c, y)$ and $r_P(c, a, z)$.

The Brady-McCammond presentation originated in [23], in which they prove that triangle Artin groups of large type (meaning that $M, N, P \geq 3$) are biautomatic through the use of the presentation complexes associated to these presentations. Furthermore, they showed that these presentation complexes are piecewise Euclidean and CAT(0).

This presentation comes from fixing an orientation on the variant of the Coxeter diagram in which we allow edges to be labeled with 2s and 3s and remove edges labeled with ∞. We start by partially orienting the Coxeter diagram Γ by orienting every edge with label ≥ 3. An example is shown in Figure 2.10.

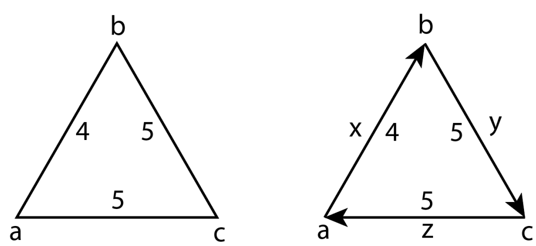

Figure 2.10.: Coxeter diagram & oriented Coxeter diagram for $A_{4,5,5}$.

Definition 2.4.3 ([2])**.** Let Γ be a simple graph (no length-1 loops or multiedges) with a partial orientation ι where $\iota(e)$ is the terminal vertex of the oriented edge e. A path γ of length ≥ 2 in Γ is called misdirected if the partial orientation on γ induced by ι can be extended to an orientation such that a maximal directed subpath of γ has length 1. A cycle γ is called almost misdirected if ignoring one edge of γ makes the resulting path misdirected.

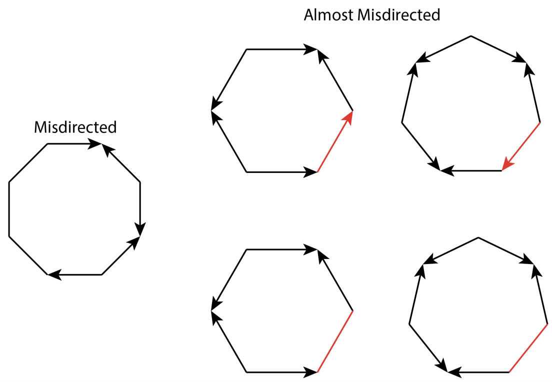

Figure 2.11.: Example of a misdirected cycle and almost misdirected cycles with subpaths of length 3 & 2 respectively

Definition 2.4.4 ([2])**.** A partial orientation ι on a simple graph Γ labelled with integers ≥ 2 is called admissible if the only oriented edges are those with labels ≥ 3 and it does not contain any almost misdirected cycles.

To obtain the Brady-McCammond presentation from the Coxeter diagram Γ of A_Γ, we start by fixing an admissible partial orientation ι on Γ. This ι allows us to define a new presentation $A_\Gamma \cong \langle a \in V(\Gamma), x \in E(\Gamma) | x = ab, r_{M_{ab}}(a, b, x)$ where $x = \{a, b\}$ and either $a = \iota(x)$ or $M_{ab} = 2\rangle$ where M_{ab} is the label on the edge between a and b. Here, ι is utilized to determine whether $x = ab$ or $x = ba$. Of course, if $M_{ab} = 2$, then $ab = ba$ and such a ι-based choice is irrelevant. When Γ is a triangle, we can choose ι so that the resulting presentation is precisely the Brady-McCammond presentation as previously described. In general, this presentation coming from an admissible partial orientation results in the following theorem.

Theorem 2.4.6 ([2])**.** *If Γ is bipartite with all labels even, then $A_\Gamma \cong A*_B$ where A and B are finite-rank free groups. Otherwise, $A_\Gamma \cong A *_C B$ where A, B, C are finite rank free groups with $rk(A) = |E(\Gamma)|$, $rk(B) = 1 - |V(\Gamma)| + 2|E(\Gamma)|$ and C is an index-2 subgroup of B with $rk(C) = 1 - 2|V(\Gamma)| + 4|E(\Gamma)|$.*

To prove this theorem, we fix the ι-induced presentation for A_Γ and construct the presentation complex of A_Γ with respect to this presentation.

Definition 2.4.5 ([23])**.** Given a group G with presentation $\langle S|R\rangle$, the presentation complex X_G is the 2 dimensional cell complex obtained in the following manner. First construct the wedge of oriented circles with one circle for each $s \in S$. Then, for each word $r \in R$, construct a 2-cell whose boundary is labelled with the word r, and attach each oriented edge to the corresponding loop in the wedge of circles.

By construction, $\pi_1(X_G) = \langle S|R\rangle$. Clearly every loop in X_G is homotopic to a combination of loops in the wedge of circles, thereby making the generators of $\pi_1(X_G) = S$. The 2-cells allow us to contract any loop on the boundary of a 2-cell through the interior of the 2-cell, thereby making every relation in $\pi_1(X_G)$ a word in R.

We now can use the presentation complex associated to the ι-induced presentation of A_Γ to arrive at our desired splitting. The case in which Γ is bipartite with even labels is not relevant to this paper since we will be focused on triangle Artin groups and Γ being a triangle precludes it from being bipartite. When Γ is not a bipartite graph with all even labels, the ι-induced presentation for A_Γ allows us to construct the relator polygons in an intentionally illuminating way. Notice that the yellow, red and orange edges in the relator polygons in Figure 2.12 only appear as vertical path-components of the boundaries, and that the green, purple & blue elements appear only in the tops and bottoms of our relator polygons, oriented horizontally. This allows us to apply Seifert Van Kampen's theorem to the complex in a natural way, as shown in the Figure 2.12.

We can then deformation retract the blue component U, the green component V, and the turquoise component $U \cap V$ to the graphs $\bar{U}$, $\bar{V}$ and $\overline{U \cap V}$, resulting in $A_\Gamma \cong \pi_1(\bar{U}) *_{\pi_1(\overline{U \cap V})} \pi_1(\bar{V})$. Since $\bar{U}$, $\bar{V}$ and $\overline{U \cap V}$ are all graphs, their respective fundamental groups are all free groups. It remains to calculate the ranks of these free groups.

Each portion of U in a relator polygon deformation retracts naturally to the horizontal paths at the nearest top/bottom of that relator polygon. Since all of the vertices in the relator polygons are all identified in the presentation complex, we must also identify the vertices in the deformation

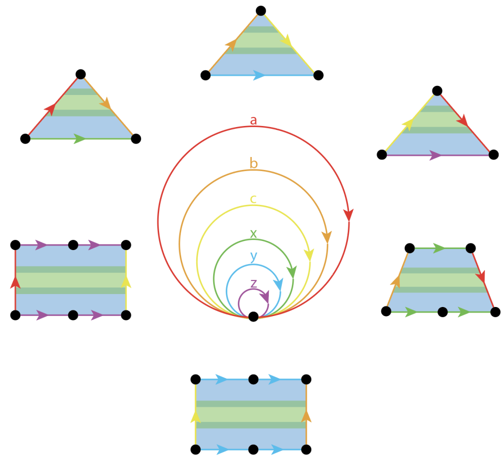

Figure 2.12.: Seifert-Van Kampen's Theorem applied to the presentation complex for $A_{4,5,5}$

retract $\bar{U}$. This will result in the wedge of many length-one loops with potentially multiple loops of the same color. Having multiple loops of the same color is redundant, so we identify these loops to realize $\bar{U}$ as the wedge of n circles, one for each color present at the tops/bottoms of the relator polygons. For example, in the case of $A_{4,5,5}$, these loops correspond to the generators x, y, z, which are precisely the generators constructed from the oriented edges of the Coxeter diagram Γ. By the construction of the Brady-McCammond complex, it will always be the case that the loops in $\bar{U}$ are in direct correspondence with these new generators added during the construction of the Brady-McCammond presentation. Therefore $\pi_1(\bar{U})$ is the free group on $E(\Gamma)$, making $rk(\pi_1(\bar{U})) = |E(\Gamma)|$.

When we deformation retract V by compressing each portion of V to its midline, we get a graph with an edge coming from each relator polygon. Notice that the relations in a Brady-McCammond presentation come in pairs: 2 relations for each edge in the Coxeter diagram Γ. One of these relations is used to define the new generator coming from an oriented edge in Γ, and the other relation is an analogue of the traditional Artin relation involving the pair of traditional Artin-generators used to define the new orientation-induced generator. Therefore $|E(\bar{V})| = 2|E(\Gamma)|$. The endpoints of an edge in $\bar{V}$ are vertices on a “vertical" edge of the relator polygon's boundary. These vertical edges correspond to the generators coming from the usual Artin presentation. These generators are represented in the Coxeter diagram Γ as vertices. Therefore $V(\bar{V}) = V(\Gamma)$ and $rk(\pi_1(\bar{V})) = 1 - |V(\bar{V})| + |E(\bar{V})| = 1 - |V(\Gamma)| + 2|E(\Gamma)|$.

In Figure 2.13 we can see that every edge in $\bar{V}$ is double-covered by the two edges in $\overline{U \cap V}$ above and below it in the relator polygon, making $\overline{U \cap V}$ a double cover of $\bar{V}$ and $[\pi_1(\bar{V}) : \pi_1(\overline{U \cap V})] = 2$.

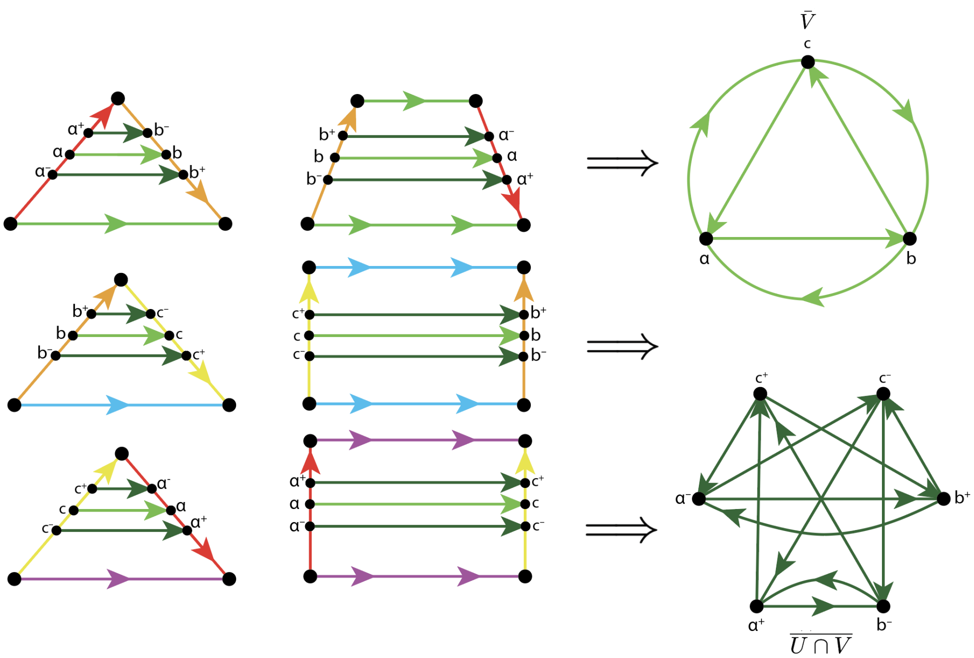

Figure 2.13.: $\overline{U \cap V} \looparrowright \bar{V}$ for $A_{4,5,5}$

We now use Lemma 2.4.3 by setting $F = \pi_1(\bar{V})$ and $H = \pi_1(\overline{U \cap V})$, which forces $r = 1 - |V(\Gamma)| + 2|E(\Gamma)|$ and $k = 2$. Putting this all together gives us that $rk(\pi_1(U \cap V)) = (1 - |V(\Gamma)| + 2|E(\Gamma)|)(2) - 2 + 1 = 1 - 2|V(\Gamma)| + 4|E(\Gamma)|$.

Having calculated the groups that arise from the application of Seifert-Van Kampen's Theorem to the Brady-McCammond complex, it remains to show that the induced maps from $\pi_1(\overline{U \cap V})$ to $\pi_1(\bar{U})$ and $\pi_1(\bar{V})$ are injective. To do so, we must define an important class of combinatorial maps that will play a major role throughout the rest of the paper.

Definition 2.4.6 ([19])**.** A combinatorial map $Y \longrightarrow X$ between graphs Y and X is a function that maps every vertex to a vertex and every edge to an edge. A combinatorial immersion $\phi : Y \looparrowright X$ is a locally injective combinatorial map.

Combinatorial immersions are important because every combinatorial immersion $\phi : Y \looparrowright X$ induces an injective homomorphism $\pi_1(Y, y) \hookrightarrow \pi_1(X, x)$ [19]. Equivalently, the existence of a combinatorial immersion $\phi : Y \looparrowright X$ guarantees that Y can be completed to a cover of X by adding trees to the vertices in Y that are keeping Y from being a cover of X. The copy of $\overline{U \cap V}$ in Figure 2.13 combinatorially immerses into $\bar{V}$ since it is already a cover. The homotopy equivalent copy of $\overline{U \cap V}$ that conforms to the deformation retract of $\bar{U}$ is shown in Figure 2.14. Notice that the three dashed edges get collapsed to vertices during the retraction.

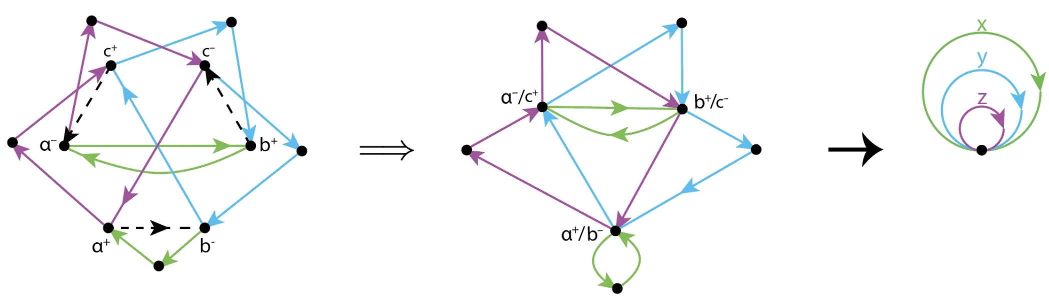

Figure 2.14.: $\overline{U \cap V} \looparrowright \bar{U}$

This graph, after the collapsing of the three dashed edges, also combinatorially immerses into $\bar{U}$, though it is not a cover. Infinite trees will have to be added to the five outer vertices for those vertices to be preimages of the lone vertex in $\bar{U}$. Nevertheless, the fact that both homotopy equivalent copies of $\overline{U \cap V}$ combinatorially immerse into $\bar{U}$ and $\bar{V}$ guarantees that the edge maps in the ensuing graph of groups will be injective. This makes our application of Seifert-Van Kampen's Theorem to the Brady-McCammond complex a splitting of each large-type Artin group into a graph of finite-rank free groups.

Once these large-type Artin groups were split as graphs of free groups, Jankiewicz uses this information to prove that the triangle Artin groups $A_{M,N,P}$ where $M, N, P \geq 4$ and $\{M, N, P\} \neq \{2m+1, 4, 4\}$ for any $m \in \mathbb{N}$ are residually finite. To do this, Jankiewicz first defines an index-2 subgroup of $A_{M,N,P} \cong A *_C B$ called the twisted double.

Definition 2.4.7 ([2])**.** Let A and C be finite rank free groups and $\beta : C \to C$ be an automorphism. The double of A along C twisted by β, denoted $D(A, C, \beta)$, is the amalgam $A *_C A$ where C is mapped into the leftmost A via the inclusion map ι, and is mapped into the rightmost A via $\iota \circ \beta$.

We construct a twisted double $D(A, C, \beta) \cong A *_C \hat{b}A\hat{b}^{-1}$ in $A_{M,N,P} \cong A *_C B$ by defining $\beta : C \to C$ to be $\beta(c) = \hat{b}c\hat{b}^{-1}$ where $\hat{b}$ is a nontrivial coset representative of $B/C \cong \{[1], [\hat{b}]\}$. We can realize $D(A, C, \beta)$ as the kernel of the homomorphism $\phi : A *_C B \to B/C \cong \mathbb{Z}/2\mathbb{Z}$. The kernel is generated by all conjugates of A, which are simply A and $\hat{b}A\hat{b}^{-1}$. Consider the action of $\ker\phi = \langle A, \hat{b}A\hat{b}^{-1}\rangle$ on the Bass-Serre tree T. The subgroup A fixes a vertex in T, as does $\hat{b}A\hat{b}^{-1}$. There are two edges between these vertices in T, one that is fixed by C, the other which is fixed by $\hat{b}C\hat{b}^{-1}$. Therefore, by The Fundamental Theorem of Bass-Serre Theory, $\ker\phi = A *_{C \cap \hat{b}C\hat{b}^{-1} = C} \hat{b}A\hat{b}^{-1}$ where $C \to A$ is the inclusion map and $C \to \hat{b}A\hat{b}^{-1}$ is sent to $\hat{b}C\hat{b}^{-1}$, making $D(A, C, \beta) = \ker\phi$. Since ϕ is obviously onto, this proves that $[A_{M,N,P} : D(A, C, \beta)] = 2$.

So, by Lemma 2.3.2, if we can prove that $D(A, C, \beta)$ is residually finite, then we have proven that $A *_C B$ is residually finite. To prove that $D(A, C, \beta)$ is residually finite, Jankiewicz first proves that it virtually splits as an algebraically clean graph of finite rank free groups.

Definition 2.4.8 ([24])**.** A graph of groups is algebraically clean if the vertex groups are free and each edge group is a free factor in its adjacent vertex group(s).

Theorem 2.4.7 ([24])**.** *If a group $\mathcal{G}$ splits as an algebraically clean graph of finite-rank free groups, then $\mathcal{G}$ is residually finite.*

Proof. Let (Γ, G) be a splitting of $\mathcal{G}$ as an algebraically clean graph of finite-rank free groups. We begin our calculation of $\pi_1(\Gamma, G)$ by collapsing a maximal tree in Γ, and amalgamating the identified vertex groups along their respective edge groups. Consider two vertices v and w

in Γ joined by an edge e. The edge group G_e is a free factor in both G_v and G_w, making the amalgamated product $G_v *_{G_e} G_w$ a free group. Therefore the iterated amalgam that appears as the vertex group after the maximal tree has been collapsed is a free group as well. Denote this collapsed version of (Γ, G) by $(\bar{\Gamma}, \bar{G})$. This new graph of groups has one vertex and a collection of loops based at that vertex, making $\pi_1(\bar{\Gamma}, \bar{G})$ is an iterated HNN extension of a free group. Furthermore, each edge group E_i is a free factor in the vertex group V since each G_v in (Γ, G) is a free factor in V and every E_i is a free factor in some such G_v by construction.

Our goal is to define a projection to the iterated HNN extension of a finite group for each $g \in \mathcal{G}$ such that g survives the projection. We begin by noting that iterated HNN extensions have a normal form for each element element $g = f_0 t_0^{\epsilon_0} f_1 t_1^{\epsilon_1} \dots f_r t_r^{\epsilon_r} f_{r+1}$ where all $f_i \in V$, $\epsilon_i \in \{-1, 0, 1\}$ and $t_1, \dots, t_r$ are stable letters. Furthermore, we can assume that this normal form for g is reduced, meaning that no combination of relations can decrease its length. In particular, this means that if g contains a subword $t_k f_k t_k^{-1}$ (resp. $t_k^{-1} f_k t_k$), then $f_k \notin E_k$ (resp. $f_k \notin \bar{E}_k$) where $E_k \cong \bar{E}_k$ is the edge group that contributes the stable variable t_k to the HNN extension. Consider the case when g contains a subword $t_k f_k t_k^{-1}$ (resp. $t_k^{-1} f_k t_k$) for some $f_k \notin E_k$ (resp. $f_k \notin \bar{E}_k$). In this case we can explicitly construct a finite index subgroup $L_k < V$ such that $E_k < L_k$ (resp. $\bar{E}_k < L_k$) and $f_k \notin L_k$. We construct such an L_k using the following covering space argument.

The vertex group V being a finitely generated free group means that we can represent $V \cong \pi_1(\bigvee_{i<\infty} S^1) = \pi_1(\mathcal{V})$. The edge group E_k being a finitely generated subgroup of V forces the existence of a cover $\mathcal{E}_k$ of $\mathcal{V}$ with a finite core. We then attach a path to the basepoint of $\mathcal{E}_k$ whose edge labels form the word for g in its normal form, and then perform any necessary Stallings folds to identify redundant edges (see Figure 2.18 for examples of Stallings folds). We now use the same algorithm as in the proof of Property 2.3.2.1 to complete this augmented $\mathcal{E}_k$ to a finite cover $\mathcal{L}_k$ of $\mathcal{V}$. Therefore $\pi_1(\mathcal{L}_k) = L_k$ is a finite index subgroup of V that does not contain g since g is an open path in $\mathcal{L}_k$. See Figure 2.15 for an example of such a construction. For every other f_k in the normal form of g that is not conjugated by a stable variable, assign L_k

to be the finite index subgroup of V that does not contain f_k, whose existence is guaranteed by the fact that V is residually finite.

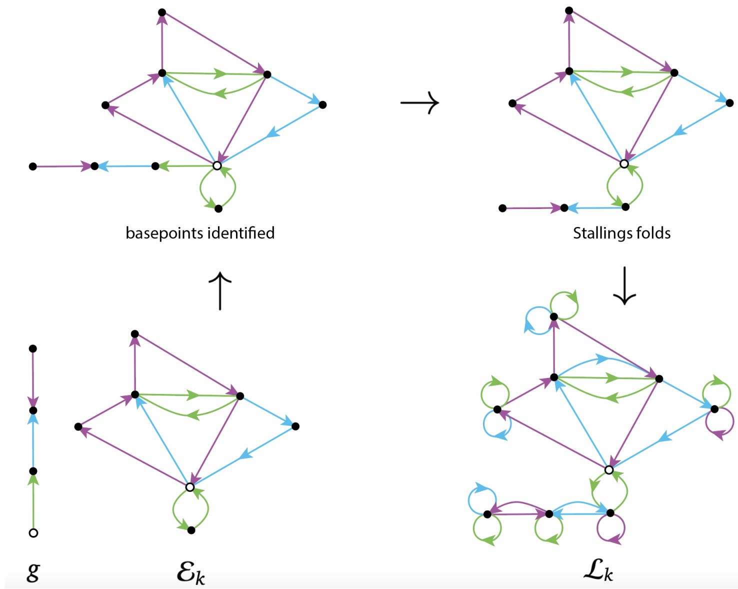

Figure 2.15.: The process of constructing a finite cover $\mathcal{L}_k$ of $S^1 \vee S^1 \vee S^1$ that contains $\mathcal{E}_k$ and g as an open path.

If $g = f_0 \neq 1$ is in normal form, then $g \in V$ and we choose N to be the finite index subgroup of V guaranteed by the residual finiteness of V to not contain g. Otherwise, let $N = \bigcap_{k=1}^{r+1} L_k$. Define $C = \bigcap_{\phi \in Aut(V)} \phi(N)$. This group C is a finite index characteristic subgroup since automorphisms preserve subgroup-index and the number of subgroups of a fixed finite index in a finitely generated group is finite. The quotient $p : V \to V/C$ extends naturally (and well-definedly) to a surjective map $\hat{p} : (...((V*_{E_1})*_{E_2})...*_{E_k}) \to (...((p(V)*_{p(E_1)})*_{p(E_2)})...*_{p(E_k)})$, which is an iterated HNN extension of a finite group since C is finite index and characteristic. If $g \in V$, then $\hat{p}(g) \neq 1$ by our choice of C. Consider $g \notin V$. Since every element f_k that is conjugated by a stable letter in g is separated from the edge group corresponding to the stable letter by L_k, none of the stable-letter-derived relations in the image can be used to reduce the length of the normal form of $\hat{p}(g)$. Therefore $\hat{p}(g)$ is an element in its reduced normal form of length $r > 1$, making $\hat{p}(g) \neq 1$. All iterated HNN extensions of finite groups are virtually free (the free group generated by the stable letters is finite index), making it residually finite by Property 2.3.2.1 and Lemma 2.3.2. Therefore $\hat{p}(g)$ will survive a further quotient to a finite group, making $\mathcal{G}$ residually finite. ∎

In order to prove that $D(A, C, \beta)$ splits as an algebraically clean graph of finite rank free groups, Jankiewicz constructs a quotient of $D(A, C, \beta)$ with respect to an oppressive set for C in A.

Definition 2.4.9 ([2]). Let $\rho : Y \to X$ be a covering map inducing the inclusion of a finite rank free group $H \cong \pi_1(Y) \to \pi_1(X) \cong G$. Let $A_\rho \subseteq G$ consist of all $g \in G$ represented by a cycle γ in X such that $\gamma = \gamma_1\gamma_2$ where

- $\gamma_1 = \rho(\mu_1)$ where μ_1 is a nontrivial simple non-closed path in Y going from the vertex y_0 to y_1.
- $\gamma_2 = \rho(\mu_2)$ and μ_2 is either trivial or a simple non-closed path in Y going from some vertex y_2 to y_0 where $y_1 \neq y_2 \neq y_0$.

We refer to A_ρ as the oppressive set for H in G with respect to ρ.

In order for a graph of groups to be algebraically clean, we need the edge groups to be free factors in the vertex groups, which requires every Y with $G_e \cong \pi_1(Y)$ to be an embedded subgraph of X where $\pi_1(X) \cong G_v$. This requires every path in Y that maps to a loop in X to have been a loop in Y as well. The oppressive set for G_e in G_v is therefore the set of all elements keeping G_e from being a free factor in G_v. If we can find a quotient ϕ of G such that the image of our oppressive set is disjoint from $\phi(G_e)$, then $\phi(G_e)$ is a free factor in $\phi(G_v)$. Finding such a quotient of $D(A, C, \beta)$ requires the use of the following lemmas.

Lemma 2.4.8. *Let $H < G$ be a free factor. Then for every finite index $G' < G$, $G' \cap H$ is a free factor in G'.*

Proof. The subgroup $H < G$ being a free factor in G, means that $G \cong H * F$ for some $F < G$. Consider the Bass-Serre tree T of this splitting. The subgroup $G' < G$ being of finite index means that the action of G' on T has a finite fundamental domain. If $G' \cap H$ is trivial, then $G' \cap H$ is trivially a free factor in G'. Assume $G' \cap H$ is nontrivial. Then we can choose the finite fundamental domain of the action of G' on T to contain the vertex stabilized by H under the action of G on T. The stabilizer of this vertex under the action of G' will therefore be $G' \cap H$. The edge stabilizers of T are trivial since our splitting is free. By The Fundamental Theorem of Bass-Serre Theory, $G' \cong \pi_1(T/G')$ with T/G' viewed as a graph of groups. Since the edge stabilizers of T are trivial, the edge groups in T/G' will be trivial. So when we amalgamate the vertex groups in T/G', we get $G' = (G' \cap H) * L$ where L is the free product of the other vertex groups in T/G' and the fundamental group of the underlying graph of T/G'. Therefore, $G' \cap H$ is a free factor in G'. ■

Lemma 2.4.9. *Let $\mathcal{A}_\rho$ be an oppressive set for C in A coming from the covering map $\rho : X_C \to X_A$. Suppose there exists a finite quotient $\Psi : D(A, C, \beta) \to K$ such that $\Psi|_A(\mathcal{A}_\rho) \cap \Psi(A) = \emptyset$. Then $D(A, C, \beta)$ virtually splits as an algebraically clean graph of finite rank free groups. In particular, $D(A, C, \beta)$ is residually finite.*

Proof. $D(A, C, \beta)$ is the fundamental group of a graph of groups and therefore inherits a natural action on its Bass-Serre tree T. The vertex-stabilizers of this action are conjugates of A, and the edge-stabilizers are conjugates of C. Since K is finite, $\ker \Psi$ is a finite index subgroup of $D(A, C, \beta)$ forcing the action of $\ker \Psi$ on T to have a finite fundamental domain. The vertex stabilizers of the action of $\ker \Psi$ are conjugates of $\ker \Psi \cap A \cong \ker \Psi|_A$, and the edge stabilizers are conjugates of $\ker \Psi|_A \cap C$. This information tells us that $T/\ker \Psi$ is the graph of groups with vertex groups isomorphic to $\ker \Psi|_A$ and edge groups isomorphic to $\ker \Psi|_A \cap C$, and thus $\ker \Psi$ splits as a graph of finite rank free groups. Our goal is to show that this splitting is algebraically clean.

Consider the edge group $\ker \Psi|_A \cap C$. Let $C' = \Psi|_A^{-1}(\Psi|_A(C))$. Since $\Psi|_A(C) \cap \Psi|_A(\mathcal{A}_\rho) = \emptyset$, $C' \cap \mathcal{A}_\rho = \emptyset$. Let $\hat{C}$ be the cover of X_A with $\pi_1(\hat{C}) = C'$. Since $C < C'$, there exists a cover $\hat{\rho} : X_C \to \hat{C}$ and ρ factors through $\hat{\rho}$. The oppressive set $\mathcal{A}_{\hat{\rho}} \subseteq \mathcal{A}_\rho \cap C' = \emptyset$ because every path in X_C that maps to a loop in $\hat{C}$ must therefore also map to a loop in X_A, since $\hat{C}$ is a cover of X_A. This makes $\hat{\rho}$ an embedding and C a free factor in C'. By Lemma 2.4.8, since $\ker \Psi|_A < C'$, $C \cap \ker \Psi|_A$ is a free factor in $\ker \Psi|_A$. Conjugates of the intersection $C \cap \ker \Psi|_A$ are the edge groups in our splitting, making the graph of groups Γ with $\pi_1(\Gamma) \cong \ker \Psi$ algebraically clean. So $\ker \Psi$ is a finite index residually finite subgroup of $D(A, C, \beta)$ by Theorem 2.4.7, making $D(A, C, \beta)$ residually finite by Lemma 2.3.2. ⧗

We now would like to use the above lemmas to show that our specific twisted double is residually finite. In our case we will be constructing a quotient of A, so we utilize the following theorem to give us the necessary ingredients to extend a quotient of A to a separating quotient of $D(A, C, \beta)$.

Theorem 2.4.10 ([2]). *Suppose there exists a quotient $\phi : A \to \bar{A}$ with the following characteristics:*

1. *$\bar{A}$ is a virtually special hyperbolic group,*
2. *$\bar{C} := \phi(C)$ is malnormal and quasiconvex in $\bar{A}$,*
3. *ϕ separates C from an oppressive set $\mathcal{A}$ of C in A,*
4. *β projects to an automorphism $\bar{\beta} : \bar{C} \to \bar{C}$.*

Then $D(A, C, \beta)$ virtually splits as an algebraically clean graph of finite rank free groups. In particular, $D(A, C, \beta)$ is residually finite.

Condition (4) allows us to extend ϕ to a projection $\Phi : D(A, C, \beta) \to D(\bar{A}, \bar{C}, \bar{\beta})$. Conditions (1)-(3) allow us to construct a quotient $\Psi : D(\bar{A}, \bar{C}, \bar{\beta}) \to K$ for some finite group K so that $\Psi \circ \Phi : D(A, C, \beta) \to K$ separates $\mathcal{A}$ from C, giving us that $D(A, C, \beta)$ is residually finite. We now set out to find such a quotient ϕ. The most difficult condition for us to satisfy in Theorem 2.4.10 is condition (2).

Definition 2.4.10. A subgroup $C < G$ is called malnormal if $gCg^{-1} \cap C = \{1\}$ for all $g \in G$.

Thus far we have only discussed our vertex groups and subgroups of $A_{M,N,P}$ in terms of the fundamental group of graphs. In order to explore malnormality in this fashion, we need to understand what subgroup intersection and conjugation look like in terms of graphs.

Definition 2.4.11. Let $\phi_i : Y_i \to X$ be a combinatorial immersion for $i = 1, 2$. The fiber product of Y_1 and Y_2 over X is the graph $Y_1 \otimes_X Y_2$ with vertex set
$\{(v_1, v_2) \in V(Y_1) \times V(Y_2) : \phi_1(v_1) = \phi_2(v_2)\}$
and edge set $\{(e_1, e_2) \in E(Y_1) \times E(Y_2) : \phi_1(e_1) = \phi_2(e_2)\}$.
There is a natural combinatorial immersion $Y_1 \otimes_X Y_2 \to X$, given by $(y_1, y_2) \mapsto \phi_1(y_1) = \phi_2(y_2)$.

Lemma 2.4.11 ([19]). *Let $H_1, H_2 \leq G \cong \pi_1(X, v)$ where X is a finite graph. For $i = 1, 2$, let $(Y_i, \hat{x}_i) \to (X, v)$ be a cover of X where $\pi_1(Y_i, \hat{x}_i) \cong H_i$. Then $H_1 \cap H_2 \cong \pi_1(Y_1 \otimes_X Y_2, (\hat{x}_1, \hat{x}_2))$.*

In our case, the graph X_A that we will be computing fiber products over is the wedge of 3 circles. Since X_A has only one vertex, the immersions ϕ_1 and ϕ_2 agree on every vertex. So the vertex set of every fiber product for the rest of this paper will be $V(Y_1) \times V(Y_2)$. For example, when we split $A_{4,5,5}$ using Seifert-Van Kampen's Theorem applied to the deformation retract of the Brady McCammond presentation in Figure 2.13, we get $A *_C B$ where $A = \pi_1(\bar{U}), B = \pi_1(\bar{V})$ and $C = \pi_1(\overline{U \cap V})$. If we denote X_C as the copy of $\overline{U \cap V}$ that is combinatorially immersed into $\bar{U}$ then the computation $X_C \otimes_{\bar{U}} X_C$ results in the connected components shown in Figure 2.16.

Notice that every element of $V(X_C) \times V(X_C)$ is present in $X_C \otimes_{\bar{U}} X_C$. There is a purple edge $(a, b) \to (c, d)$ only when there is a purple edge $a \to c$ and a purple edge $b \to d$ in the respective components of the fiber product. The same is true for blue and green edges. For example, we have a purple edge $(1, 2) \to (2, 4)$ in the bottom connected component since there is a purple edge $1 \to 2$ in the first copy of X_C and a purple edge $2 \to 4$ in the second copy of X_C. This example also demonstrates that fiber products need not be connected, making choice of basepoint very important when calculating the fundamental group. We will be primarily interested in cores of graphs, and thus will omit the trees needed to complete the connected

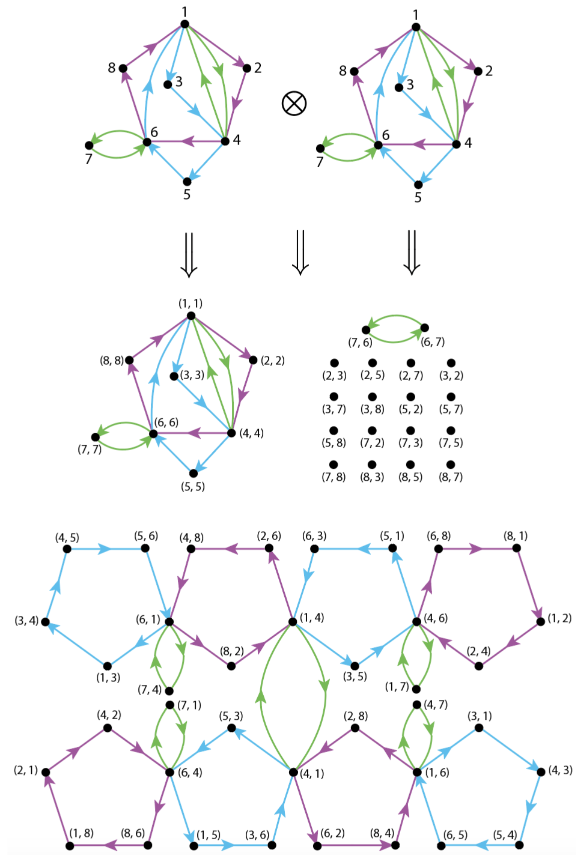

Figure 2.16.: $X_C \otimes_{\bar{U}} X_C$

components of fiber products to actual covers, as well as any maximal acyclic paths that arise during fiber product computations.

When we conjugate $C < A$ by an element $g \in A$, this naturally corresponds to translating the basepoint b along the path labelled by g^{-1} in X_C to the vertex b_g. Such a path is incident with every vertex in X_C since X_C is a cover of X_A. We will see that $gCg^{-1} \cong \pi_1(X_C, b_g)$. Figure 2.17 demonstrates such a basepoint translation pictorially.

In Figure 2.17 we have initially chosen the basepoint of X_C to be the vertex labeled 1. If we conjugate $\pi_1(X_C, 1)$ by z (the algebraic element corresponding to traveling once around the purple loop in A), this corresponds to shifting the basepoint backwards along the purple edge to

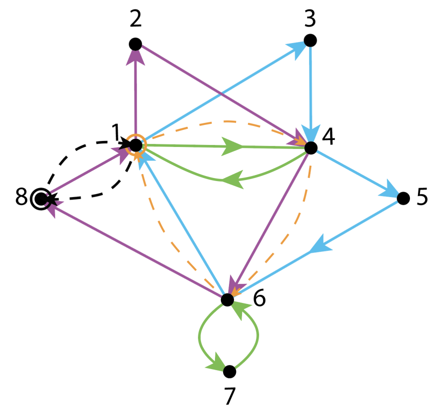

Figure 2.17.: Change of basepoint

the vertex labeled 8. This is because every loop that begins at 8 can be realized by first traversing the purple edge to 1, performing a loop at 1, and then returning backwards along the same purple edge back to 8. In the figure, a loop at 1 is shown in orange and the traversal of the purple path is shown in black. This pictorially demonstrates that $z\pi_1(X_C, 1)z^{-1} \cong \pi_1(X_C, 8)$. We may also translate the basepoint along a path that leaves the core, i.e. translating the basepoint 1 along a blue edge in the tree (not shown in the figure) attached to 1. In both cases, translating the basepoint does not affect the structure of the underlying graph itself.

Therefore if we want to find a quotient ϕ such that $\phi(C)$ is malnormal in $\phi(A)$, such a quotient would have to introduce 2-cells to the graphs that arise as connected components of fiber products so as to make them all contractible, thereby making $\phi(C^g \cap C) \cong \{1\}$ as desired.

Theorem 2.4.12 ([2])**.** *Suppose $M, N, P \geq 4$ with at least one of M, N, P even. Then $A_{M,N,P} \cong A *_C B$ for A, B, C finite rank free groups with $[B : C] = 2$ and $A \cong \langle x, y, z \rangle$. There exists a quotient $\phi : A \longrightarrow \bar{A}$ defined in the following way. Let $\bar{k} = \begin{cases} \frac{k}{2} & k \text{ even} \\ k & k \text{ odd} \end{cases}$ and $\bar{A} = \langle x, y, z | x^{\bar{M}}, y^{\bar{N}}, z^{\bar{P}} \rangle \cong \mathbb{Z}/\bar{M}\mathbb{Z} * \mathbb{Z}/\bar{N}\mathbb{Z} * \mathbb{Z}/\bar{P}\mathbb{Z}$ with the associated quotient $\phi : A \longrightarrow \bar{A}$. The map ϕ is a projection that satisfies the four conditions in Theorem 2.4.10.*

The existence of such a quotient ϕ proves that $D(A, C, \beta)$ is residually finite by Theorem 2.4.10, and thus $A_{M,N,P}$ is residually finite as well by Lemma 2.3.2. This same method was used to prove that $A_{M,N,P}$ is residually finite when $(M, N, P) = (2m + 1, 2n + 1, 2p + 1)$ with $m, n, p \geq 2$, albeit with a different quotient, owing to the fact that the graphs representing $C^g \cap C$ contain non-monochrome simple loops in those cases.

But what about $A_{2,N,P}$? When $M = 2$, the fourth relation in the Brady-McCammond presentation becomes $x = ba$, which results in a fourth triangle being present in the presentation complex. When we deformation retract $U \cap V$ along the deformation retract of U, the vertices b^+, c^-, a^- and c^+ are all identified to one vertex. This four-fold identification forces multiple closed Stallings folds to occur.

Stallings folds are an important part of the process of converting a graph into a cover of another graph. In the case of $\bar{U}$, the deformation retract of U, there is only one vertex, and thus any vertex in a cover of $\bar{U}$ must locally look identical to this lone vertex. As discussed before, if a vertex in a combinatorially immersed graph does not have the same edges adjacent to it as the lone vertex in $\bar{U}$, we can simply attach an infinite tree to this vertex so that every vertex in the tree has the same adjacent edges as the lone vertex in $\bar{U}$. But we also must consider the case in which there are redundant edges entering or exiting a vertex. Stallings folds remedy precisely these redundancies. Stallings folds come in two varieties, open and closed, as seen in Figure 2.18.

The Stallings fold at the top of Figure 2.18 is an example of an open Stallings fold. The other two Stallings folds are examples of closed Stallings folds. A closed Stallings fold is a Stallings fold that results in an entire loop being collapsed in the graph. In the context of constructing amalgamated products, closed Stallings folds are to be avoided at all costs since they result in the loss of a nontrivial loop, thereby creating a kernel in the induced map on the fundamental groups of the graphs. If closed Stallings folds were to occur during the course of a deformation retract like in the previous examples, then Seifert-Van Kampen's theorem applied to the Brady-

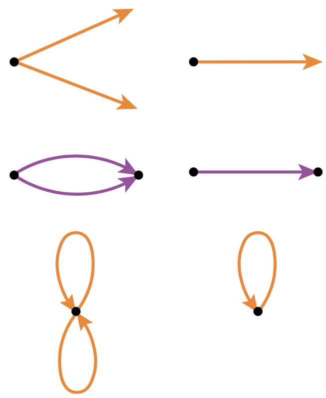

Figure 2.18.: Stallings Folds

McCammond complex cannot be an amalgamated free product. To remedy the occurrence of closed Stallings folds when the previous procedures are applied to $A_{2,N,P}$, Jankiewicz defines the new presentation for $A_{2,N,P}$ described below.

Lemma 2.4.13 ([1]). *$A_{2,N,P} \cong \langle b, x, y | r_N(b, x), r_P(b, y), bx^{-1}yb^{-1} = yx^{-1}\rangle$ where*

- $r_N(b, x) \implies bx^n b^{-1} = x^n$ *when* $N = 2n$
- $r_N(b, x) \implies bx^n b = x^{n+1}$ *when* $N = 2n + 1$

and similarly for $r_P(b, y)$.

To prove that this is indeed a valid presentation for $A_{2,N,P}$ one only needs to identify $x = ab$ and $y = cb$. The presentation complex of $A_{2,N,P}$ where $N > 3$ and one of N, P is odd is shown in Figure 2.19.

We can again apply Seifert Van-Kampen's theorem to this presentation complex, and then deformation retract U, V and $U \cap V$ to obtain the graphs in Figure 2.20.

The general form of these graphs can be used to prove that $A_{2,N,P} \cong F_3 *_{F_7} F_4$ when $N > 3$ and one of N, P is odd. An example of what the presentation complex of $A_{2,N,P}$ when $N > 3$ and both N and P are even is shown in Figure 2.21.

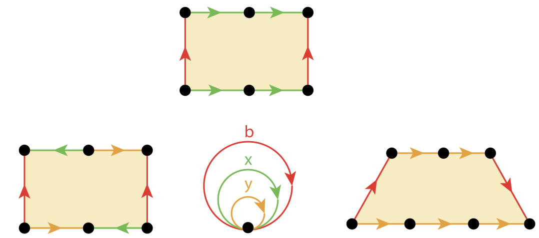

Figure 2.19.: Presentation complex for
$A_{2,4,5} = \langle b, x, y | bx^2b^{-1} = x^2, by^2b = y^3, bx^{-1}yb^{-1} = yx^{-1} \rangle$

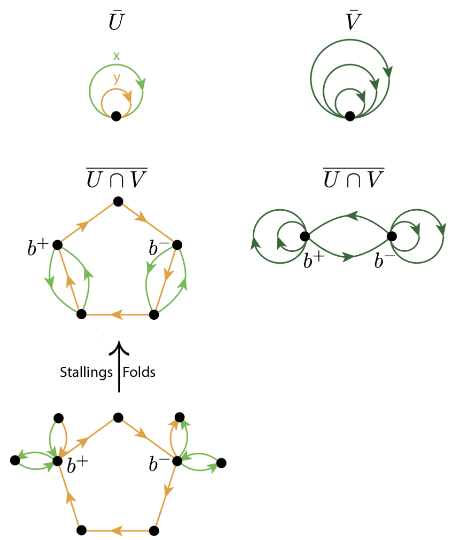

Figure 2.20.: Deformation retracts of U, V and $U \cap V$

In Figure 2.21 we attempt to apply Seifert Van-Kampen's Theorem in the same way as in the previous examples, but this unfortunately cannot be done since the blue regions are no longer connected. Instead, we can forego using Seifert Van-Kampen's Theoreom by instead noticing that the presentation $A_{2,2n,2p} = \langle b, x, y | bx^nb^{-1} = x^n, by^pb^{-1} = y^p, bx^{-1}yb^{-1} = yx^{-1} \rangle$ is precisely that of an HNN extension with stable letter b. So $A_{2,2n,2p} \cong \langle x, y \rangle *_{\langle x^n, y^p, \alpha \rangle}$ where the two injective homomorphisms $\phi_1 : \langle x^n, y^p, \alpha \rangle \longrightarrow \langle x, y \rangle$ and $\phi_2 : \langle x^n, y^p, \alpha \rangle \longrightarrow \langle x, y \rangle$ are defined by $\phi_1(\alpha) = x^{-1}y$ and $\phi_2(\alpha) = yx^{-1}$.

But when we impose Jankiewicz's presentation on $A_{2,3,2m}$ with $2m \geq 6$, a closed Stallings

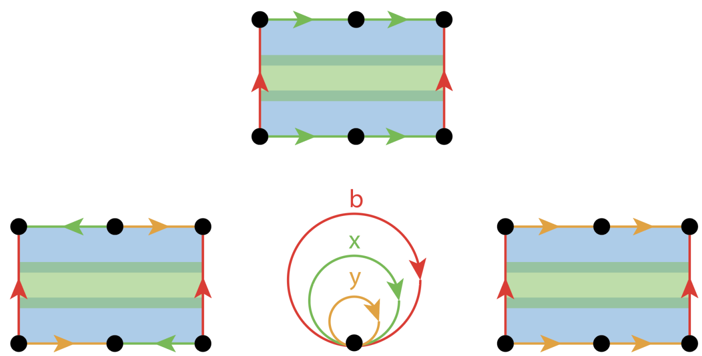

Figure 2.21.: An attempt at applying Seifert-Van Kampen's Theorem to the presentation complex of $A_{2,4,4} = \langle b, x, y | bx^2b^{-1} = x^2, by^2b^{-1} = y^2, bx^{-1}yb^{-1} = yx^{-1} \rangle$.

fold occurs, as seen in Figure 2.22. The final graph in Figure 2.22 has two orange paths of length $m-2$ going between the topmost triangular vertices, which, after $m-3$ open Stallings folds, will force a final closed Stallings fold to occur. This graph is the copy of $\overline{U \cap V}$ that we are attempting to complete to a combinatorial immersion into $\bar{U}$. The guaranteed presence of this closed Stallings fold forces us to look for yet another presentation that we can use to split $A_{2,3,P}$ as an amalgamated product. Wu & Ye utilized the following presentation in [25] to prove that $A_{2,3,2m}$ for $2m > 6$ splits as a graph of free groups.

Theorem 2.4.14 ([25])**.** *$A_{2,3,2m} \cong \langle b, c, x, y, d, \delta | d = xc, d = bx, y = bc, yb = cy, \delta b = c\delta, \delta = dx^{m-2}d \rangle \cong F_3 *_{F_7} F_4$ for $m > 3$.*

The details of the above isomorphism along with the presentation complexes associated to these presentations are analyzed in detail in the next section. We have now split almost every hyperbolic triangle Artin group. All that remain are $A_{2,3,P}$ for $P > 6$ and odd. Wu & Ye proved in [25] that such Artin groups actually cannot split as a graph of free groups.

Theorem 2.4.15 ([25])**.** *When $m \neq 3k+1$, $A_{2,3,2m+1}$ does not split as a graph of free groups.*

Proof. The details of this proof will take us too far off course from the rest of the paper, but we provide here a brief sketch of the foundational arguments. We begin by proving that every action

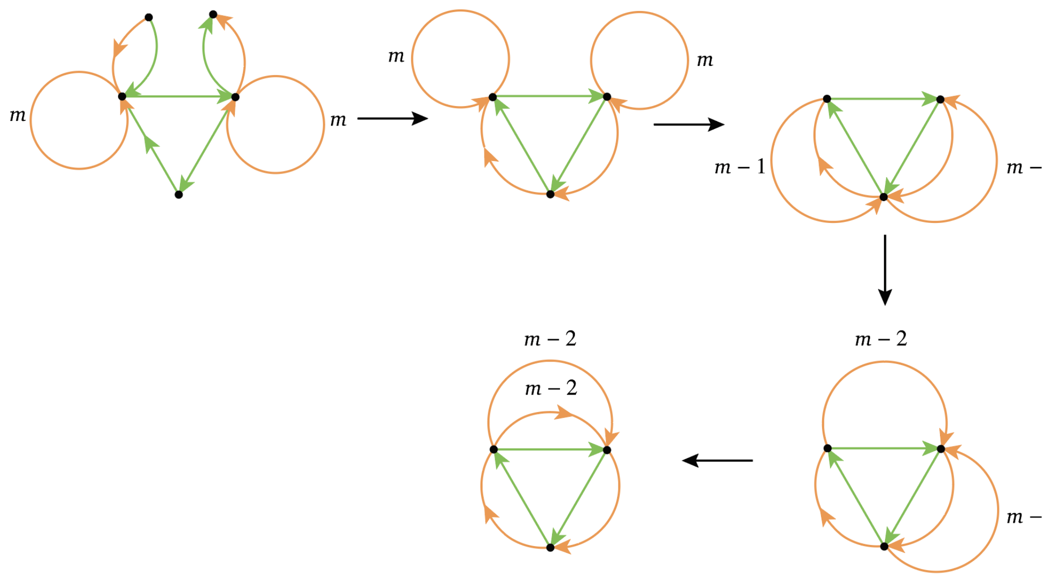

Figure 2.22.: Stallings folds in the $A_{2,3,2m}$ case

by isometries of $A_{2,3,2m+1}$ on a simplicial tree T either has a global fixed point or a geodesic line that is invariant under that action. We prove this by assuming to the contrary that there exists an action on a tree T without a global fixed point or invariant geodesic line. This forces the existence of a geodesic path with specially constructed vertex stabilizers that in turn forces one of the generating elements $a \in A_{2,3,2m+1}$ to have a nonempty fixed point set. This generator a is special in that it commutes with c, and we use this nonempty fixed point set to prove that $Fix(a) \neq \emptyset$ forces the existence of a global fixed point, providing us with a contradiction.

Now assume that $A_{2,3,2m+1}$ splits as a graph of free groups. Then $A_{2,3,2m+1}$ acts on the Bass-Serre tree, T, of this splitting. The elements that fix the vertices of T are conjugates of the vertex groups. If there were to be a global fixed point of the action on T, this would force $A_{2,3,2m+1}$ to consist only of elements from a conjugate of a vertex group, which would make the splitting trivial and $A_{2,3,2m+1}$ a free group, which it clearly is not. Therefore there must be a geodesic line l in T that is invariant under the action of $A_{2,3,2m+1}$.

The Artin group $A_{2,3,2m+1}$ acts on l by translations, and translations of a line form a copy of $\mathbb{Z}$. This allows us to define a map $f : A_{2,3,2m+1} \longrightarrow \mathbb{Z}$ that factors through $A^{Ab}_{2,3,2m+1}$. The Abelianization $A^{Ab}_{2,3,2m+1} = \mathbb{Z}$ due to the Artin relations and the newfound Abelian relations forcing all of the generators to be identified. Therefore $\ker f$ is the kernel of the Abelianization map, which is the commutator subgroup, A, of $A_{2,3,2m+1}$. Therefore there is a vertex in l with A as its stabilizer. But the vertices in l are also vertices in T, meaning that the vertex stabilizers are conjugates of free groups. If A were to be free, then this would make $A_{2,3,2m+1}$ free-by-cyclic, and thus also coherent [26]. A combination of results from [27] and [28] forces the only coherent triangle Artin groups to be of the form $A_{2,2,n}$. Therefore $A_{2,3,2m+1}$ cannot be coherent and thus cannot split as a graph of free groups. ⧗

The case when $m = 3k + 1$ is proven in a similar fashion, except there is now a third condition to consider where $A_{2,3,6k+3}$ has a vertex stabilizer that contains specific elements that keep the vertex stabilizer from being a free group, and the group itself from splitting as a graph of free groups. The goal of the next section is to use the splitting of Wu & Ye from Theorem 2.4.14 to prove that all Artin groups $A_{2,3,2n}$ for $n > 3$ are residually finite using similar methods as Jankiewicz.

Chapter 3.

Main Work

§ 3.1. Splitting $A_{2,3,2n}$ For $n > 3$

We begin by fixing the presentation of $A_{2,3,2n}$ from Lemma 2.4.14, which will be used to split $A_{2,3,2n}$ as a graph of free groups. An isomorphism between

$$A_{2,3,2n} = \langle a, b, c | ac = ca, bcb = cbc, (ab)^{2n} = (ba)^{2n} \rangle$$

and the presentation in Lemma 2.4.14 can be defined by mapping, in one direction:

$$\phi(b) = b, \phi(c) = c, \phi(x) = cbac^{-1}$$

$$\phi(y) = bc, \phi(d) = bcbac^{-1}, \phi(\delta) = bc(ba)^n$$

and in the other direction:

$$\psi(a) = b^{-1}c^{-1}xc, \psi(b) = b, \psi(c) = c$$

The presentation complex associated with the presentation in Lemma 2.4.14 can be seen in Figure 3.1. In order to realize $A_{2,3,2n}$ as an amalgamated product, we apply Seifert-Van Kampen's Theorem to X as shown in Figure 3.2.

The red band, U, and the blue band, V, are path connected open subsets of X with nonempty intersection, namely the purple band. There is a natural deformation retract that we then perform on U and V respectively to obtain the graphs $\bar{U}$ and $\bar{V}$, as shown in Figures 3.3 & 3.4.

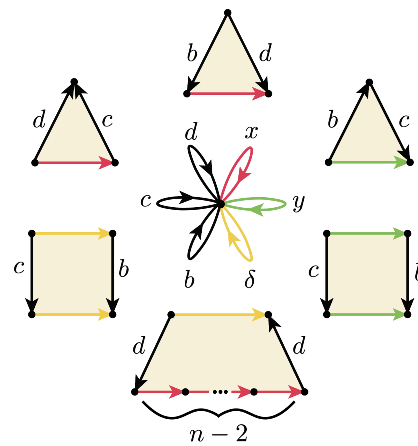

Figure 3.1.: Presentation complex X

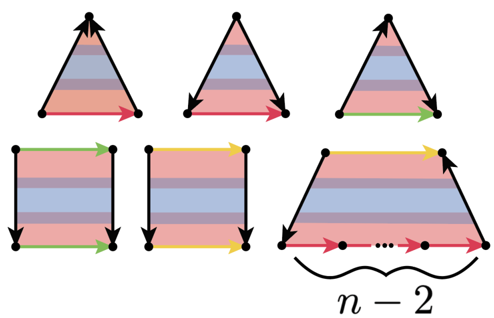

Figure 3.2.: Seifert Van Kampen's Theorem applied to the 2-cells in X

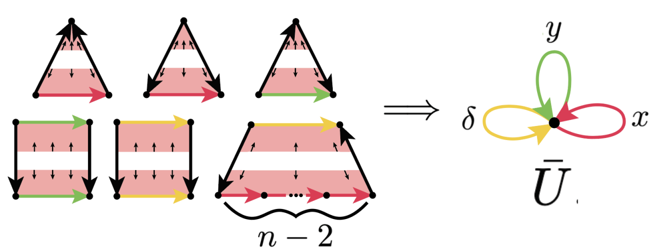

Figure 3.3.: U deformation retracts to the wedge of 3 circles, $\bar{U}$.

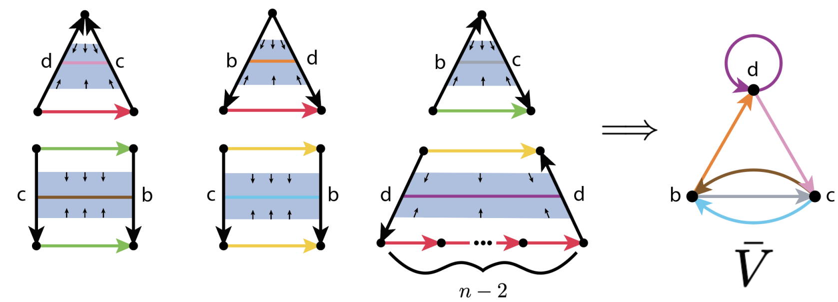

Figure 3.4.: V deformation retracts to the graph $\bar{V}$ with $\pi_1(\bar{V}) \cong F_4$.

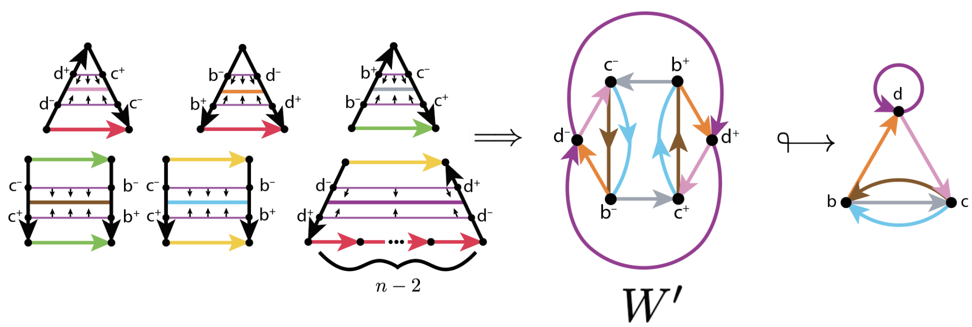

Figure 3.5.: W', the deformation retract of $U \cap V$ when viewed as a graph that can be combinatorially immersed into $\bar{V}$.

Next, we deformation retract $U \cap V$ to the graph $\overline{U \cap V}$ constructed from the purple horizontal lines in Figures 3.5 & 3.6. We will soon see that the maps $\overline{U \cap V} \to \bar{U}$ and $\overline{U \cap V} \to \bar{V}$ are combinatorial immersions.

Lemma 3.1.1. *The induced map $\overline{U \cap V} \looparrowright \bar{V}$ is a combinatorial immersion.*

Proof. The graph W' in Figure 3.5 is the image of $\overline{U \cap V}$ under the composition $\overline{U \cap V} \hookrightarrow U \cap V \hookrightarrow V \to \bar{V}$ where the final map is the deformation retract of V. The combinatorial immersion from $W' \looparrowright \bar{V}$ is represented by the coloring of the edges in W'.

Notice further that W' is a double cover of $\bar{V}$, making obvious the fact that $W' \looparrowright \bar{V}$ is a combinatorial immersion. The mapping $\overline{U \cap V} \to \bar{U}$ is not as simple.

Lemma 3.1.2. *The induced map $\overline{U \cap V} \looparrowright \bar{U}$ is a combinatorial immersion.*

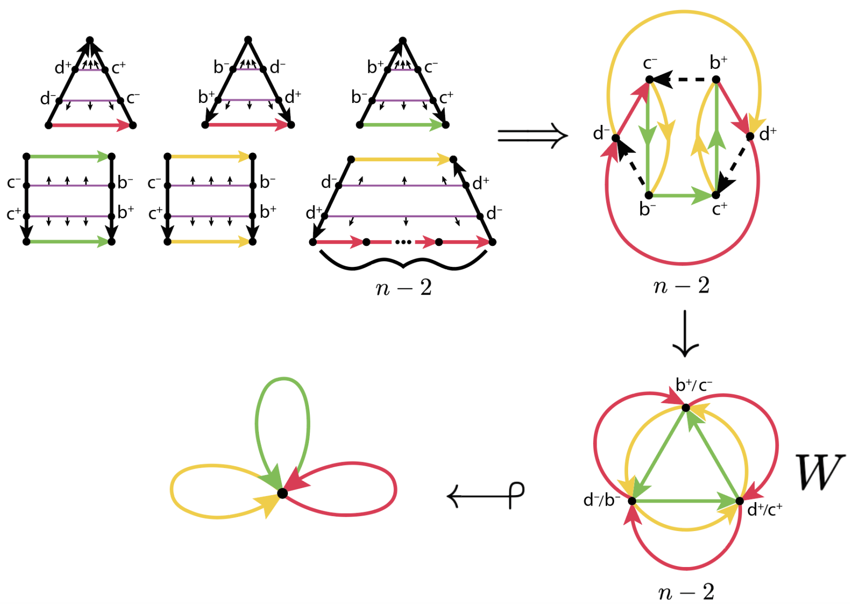

Figure 3.6.: W, the deformation retract of $U \cap V$ when viewed as a graph that can be combinatorially immersed into $\bar{U}$.

Proof. The graph W in Figure 3.6 is the image of $\overline{U \cap V}$ under the composition $\overline{U \cap V} \hookrightarrow U \cap V \hookrightarrow U \longrightarrow \bar{U}$. The graph above W is an intermediate step in the construction of W where the dashed edges are collapsed to vertices in W under the deformation retract of $\bar{U}$. Once again, the combinatorial immersion $\overline{U \cap V} \looparrowright \bar{U}$ is defined by the coloring of the edges in W. ■

We utilize these combinatorial immersions to induce the injections $\pi_1(W) \hookrightarrow \pi_1(\bar{U})$ and $\pi_1(W') \hookrightarrow \pi_1(\bar{V})$ respectively. Since we would need infinitely many vertices to complete W to a cover of $\bar{U}$, this makes $\pi_1(W)$ an infinite-index subgroup of $\pi_1(\bar{U})$. Since W' is a double cover of $\bar{V}$, this makes $\pi_1(W')$, an index-2 subgroup of $\pi_1(\bar{V})$. Both $\pi_1(W)$ and $\pi_1(W')$ are isomorphic to F_7, making W and W' homotopy equivalent. We denote this homotopy equivalence by $\sigma : W' \longrightarrow W$. Intuitively, σ maps each edge in W' to the closest horizontal component at the top/bottom of the relator polygon from which that edge is derived. For details about the behavior of σ, see Figure 3.9. The deformation retracts, combinatorial immersions and σ all fit into the following diagram that commutes up to homotopy:

$$\begin{tikzcd}
& & W \arrow[d, hook] \\
\bar{V} \arrow[d, hook] & W' \arrow[l, hook'] \arrow[r] \arrow[ur, "\sigma"] \arrow[d, hook] & \bar{U} \arrow[d, hook] \\
V & U \cap V \arrow[l, hook'] \arrow[r, hook] & U
\end{tikzcd}$$

This commutative diagram coming from Lemmas 3.1.1 & 3.1.2, along with Seifert-Van Kampen's Theorem being applied to the presentation complex for $A_{2,3,2n}$, proves the following theorem.

Theorem 3.1.3. *For $n \geq 3$, $A_{2,3,2n} \cong \pi_1(U) *_{\pi_1(U \cap V)} \pi_1(V)$ where*

- $\pi_1(U) = \pi_1(\bar{U}) \cong F_3$
- $\pi_1(V) = \pi_1(\bar{V}) \cong F_4$
- $\pi_1(U \cap V) = \pi_1(W) \cong \pi_1(W') \cong F_7$

Since this graph of groups is a graph of free groups, the vertex and edge groups implicitly satisfy the first 2 criteria in [29, Theorem 2.3], thereby reducing [29, Theorem 2.3] to the more approachable theorem below.

Theorem 3.1.4 (Huang-Wise [29])**.** *Let G be a graph of finite rank free groups. If G has finite stature with respect to its vertex groups, then G is residually finite.*

We will discuss finite stature in greater detail in the following subsection. For now we end this section with an explicit generating set for $\pi_1(W, 1)$, which will be used in later sections.

Lemma 3.1.5. $\pi_1(W, 1) = \langle x^n, y^3, y\delta^{-1}, y(y\delta^{-1})y^{-1}, \delta^{-1}y, yx, y(yx)y^{-1} \rangle$.

Proof. In order to see that this is a generating set for $\pi_1(W, 1)$, collapse the maximal tree shown in Figure 3.7. Recall that each red edge corresponds to the algebraic element x, each green edge corresponds to y and each yellow edge corresponds to δ. The loops present after the collapse of this maximal tree are in one-to-one correspondence with the generating set in the lemma. ■

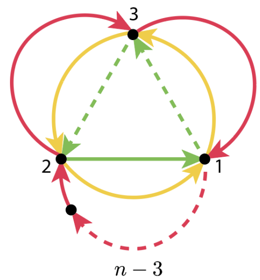

Figure 3.7.: Collapsing the maximal tree.

§ 3.2. Finite Stature Procedure

In order to prove that $A_{2,3,2n}$ is residually finite, we first prove that $A_{2,3,2n}$ has finite stature with respect to $\{\pi_1(\bar{U}), \pi_1(\bar{V})\}$, the vertex groups of the splitting.

Definition 3.2.1 ([29])**.** Let G be a group and $\{H_\lambda\}_{\lambda\in\Lambda}$ be a collection of subgroups of G. Then G has finite stature with respect to $\{H_\lambda\}_{\lambda\in\Lambda}$ if for each $\mu \in \Lambda$, there are finitely many H_μ-conjugacy classes of infinite subgroups of the form $H_\mu \cap D$ where D is an intersection of G-conjugates of elements in $\{H_\lambda\}_{\lambda\in\Lambda}$.

In our case the set $\{H_\lambda\}_{\lambda\in\Lambda}$ is the set of vertex groups of our splitting, namely $\{\pi_1(\bar{U}), \pi_1(\bar{V})\}$. Therefore each D is an intersection of $A_{2,3,2n}$-conjugates of the vertex groups. Conjugates of vertex groups are stabilizers of vertices in the Bass-Serre Tree T associated with the graph of groups. Therefore the intersection of conjugates of vertex groups is a group element that stabilizes multiple vertices in T. Since the action of elements of G on T preserves vertex-adjacency, the action of D on T maps paths to paths. If D fixes two vertices in T, then the path between the vertices is unique, forcing D to fix the entire path in T between the vertices. Edge stabilizers in T are conjugates of $\pi_1(\bar{U}) \cap \pi_1(\bar{V}) = \pi_1(W) \cong \pi_1(W')$. Therefore D is either $\pi_1(\bar{U})$, $\pi_1(\bar{V})$ or the intersection of $A_{2,3,2n}$-conjugates of $\pi_1(W) \cong \pi_1(W')$.

Thus far we have only discussed the groups in our graph of groups in terms of the fundamental group of graphs. We would like to continue to proceed in this fashion. In order to do so, we need to understand what subgroup intersection and conjugation look like in the context of graphs.

3.2.1. Fiber Products

We will be performing fiber products of covers of $W \sim W'$ throughout the rest of the paper. We choose to work with covers of W instead of W' since $\bar{U}$ has only one vertex, guaranteeing that the immersions ϕ_1 and ϕ_2 agree on every vertex. This results in $V(Y_1) \times V(Y_2)$ being the vertex set of every fiber product for the rest of this paper. Figure 3.8 shows an example of a fiber product calculation that will be used later on.

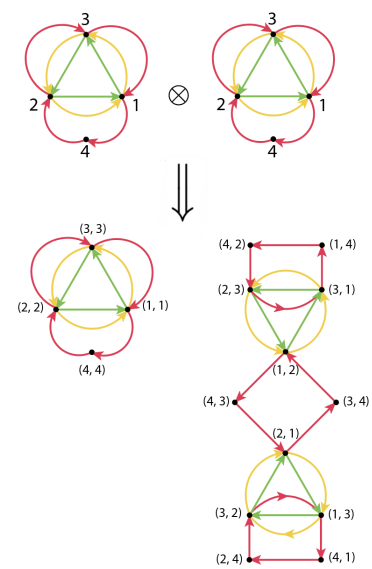

Figure 3.8.: $W \otimes_{\bar{U}} W$ for $A_{2,3,8}$

Lemma 3.2.1. *Let $Y_1 \subseteq X_1$ and $Y_2 \subseteq X_2$ be subgraphs. Then $Y_1 \otimes_{\bar{U}} Y_2 \subseteq X_1 \otimes_{\bar{U}} X_2$.*

This lemma is immediate by the definition of fiber products. Also, algebraically, it is intuitive that the intersection of subgroups will always be a subgroup of the intersection of the supergroups.

3.2.2. Basepoint Translation

We now must discuss how to understand conjugation of subgroups when our subgroups are being considered as the fundamental group of graphs. We begin by exploiting the following property of $A_{2,3,2n}$, $n > 3$.

Lemma 3.2.2. *Every element $g \in A_{2,3,2n}$ with $n > 3$ can be written as $g = g_1 v g_2 v \ldots g_{m-1} v g_m$ where $g_1, \ldots, g_k \in \pi_1(\bar{U})$ and $v \in \pi_1(\bar{V})$ is a fixed coset representative coming from $\pi_1(\bar{V})/\pi_1(W') = \{[1], [v]\}$.*

Proof. $A_{2,3,2n} \cong \pi_1(\bar{U}) *_{\pi_1(W)=\pi_1(W')} \pi_1(\bar{V}) = (\pi_1(\bar{U}) * \pi_1(\bar{V}))/(\pi_1(W) = \pi_1(W'))$. Therefore every element $g = h_1 k_1 h_2 k_2 \ldots h_{m-1} k_{m-1} h_m$ where each $h_i \in \pi_1(\bar{U})$ and $k_i \in \pi_1(\bar{V})$. Since $\pi_1(W')$ is an index-2 subgroup of $\pi_1(\bar{V})$, each $k_i = l_i v$ for some $l_i \in \pi_1(W') = \pi_1(W)$. Therefore $g = h_1 k_1 \ldots h_{m-1} k_{m-1} h_m = h_1 l_1 v h_2 l_2 v \ldots h_{m-1} l_{m-1} v h_m$. Setting $h_i l_i = g_i$ for $1 < i < m$ and $h_m = g_m$ results in the desired expansion of g. ◼

Consider a subgroup $H < \pi_1(W)$. We can represent $H \cong \pi_1(Y, p)$ where $Y \looparrowright W$. By Lemma 3.2.2 it suffices to understand how conjugation by elements of $\pi_1(\bar{U})$, and conjugation by the coset representative $v \in \pi_1(\bar{V})$, affect H.

When we conjugate H by an element $g \in \pi_1(\bar{U})$, this naturally corresponds to translating the basepoint p along the path labelled by g in Y to the vertex p_g just like Figure 2.18. But since $v \notin \pi_1(\bar{U})$, we cannot represent conjugation by v of a subgroup of $\pi_1(W) < \pi_1(\bar{U})$ by simply shifting the basepoint along a path in the corresponding cover of $\bar{U}$. Consider $v^{-1}\pi_1(W')v$. Since $\pi_1(W') \trianglelefteq \pi_1(\bar{V})$, $v^{-1}\pi_1(W')v \cong \pi_1(W')$. Since W' is a cover of $\bar{V}$ and $v \in \pi_1(\bar{V})$, conjugating $\pi_1(W')$ by v corresponds to shifting the basepoint of W' along the path v. The basepoint of W' being one of $\{d^+, d^-, b^+, b^-, c^+, c^-\}$ and v being a loop in $\bar{V}$ means that it must take the basepoint to its partner, meaning $d^+ \leftrightarrow d^-, b^+ \leftrightarrow b^-$ or $c^+ \leftrightarrow c^-$ depending on the basepoint. This is all to say that $\pi_1(W', p_v) \cong v^{-1}\pi_1(W', p)v$, and the isomorphism is defined on the vertices by $d^+ \leftrightarrow d^-, b^+ \leftrightarrow b^-, c^+ \leftrightarrow c^-$. This isomorphism extends naturally to the edges.

Consider the homotopy equivalence $\sigma : W' \longrightarrow W$ introduced in the previous section. Since σ is a homotopy equivalence, it extends naturally to all covers of W' as well. Figure 3.9 shows how σ behaves locally on edges. If we want to translate the basepoint of a cover Y of W by v, we start by considering the graph Y' that is the cover of W' such that $\sigma(Y') = Y$. The v-translated copy of Y', $(Y')^v$, is a cover of the v-translated copy of W', which is a copy of W' with $+ \leftrightarrow -$ swapped in every superscript. Therefore $(Y')^v$ must likewise be obtained by swapping $+ \leftrightarrow -$ in the superscripts of every vertex in Y'. Define $\bar{\beta} : Y' \longrightarrow (Y')^v$ to be defined by this swapping of superscripts.

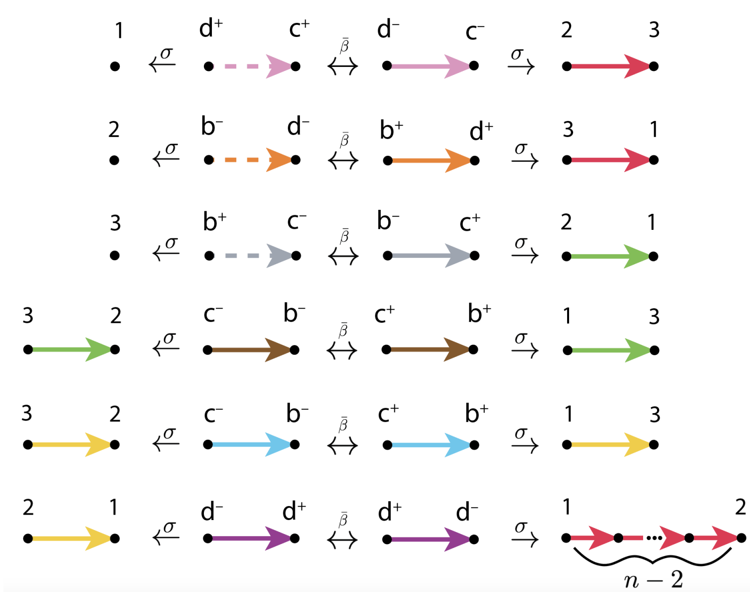

Figure 3.9.: $\beta = \sigma \circ \bar{\beta}$ applied to every applicable vertex & edge.

Putting these pieces together, we can calculate Y^v, for any cover Y of W by applying $\beta = \sigma \circ \bar{\beta}$ to Y'. Pictorially, we can think of $\beta(Y')$ as swapping the tops and bottoms of the relator polygons in the presentation complex. Figure 3.9 shows the details of how β affects the vertices and edges of the core of every cover of W. In the figure, the dotted arrows are meant to represent edges that are being mapped to a vertex by σ. An example of a β calculation can be found in Figure 3.13.

To recap, the basepoint translation of a graph $Y \looparrowright W$ can either result in:

1. Y with a new basepoint chosen, when Y is basepoint translated by an element of $\pi_1(\bar{U})$ or

2. A (potentially) new graph $\beta(Y)$ with a (potentially) new basepoint, when Y is basepoint translated by the nontrivial coset representative $\upsilon \in \pi_1(\bar{V})$.

3.2.3. Defining the Set S

Lemma 3.2.3. *If there exists a finite set S of finite cores of covers of W such that S contains:*

- *W*
- *Every connected component of every fiber product of elements in S*
- *The image under β of every element in S*

Then $A_{2,3,2n}$ has finite stature with respect to its vertex groups.

Proof. Such a set S corresponds to a collection of $\pi_1(\bar{U})$-conjugacy classes & $\pi_1(\bar{V})$-conjugacy classes of subgroups of $\pi_1(W)$, and these conjugacy classes are closed under intersection. Therefore S contains every graph Y_D corresponding to a subgroup D as described in Definition 3.2.1. By construction, $\pi_1(\bar{U}) \cap \pi_1(\bar{V}) = \pi_1(W) \cong \pi_1(W')$. Assume $D \ncong \pi_1(\bar{U})$ or $\pi_1(\bar{V})$. Then $D < \pi_1(W) = \pi_1(W')$ and $\pi_1(\bar{U}) \cap D \cong D < \pi_1(\bar{U})$. Similarly, $\pi_1(\bar{V}) \cap D \cong D < \pi_1(\bar{V})$. Therefore, if there are finitely many elements in S, there are also finitely many subgroups $\pi_1(\bar{U}) \cap D$ and $\pi_1(\bar{V}) \cap D$. Since each Y_D has finitely many vertices in its core and Y_D is closed under basepoint translation by all elements of $\pi_1(\bar{U})$ and $\pi_1(\bar{V})$, there are finitely many conjugacy classes of $\pi_1(\bar{U}) \cap D$ and $\pi_1(\bar{V}) \cap D$. Therefore $A_{2,3,2n}$ has finite stature by definition.

Connected components coming from fiber products have tuples for vertex labels. The maps σ and β are defined with respect to vertices that have integer labels. We will need to calculate β of connected components coming from fiber products, so we choose to project each vertex-tuple to its second component. This will allow us to apply β to the connected component, and to decide which graph in S (if any) that the connected component is a subgraph of. Because we are making this choice, we must perform $H \otimes_{\bar{U}} K$ and $K \otimes_{\bar{U}} H$ for all pairs $H, K \in S$, despite the fact that $(H \otimes_{\bar{U}} K) \cong (K \otimes_{\bar{U}} H)$ by swapping the tuple-values at each vertex.

3.2.4. q-Contractibility

In Lemma 2.4.14, we fixed the presentation

$$\pi_1(W, 1) = \langle x^n, y^3, y\delta^{-1}, y(y\delta^{-1})y^{-1}, \delta^{-1}y, yx, y(yx)y^{-1} \rangle$$

Consider the quotient $\bar{q} : \pi_1(W, 1) \longrightarrow \pi_1(W, 1)/\langle\langle N \rangle\rangle$ where

$$N = \langle x^n, y^3, y\delta^{-1}, y^2\delta^{-1}y^{-1}, \delta^{-1}y \rangle$$

Definition 3.2.2. For every $K \looparrowright W$, define $q(K)$ to be the minimal 2-complex with K as its 1-skeleton and 2-cells attached to all simple loops whose edges form a word in $\langle\langle N \rangle\rangle$. The complex $q(K)$ satisfies $\bar{q}(\pi_1(K)) = \pi_1(q(K))$.

Definition 3.2.3. Let K be the core of a cover of W. The graph K is q-fillable if every path in K whose edges form a generator of N with the above presentation is a simple loop.

A graph K being q-fillable means that every loop in K that corresponds to an element of N is built from simple loops, each of which corresponds to a generator of N.

Lemma 3.2.4. *Suppose that $H, K \looparrowright W$ are q-fillable. Then:*

- *$H \otimes_{\bar{U}} K$ is q-fillable,*
- *$\beta(H)$ is q-fillable.*

Proof. Let L be a connected component of $H \otimes_{\bar{U}} K$. Then L combinatorially immerses into both H and K. Let ℓ be a path based at (h, k) in L whose edge labels form a generating element of N. The fiber product definition allows us to view ℓ also as a path based at h in H and a path based at k in K. Since H and K are q-fillable, this makes ℓ a simple loop in both H and K. Therefore the only vertex in H reached twice by ℓ is h and the only vertex in K reached twice by ℓ is k. Therefore the only vertex in L reached twice by ℓ is (h, k), making ℓ a simple loop in L and $H \otimes_{\bar{U}} K$ q-fillable by definition.

Let ℓ be a path in $\beta(H)$ whose edge labels form a generating element of N. Figure 3.10 shows that β maps every generator of N to a generator of N. Therefore $\beta^{-1}(\ell)$ is a collection of paths in H whose edge-labels form the same generating element of N. Since H is q-fillable, this makes every path in $\beta^{-1}(\ell)$ a simple loop. Therefore ℓ itself is also forced to be a simple loop, making $\beta(H)$ q-fillable. ▮

Lemma 3.2.5. *Let H, K be q-fillable graphs such that $q(K)$ is contractible. Then every connected component of $K \otimes_{\bar{U}} H$ and $H \otimes_{\bar{U}} K$ is a subgraph of K.*

Proof. Let L be a connected component of $K \otimes_{\bar{U}} H$. Then there is a combinatorial immersion $\phi : L \to K$. In order for ϕ to not be an embedding, there needs to exist an open path p in L such that $\phi(p)$ is a loop in K. Assume such an open path p exists. Since K is q-contractible, this makes $\phi(p)$ a loop whose edge-labels form an element in N. Therefore p is a path in L whose edge labels form an element in N. Lemma 3.2.4 guarantees that L is q-fillable, forcing every every path in L whose edge labels form an element of N to be a loop. Therefore p is a loop, forcing ϕ to be an embedding. ▮

Lemma 3.2.6. *Let K be a q-fillable graph such that $q(K)$ is contractible. Then $q(\beta(K))$ is contractible as well.*

Proof. $q(K)$ contractible means that every loop in K is a combination of path-translations of red cycles of length n, green triangles and yellow-green bigons. As shown in Figure 3.10, this collection of loops is closed under β. Therefore $\beta(K)$ also consists only of loops that arise as combinations of path-translations of red cycles of length n, green triangles and yellow-green bigons, making $q(\beta(K))$ contractible as well. ▮

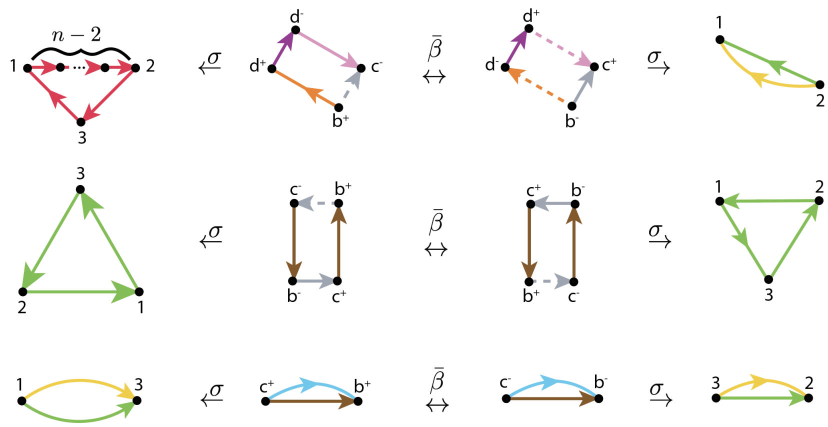

Figure 3.10.: β applied to the boundaries of the 2-cells.

§ 3.3. Residual Finiteness of $A_{2,3,2n}$ For $n > 4$

The goal of this section is to prove the following theorem:

Theorem 3.3.1. *$A_{2,3,2n}$ for $n > 4$ is residually finite.*

Combining Lemma 3.2.3 with Theorem 3.1.4 reduces the proof of this theorem to proving that a finite set S, as described in Lemma 3.2.3, exists.

3.3.1. Iterative Construction of S

The general procedure for constructing S is as follows. Begin with $S = \{W\}$, then:

- Perform $H \otimes_{\bar{U}} H$, $H \otimes_{\bar{U}} K$ and $K \otimes_{\bar{U}} H$ for each $H, K \in S$. Project each vertex-tuple to its second component and add to S any resulting connected component that is not a subgraph of an element already in S.
- Calculate $\beta(H)$ for every $H \in S$. Add $\beta(H)$ to S if it is not a subgraph of any element of S.
- Repeat the above two steps until no new graphs can be added to S in this fashion.

- Add all subgraphs of elements of S to S.

We will prove that such a process will terminate, resulting in a finite set S of finite graphs. We begin with $S = \{W\}$ since S must, at minimum, contain W. Conjugating W by any element of $\pi_1(\bar{U})$ does not change W, it just shifts the basepoint. Also, since $\pi_1(W) \cong \pi_1(W')$ and $\pi_1(W') \trianglelefteq \pi_1(\bar{V})$, $v\pi_1(W)v^{-1} \cong v\pi_1(W')v^{-1} \cong \pi_1(W') \cong \pi_1(W)$. Pictorially, this is shown in Figure 3.11.

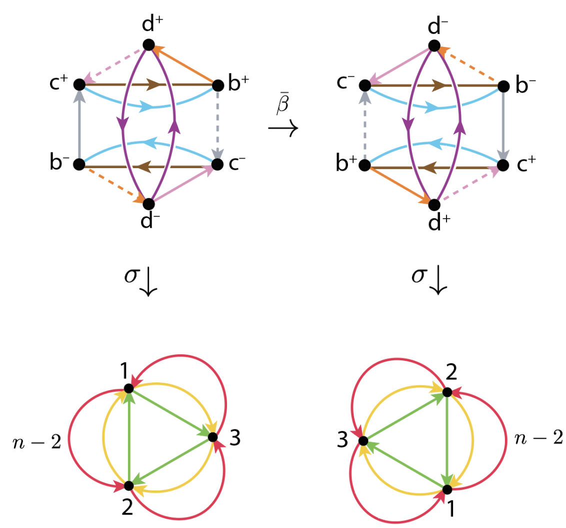

Figure 3.11.: $\beta(W) \cong W$

We end this subsection with the statement of Lemma 3.3.2, which we will spend the remainder of the section proving.

Lemma 3.3.2. *The set S is finite for $A_{2,3,2n}$ with $n > 4$.*

So far $S = \{W\}$ is closed under basepoint translation, but we also S need to be closed under fiber product.

Lemma 3.3.3. *For $A_{2,3,2n}$ with $n > 4$, $W \otimes_{\bar{U}} W$ has the following connected components:*

- *One copy of W*

- *$n - 5$ copies of an n-gon with red edges*
- *One of each of the graphs in Figure 3.12, to be denoted henceforth by X_1 (leftmost) and X_2 (rightmost).*

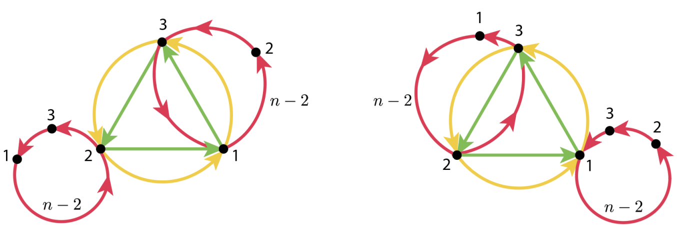

Figure 3.12.: X_1 & X_2

Proof. W has three yellow edges, three green edges and n red edges. Therefore $W \otimes_{\bar{U}} W$ must have nine yellow edges, nine green edges and n^2 red edges. By direct computation, these yellow edges and green edges arise as the triangles $(1,1) \to (3,3) \to (2,2) \to (1,1)$ in the copy of W, $(1,2) \to (3,1) \to (2,3) \to (1,2)$ and $(1,3) \to (3,2) \to (2,1) \to (1,3)$. After projecting to the second component of each tuple, these triangles become the triangles present in X_1 and X_2 respectively. The placement of the red edges in W, X_1 and X_2 follow directly from these vertex-labelings. W has n vertices, X_1 has $2n$ vertices and X_2 has $2n$ vertices. Therefore there are $n^2 - 5n$ vertices unaccounted for. Since we have identified all of the yellow and green edges in $W \otimes_{\bar{U}} W$, only red edges can connect the remaining $n^2 - 5n$ vertices. Every vertex in W is contained in a red loop of length n, therefore every vertex in $W \otimes_{\bar{U}} W$ must also be contained in a red loop of length n. Therefore the remaining $n^2 - 5n$ vertices must occur in $n - 5$ red loops of length n. ∎

3.3.2. The Proof of Lemma 3.3.2

Proof. Lemma 3.3.3 describes the connected components of $W \otimes_{\bar{U}} W$ as being a copy of W, X_1 and X_2, along with a collection of red n-cycles. These red n-cycles are all subgraphs of W, and so by Lemma 3.2.1 they are only capable of producing subgraphs of fiber products involving W to S. At the end we will include all unique subgraphs of elements of S, but for the first two steps of our construction we are focused on new maximal elements of S.

So far $S = \{W, X_1, X_2\}$. We need S to be closed under β, so we must calculate $\beta(X_1)$ and $\beta(X_2)$. These calculations are carried out in Figures 3.13 & 3.14 by applying $\bar{\beta}$ and σ to each edge and referring to Figure 3.9 for the image of each edge. These calculations result in two new graphs that must be included in S. The vertices $\bar{1}$ (resp. $\bar{2}$) in $\beta(X_1)$ and $\beta(X_2)$ is a preimage of the vertex labelled 1 (resp. 2) in W, and the notation is used to differentiate the two preimages for ease of computation later on.

We now have $S = \{W, X_1, X_2, \beta(X_1), \beta(X_2)\}$, a set closed under β. It remains to show that every connected component of each pairwise fiber product of elements in this set is either a subgraph of a graph in S or is q-contractible. To prove this claim we proceed in a similar fashion as in Lemma 3.3.3. As n increases, the number of red edges increases, but the number of yellow and green edges remains the same, so it suffices to focus on how the yellow and green edges are dispersed among the connected components. The collection of connected components of a fiber product that contain yellow & green edges will be termed the relevant portion for the remainder of the proof. The other connected components will be dealt with at the end of the proof.

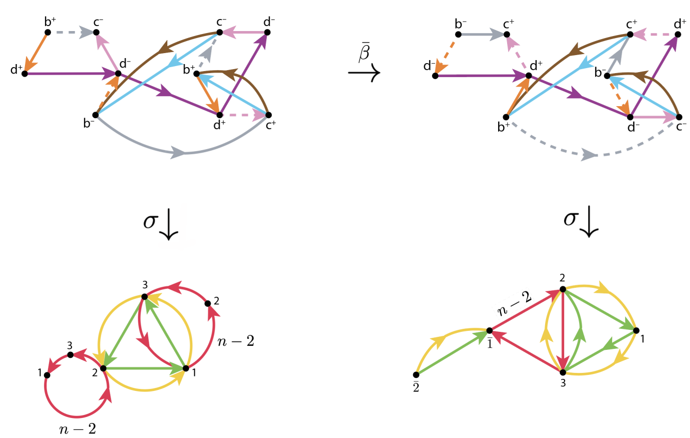

Figure 3.13.: $\beta(X_1)$

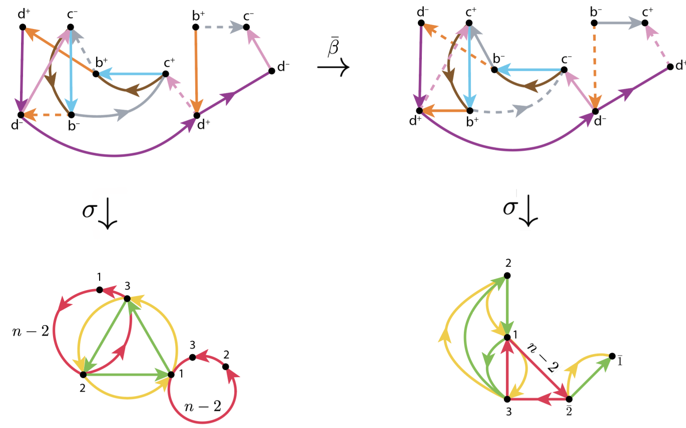

Figure 3.14.: $\beta(X_2)$

All fiber products involving only W, X_1 and X_2

We begin by calculating the relevant portion of $W \otimes_{\bar{U}} X_1$. By direct computation, these components are two copies of X_1 and a graph that is q-contractible, as seen in Figure 3.15. The presence of q-contractible graphs in the fiber product computations will be addressed at the end of the proof. For now, we focus on adding to S only the non q-contractible graphs that arise.

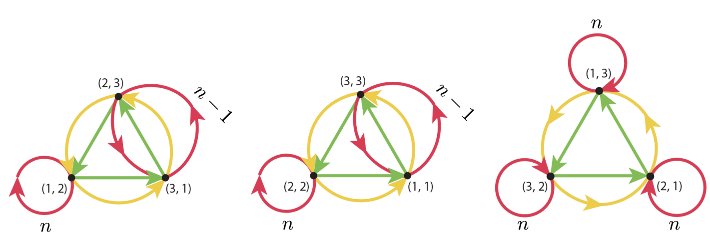

Figure 3.15.: The relevant portion of $W \otimes_{\bar{U}} X_1$

We compute the relevant portion of $X_1 \otimes_{\bar{U}} W$ by swapping the entries in each vertex-tuple in Figure 3.15, resulting in X_1, X_2 and a q-contractible graph.

The fiber products used to build S are with respect to $\bar{U}$, meaning that we view each component of the fiber product as being combinatorially immersed in $\bar{U}$. The graph $\bar{U}$ has only one vertex, so the vertex labellings of the graphs in S are therefore irrelevant in the fiber product calculations.

Notice that X_1 and X_2 are identical graphs when the vertices are not labelled. We denote this property by $X_1 \cong X_2$. Since $X_1 \cong X_2$, the connected components that arise from $Y \otimes_{\bar{U}} X_1$ (resp. $X_1 \otimes_{\bar{U}} Y$) will be the same graphs, up to vertex relabelling, as the connected components in $Y \otimes_{\bar{U}} X_2$ (resp. $X_2 \otimes_{\bar{U}} Y$) for all $Y \looparrowright \bar{U}$.

In particular, the connected components of $W \otimes_{\bar{U}} X_2$ are identical to the connected components of $W \otimes_{\bar{U}} X_1$ up to vertex labels. Since a graph is q-contractible regardless of its vertex labels, it suffices to calculate the other two relevant connected components of $W \otimes_{\bar{U}} X_2$. This results in two copies of X_2 as seen in Figure 3.16.

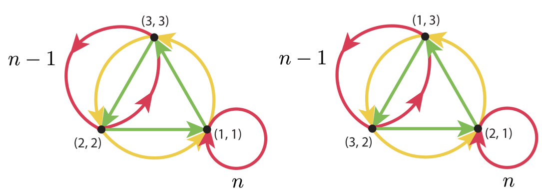

Figure 3.16.: The non q-contractible relevant portion of $W \otimes_{\bar{U}} X_2$

To calculate the non q-contractible relevant portion of $X_2 \otimes_{\bar{U}} W$, we swap the tuple-entries at each vertex in Figure 3.16, resulting in a copy of X_2 and X_1.

The relevant portion of $X_1 \otimes_{\bar{U}} X_2$ is 2 q-contractible graphs and X_2, as seen in Figure 3.17.

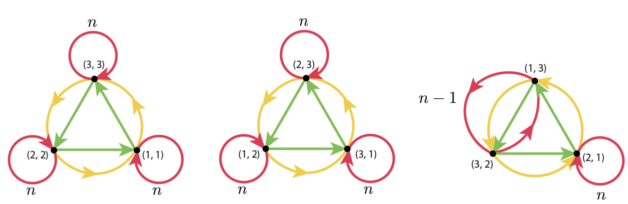

Figure 3.17.: The relevant portion of $X_1 \otimes_{\bar{U}} X_2$

To obtain the non q-contractible relevant portion of $X_2 \otimes_{\bar{U}} X_1$, we swap the tuple-entries at every vertex in the rightmost graph in Figure 3.17, resulting in a copy of X_1.

Since $X_1 \cong X_2$, there is only one non q-contractible relevant connected component in both $X_1 \otimes_{\bar{U}} X_1$ and $X_2 \otimes_{\bar{U}} X_2$ respectively. For $X_1 \otimes_{\bar{U}} X_1$, that connected component is the copy of X_1

shown in Figure 3.18.

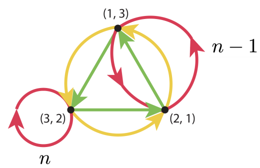

Figure 3.18.: The non q-contractible relevant portion of $X_1 \otimes_{\bar{U}} X_1$

For $X_2 \otimes_{\bar{U}} X_2$ the non q-contractible relevant connected component is the copy of X_2 shown in Figure 3.19.

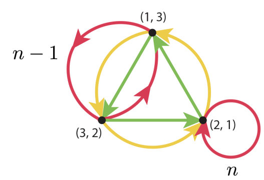

Figure 3.19.: The non q-contractible relevant portion of $X_2 \otimes_{\bar{U}} X_2$

$Z \otimes_{\bar{U}} \beta(X_1)$ **where** $Z \in \{W, X_1, X_2\}$

The relevant portion of $W \otimes_{\bar{U}} \beta(X_1)$ is a copy of $\beta(X_1)$, two q-contractible graphs and a subgraph of X_2, as shown in Figure 3.20.

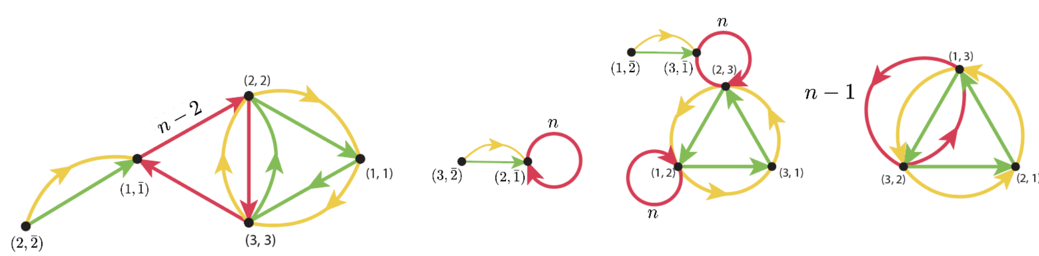

Figure 3.20.: The relevant portion of $W \otimes_{\bar{U}} \beta(X_1)$

Swapping the tuple-values at each vertex in leftmost and rightmost graphs in Figure 3.20 results in the non q-contractible relevant portion of $\beta(X_1) \otimes_{\bar{U}} W$ being $\beta(X_1)$ and a subgraph of

X_1. The relevant portion of $X_1 \otimes_{\bar{U}} \beta(X_1)$ consists of 4 q-contractible connected components and a subgraph of X_2 as shown in Figure 3.21.

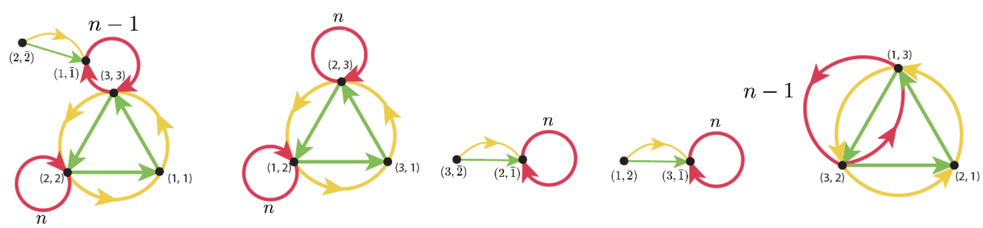

Figure 3.21.: The relevant portion of $X_1 \otimes_{\bar{U}} \beta(X_1)$

Swapping the tuple-entries at each vertex in the rightmost graph in Figure 3.21 results in the only non q-contractible relevant component of $\beta(X_1) \otimes_{\bar{U}} X_1$ being a subgraph of X_1.

Since $X_1 \cong X_2$, we know from the above calculations that there is only one non q-contractible relevant component of $X_2 \otimes_{\bar{U}} \beta(X_1)$. This component is the subgraph of X_2 shown in Figure 3.22. Swapping the tuple-entries at each vertex leaves the labels unchanged, so the only non q-contractible relevant component of $\beta(X_1) \otimes_{\bar{U}} X_2$ is the same subgraph of X_2.

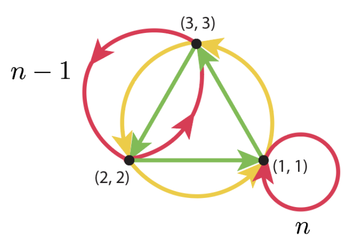

Figure 3.22.: The relevant portion of both $X_2 \otimes_{\bar{U}} \beta(X_1)$ & $\beta(X_1) \otimes_{\bar{U}} X_2$

$Z \otimes_{\bar{U}} \beta(X_2)$ **where** $Z \in \{W, X_1, X_2\}$

The relevant portion of $W \otimes_{\bar{U}} \beta(X_2)$ is a copy of $\beta(X_2)$, a subgraph of X_1 and two q-contractible graphs as shown in Figure 3.23.

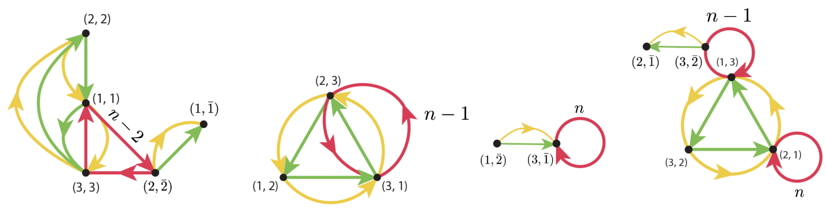

Figure 3.23.: The relevant portion of $W \otimes_{\bar{U}} \beta(X_2)$

Swapping the tuple-entries at each vertex of the two leftmost graphs in Figure 3.23 results in the non q-contractible relevant portion of $\beta(X_2) \otimes_{\bar{U}} W$ being $\beta(X_2)$ and a subgraph of X_2.

The relevant portion of $X_1 \otimes_{\bar{U}} \beta(X_2)$ consists of 4 q-contractible graphs and a subgraph of X_1 as shown in Figure 3.24.

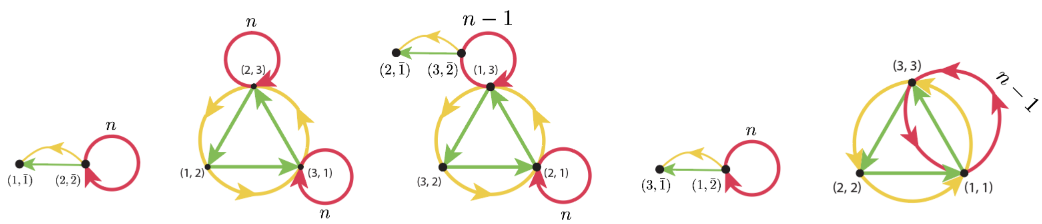

Figure 3.24.: The relevant portion of $X_1 \otimes_{\bar{U}} \beta(X_2)$

Swapping the tuple-entries at each vertex in the rightmost graph in Figure 3.24 results in the same graph being the only non q-contractible relevant component of $\beta(X_2) \otimes_{\bar{U}} X_1$.

Since $X_1 \cong X_2$, the only non q-contractible relevant component of $X_2 \otimes_{\bar{U}} \beta(X_2)$ is the subgraph of X_1 shown in Figure 3.25.

Swapping the tuple-entries at each vertex in the graph in Figure 3.25 results in the only non q-contractible relevant component of $\beta(X_2) \otimes_{\bar{U}} X_2$ being a subgraph of X_2.

$\beta(X_1) \otimes_{\bar{U}} \beta(X_1), \beta(X_1) \otimes_{\bar{U}} \beta(X_2)$ **and** $\beta(X_2) \otimes_{\bar{U}} \beta(X_2)$

The relevant portion of $\beta(X_1) \otimes_{\bar{U}} \beta(X_1)$ consists of a copy of $\beta(X_1)$ and 6 q-contractible graphs as shown in Figure 3.26.

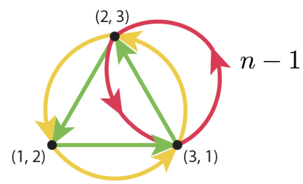

Figure 3.25.: The relevant portion of $X_2 \otimes_{\bar{U}} \beta(X_2)$

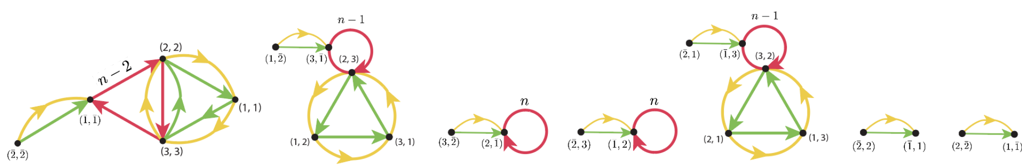

Figure 3.26.: The relevant portion of $\beta(X_1) \otimes_{\bar{U}} \beta(X_1)$

The relevant portion of $\beta(X_1) \otimes_{\bar{U}} \beta(X_2)$ is a subgraph of X_1 and 6 q-contractible graphs as shown in Figure 3.27.

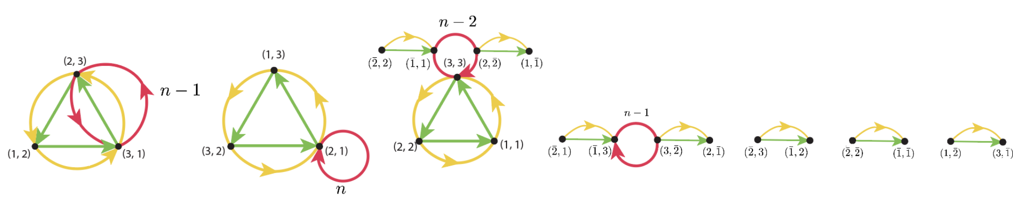

Figure 3.27.: The relevant portion of $\beta(X_1) \otimes_{\bar{U}} \beta(X_2)$

Swapping the tuple-entries of the leftmost graph in Figure 3.27 results in the only non q-contractible relevant component of $\beta(X_2) \otimes_{\bar{U}} \beta(X_1)$ being a subgraph of X_2.

The relevant portion of $\beta(X_2) \otimes_{\bar{U}} \beta(X_2)$ is a copy of $\beta(X_2)$ and 6 q-contractible graphs shown in Figure 3.28.

Assembling the Pieces To Finalize S

We now have a complete enumeration of every relevant connected component in every fiber product coming from pairs of graphs in $S = \{W, X_1, X_2, \beta(X_1), \beta(X_2)\}$. Every such component

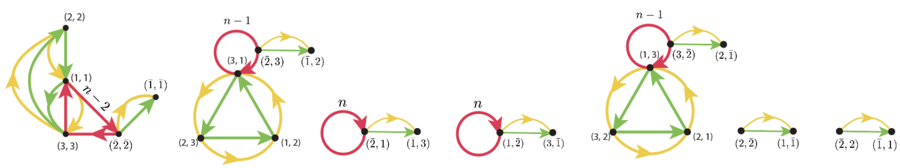

Figure 3.28.: The relevant portion of $\beta(X_2) \otimes_{\bar{U}} \beta(X_2)$

is either a subgraph of a graph in $S = \{W, X_1, X_2, \beta(X_1), \beta(X_2)\}$ or is q-contractible. In each fiber product computed in the last subsection, there are also non-relevant connected components that arise, consisting of unaccounted for vertices and red edges. Note that every red edge in every graph in $S = \{W, X_1, X_2, \beta(X_1), \beta(X_2)\}$ is contained in a red cycle of length n, which forces every red edge in every fiber product to also be contained in a red cycle of length n. Therefore the remaining vertices and red edges in each of the fiber products must arise as connected components consisting solely of one cycle of length n, which is a subgraph of W, and will therefore be included in the final step in the construction of S.

If a graph $K \in S$ is q-contractible, Lemma 3.2.5 tells us that every connected component of $K \otimes_{\bar{U}} H$ and $H \otimes_{\bar{U}} K$ is a subgraph of K for all $H \in S$. While $\beta(K)$ may not already be in S (in which case we add $\beta(K)$ to S), $\beta(K)$ also has $q(\beta(K))$ contractible by Lemma 3.2.6, and therefore every connected component of $\beta(K) \otimes_{\bar{U}} H$ and $H \otimes_{\bar{U}} \beta(K)$ is a subgraph of $\beta(K)$ for all $H \in S$ as well. Since $\beta^2 = 1$, this guarantees that any q-contractible K can contribute at most one new maximal graph to S, namely $\beta(K)$. In the above calculations we encountered a finite number of q-contractible relevant graphs, so including these graphs and their images under β in S allows S to remain finite. Finally, since every element of S is finite, each graph contains finitely many subgraphs. We include all subgraphs of every element of S to finalize our set S guaranteeing that it is both finite and satisfies the properties described in Lemma 3.2.3.

§ 3.4. Residual Finiteness of $A_{2,3,8}$

Theorem 3.4.1. *$A_{2,3,8}$ is residually finite.*

The goal of this section is to prove Theorem 3.4.1 in a similar manner as the proof of Theorem 3.3.1. To do so, we will again construct a finite set S that will satisfy Lemma 3.2.3.

Proof. The reason that $A_{2,3,8}$ is a special case comes from the very first calculation, $W \otimes_{\bar{U}} W$. This calculation is carried out in Figure 3.8 and results in two connected components. The first connected component is, of course, a copy of W with vertices $(1,1)$, $(2,2)$, $(3,3)$ and $(4,4)$. The second connected component is shown in Figure 3.29, and will be denoted Y_1.

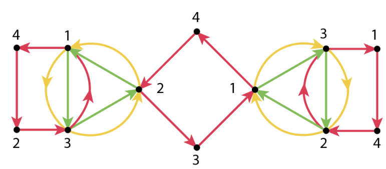

Figure 3.29.: Y_1

So far $S = \{W, Y_1\}$. The next step is to compute $\beta(Y_1)$. This calculation is performed in Figure 3.30.

Since $\beta(Y_1)$ is a new graph, we include $\beta(Y_1)$ in S and proceed by performing the fiber product of every pair in $S = \{W, Y_1, \beta(Y_1)\}$. A full enumeration of all of the connected components that appear in these fiber product computations is shown in Table 3.1 at the end of the section. Readers who would like to perform the fiber product computations should refer to the GitHub link at the end of the paper for a Python program that will aid in these calculations. Here we describe the new graphs that arise from these fiber products. One of the connected components of $W \otimes \beta(Y_1)$ is a new graph that we will denote Y_2 and is shown in Figure 3.31. One of the connected components of $\beta(Y_1) \otimes W$ is a new graph that we will denote Y_3 and is shown in Figure 3.32.

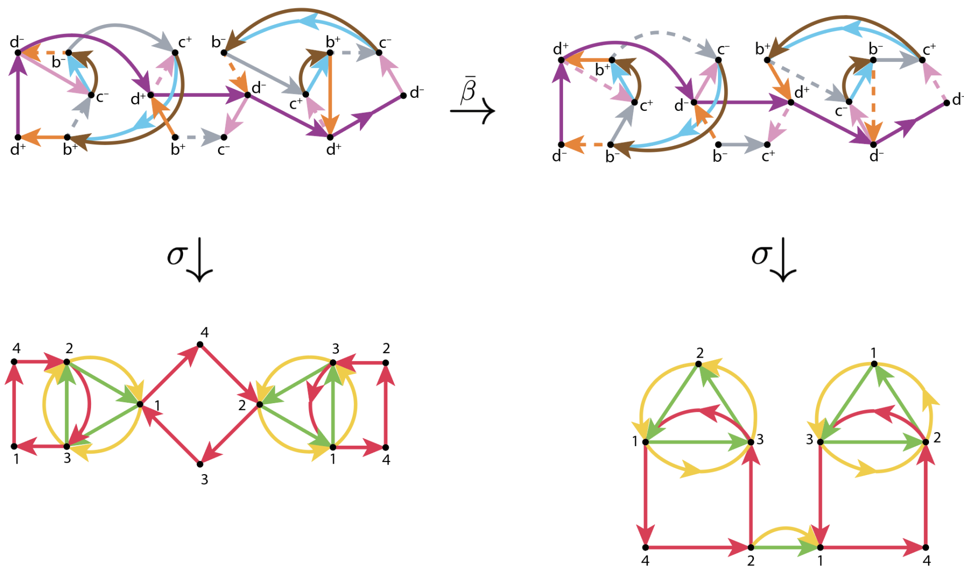

Figure 3.30.: $\beta(Y_1)$

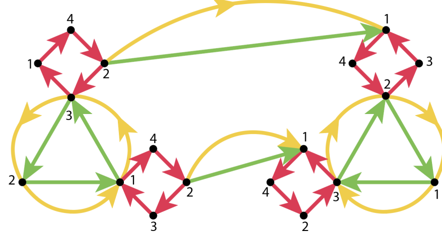

Figure 3.31.: Y_2

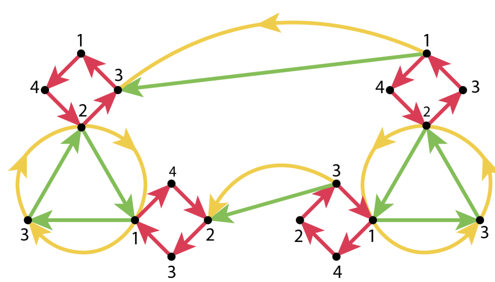

Figure 3.32.: Y_3

Again, we need S to be closed under β, so we must compute $\beta(Y_2)$ and $\beta(Y_3)$. These calculations are shown in Figures 3.33 & 3.34. Luckily, $\beta(Y_2)$ is a rotated copy of Y_2, so $\beta(Y_3)$ is the only other new graph that we must add to S after this step. Table 3.1 contains information about every connected component that arises from the fiber products of every pair of graphs in $S = \{W, Y_1, \beta(Y_1), Y_2, Y_3, \beta(Y_3)\}$. We use $Y_4 = Y_2 \sqcup Y_3 \sqcup \beta(Y_3)$ to streamline computations. Some connected components that arise over the course of these computations are subgraphs of multiple of $W, Y_1, \beta(Y_1), Y_2, Y_3, \beta(Y_3)$. In these situations, a choice was made regarding which column this component is counted in. This choice is arbitrary and does not affect the finiteness of S.

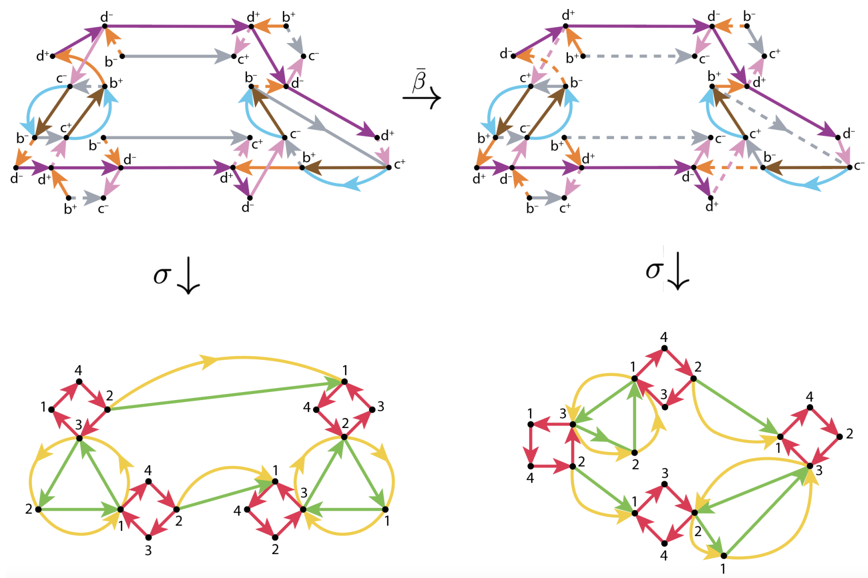

Figure 3.33.: $\beta(Y_2) \cong Y_2$

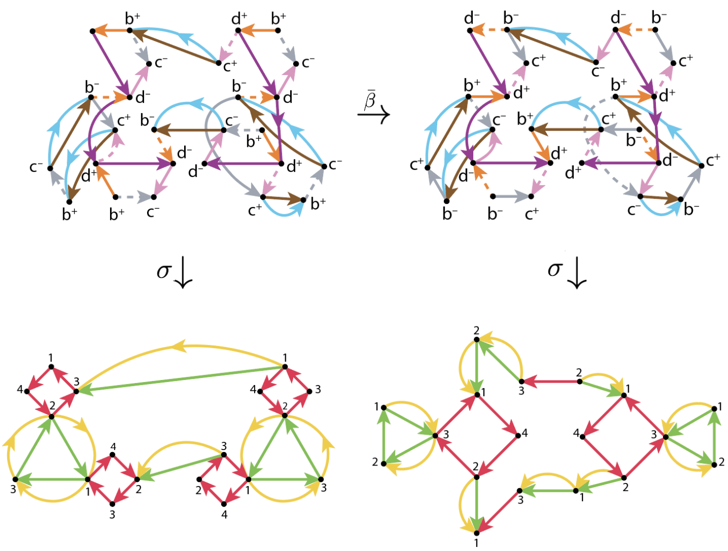

Figure 3.34.: $\beta(Y_3)$

Number of Subgraphs of Connected Components In S							
	W	Y_1	$\beta(Y_1)$	Y_2	Y_3	$\beta(Y_3)$	q-contractible
$W \otimes_{\bar{U}} W$	1	1	0	0	0	0	0
$W \otimes_{\bar{U}} Y_1$	0	2	0	0	0	0	2
$W \otimes_{\bar{U}} \beta(Y_1)$	0	2	1	1	0	0	0
$W \otimes_{\bar{U}} Y_4$	0	0	0	2	2	2	12
$Y_1 \otimes_{\bar{U}} W$	0	2	0	0	0	0	2
$Y_1 \otimes_{\bar{U}} Y_1$	0	2	0	0	0	0	14
$Y_1 \otimes_{\bar{U}} \beta(Y_1)$	0	4	0	2	0	0	6
$Y_1 \otimes_{\bar{U}} Y_4$	0	0	0	2	2	2	72
$\beta(Y_1) \otimes_{\bar{U}} W$	0	2	1	0	1	0	0
$\beta(Y_1) \otimes_{\bar{U}} Y_1$	0	4	0	1	1	0	6
$\beta(Y_1) \otimes_{\bar{U}} \beta(Y_1)$	0	2	1	1	0	0	8
$\beta(Y_1) \otimes_{\bar{U}} Y_4$	0	0	0	3	1	1	68
$Y_4 \otimes_{\bar{U}} W$	0	0	0	2	3	1	12
$Y_4 \otimes_{\bar{U}} Y_1$	0	0	0	2	3	1	72
$Y_4 \otimes_{\bar{U}} \beta(Y_1)$	0	0	0	4	0	1	68
$Y_4 \otimes_{\bar{U}} Y_4$	0	0	0	6	3	4	512

Table 3.1.: An enumeration of the connected components that occur in the fiber products in the construction of S for $A_{2,3,8}$.

Applying the same reasoning as in the end of the proof in Section 3.3.3 to the above table proves that S is finite. By Lemma 3.2.3, this proves that $A_{2,3,8}$ has finite stature with respect to its vertex groups. Therefore, by Theorem 3.1.4, $A_{2,3,8}$ is residually finite. ⌧

Appendix

The program used to analyze the connected components that arise in the $A_{2,3,8}$ fiber products is available at

https://github.com/GreysonPMeyer/Triangle-Artin-Groups

The program is written in Python, and you will need a Python interpreter to run it. These interpreters are available, for free and for almost all platforms, from http://python.org. To build any of the graphs in the code, use the “build_F3" function. To view the connected components of a fiber product, use the “check_fiber" function and follow the prompts. To verify the rows in the table, use the “check_row" function. The $Y_4 \otimes_{\bar{U}} Y_4$ calculation takes a lot of time, so the full dictionary containing the connected components of this calculation is included at the end of the code. Checking the $Y_4 \otimes_{\bar{U}} Y_4$ row of the table can be done by applying the “check_row_Y4xY4" function to the “Y4_x_Y4" dictionary.

Bibliography

[1] Kasia Jankiewicz, *Splittings of triangle artin groups*, Groups Geom. Dyn. 18 (2024), no. 1 (2021), 91–108, available at 2103.01343.

[2] ______, *Residual finiteness of certain 2-dimensional artin groups*, Advances In Mathematics (2022), available at 2006.04717.

[3] ______, *Finite stature in artin groups*, 2023.

[4] Ruth Charney, *An introduction to right-angled artin groups*, Geometriae Dedicata **125** (2007), 141–158, available at math/0610668.

[5] Jon Mccammond, *The mysterious geometry of artin groups*, Winter Braids Lecture Notes **4** (201701), 1–30.

[6] Ruth Charney, *Artin groups of finite type are biautomatic*, Mathematische Annalen **292** (1992Mar), no. 1, 671–683.

[7] ______, *Geodesic automation and growth functions for artin groups of finite type*, Mathematische Annalen **301** (1995Jan), no. 1, 307–324.

[8] Joan S. Birman, K. H. Ko, and J. S. Lee, *A new approach to the word and conjugacy problems in the braid groups*, Advances In Mathematics **139** (1998), 322–353, available at math/9712211.

[9] Jamie Mulholland and Dale Rolfsen, *Local indicability and commutator subgroups of artin groups* (2006), available at math/0606116.

[10] Joseph A. Altobelli, *The word problem for artin groups of fc type*, Journal of Pure and Applied Algebra **129** (1998), no. 1, 1–22.

[11] J. Stilwell and J.P. Serre, *Trees*, Springer Monographs in Mathematics, Springer Berlin Heidelberg, 2012.

[12] Bogdan Nica, *Linear groups - Malcev's theorem and Selberg's lemma*, arXiv e-prints (June 2013), arXiv:1306.2385, available at 1306.2385.

[13] Daan Krammer, *Braid groups are linear* (2004), available at math/0405198.

[14] Stephen J. Bigelow, *Braid groups are linear*, Journal of the American Mathematical Society **14** (2000), no. 2, 471–486, available at math/0005038.

[15] Meenaxi Bhattacharjee, *Constructing finitely presented infinite nearly simple groups*, Communications in Algebra **22** (1994), no. 11, 4561–4589, available at https://doi.org/10.1080/00927879408825087.

[16] Daniel T Wise, *Non-positively curved squared complexes: Aperiodic tilings and non-residually finite groups*, Ph.D. Thesis, 1996 (English). Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2023-02-21.

[17] Marc Burger and Shahar Mozes, *Finitely presented simple groups and products of trees*, Comptes Rendus de l'Académie des Sciences - Series I - Mathematics **324** (1997), no. 7, 747–752. Serie I: Mathematique.

[18] Craig C. Squier, *On certain 3-generator artin groups*, Transactions of the American Mathematical Society **302** (1987), no. 1, 117–124.

[19] John R. Stallings, *Topology of finite graphs.*, Inventiones mathematicae **71** (1983), 551–566.

[20] R.C. Lyndon and P.E. Schupp, *Combinatorial group theory*, Classics in Mathematics, Springer Berlin Heidelberg, 2001.

[21] Allen Hatcher, *Algebraic topology*, Cambridge Univ. Press, Cambridge, 2000.

[22] Luis Paris, *Lectures on artin groups and the $$k(\pi, 1)$$conjecture*, Groups of exceptional type, coxeter groups and related geometries, 2014, pp. 239–257.

[23] Thomas Brady and Jonathan P. McCammond, *Three-generator artin groups of large type are biautomatic*, Journal of Pure and Applied Algebra **151** (2000), no. 1, 1–9.

[24] Daniel T. Wise, *The residual finiteness of negatively curved polygons of finite groups*, Inventiones mathematicae **149** (2002Sep), no. 3, 579–617.

[25] Xiaolei Wu and Shengkui Ye, *Splittings and poly-freeness of triangle artin groups*, 2023.

[26] Mark Feighn and Michael Handel, *Mapping tori of free group automorphisms are coherent*, Annals of Mathematics **149** (1999), 1061–1077, available at math/9905209.

[27] Cameron McA. Gordon, *Artin groups, 3-manifolds and coherence*, 2004.

[28] Daniel T. Wise, *The last incoherent artin group*, 2012.

[29] Jingyin Huang and Daniel T. Wise, *Stature and separability in graphs of groups* (2019), available at 1904.06021.