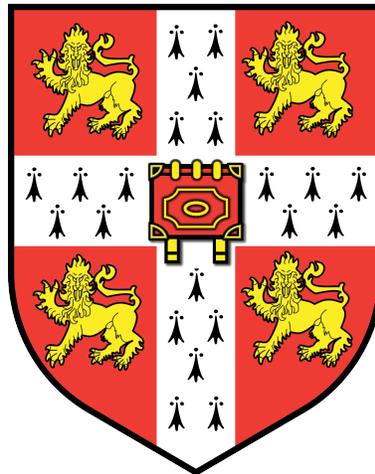


Relating Thompson's group V to graphs of groups and Hecke algebras

Richard Leo Freeland



Trinity College
Department of Pure Mathematics and Mathematical Statistics
University of Cambridge
This dissertation is submitted for the degree of Doctor of Philosophy.

September 2019

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Abstract

This thesis is in two main sections, both of which feature Thompson's group V , relating it to classical constructions involving automorphism groups on trees or to representations of symmetric groups. In the first section, we take \mathcal{G} to be a graph of groups, which acts on its universal cover, the Bass-Serre tree, by tree automorphisms. Brownlowe, Munday, Pask, Spielberg and Thomas constructed a C^* -algebra for a graph of groups, written $C^*(\mathcal{G})$, which bears many similarities to the C^* -algebra of a directed graph G . Inspired by the fact that directed graph C^* -algebras $C^*(G)$ have algebraic analogues in Leavitt path algebras $L_K(G)$, we define a Leavitt graph-of-groups algebra $L_K(\mathcal{G})$ for \mathcal{G} . We extend Leavitt path algebra results to $L_K(\mathcal{G})$, including uniqueness theorems describing homomorphisms out of $L_K(\mathcal{G})$, and establish a wider context for the algebras by showing they are Steinberg algebras of a particular étale groupoid. Finally we show that certain unitaries in $L_K(\mathcal{G})$ form a group we can understand as a variant of Thompson's V , combining features of both Nekrashevych-Röver groups and Matui's topological full groups of one-sided shifts. We prove finiteness and simplicity results for these Thompson variants. The latter section of this thesis turns to representation theory. We briefly state some results about representations of V (due to Dudko and Grigorchuk) which we generalize to the new family of Thompson groups, including a discussion of representations of finite factor type and Koopman representations. Then, we describe how one would try to construct a Hecke algebra for V , built from copies of the Iwahori-Hecke algebra of \mathfrak{S}_n in a way inspired by how V can be constructed from copies of the symmetric group. We survey attempts to construct this and demonstrate what we believe to be the closest possible analogue to the \mathfrak{S}_n theory. We discuss how this construction could prove useful for understanding further representation theory.

Dedication

In memory of Ben Elliott.

Acknowledgements

Writing this thesis would not have been possible without the help of many people who deserve appreciation here. First, I would like to thank my supervisor Chris Brookes for introducing me to so many of the topics studied here; for his deep insight into which problems should be solvable; and for his continued patience and encouragement, without which I would have given up on research long ago.

Thank you to my fellow algebra PhD students: Amit, Andreas, Nicolas and Stacey, for their friendship, the helpful mathematical conversations, and for their example to show me that finishing a PhD is possible. Thank you also to the EPSRC Research Council for my funding and to Trinity College for my accommodation during the last four years.

Finally I would like to thank my family - Kevin, Sue, Alex, and Katy - whose love and faith has been invaluable.

Chapter 1

Introduction

The majority of this thesis is a study of groups acting on trees, or on their ends. One of the main purposes is to unite two different well-studied groups defined in terms of trees. The first is the fundamental group of Bass-Serre theory, which is an automorphism group of a tree defined in terms of a graph of groups \mathcal{G} . The second is Thompson's group V and its many relatives, a group defined by permuting the ends of a binary tree, with many unusual properties. We will produce a family of groups combining features of both of these objects, and analyse it to see how properties of V and \mathcal{G} transfer to these groups. In the final section of this thesis, we study representation theory of Thompson's group. We show how to generalize some known theorems about representations of V to the new family of Thompson-type groups. Finally, we discuss how the similarities between V and symmetric groups \mathfrak{S}_n might let us say more about representations of V , in particular trying as far as possible to extend a Hecke algebra construction from \mathfrak{S}_n to V .

In this introduction we define the main objects of our consideration. We first describe the theory of automorphism groups of trees, which Bass and Serre described as fundamental groups of graphs of groups. Next, we exhibit a variety of related algebraic objects that act on paths in a directed graph, beginning with the Leavitt path algebra and finishing with Thompson's group V and its variants. We will combine both Leavitt path algebras and Thompson's groups with graphs of groups to form new objects.

Notation: We begin by fixing some notation on directed graphs. A *directed graph* Γ is a quadruple $(\Gamma^0, \Gamma^1, s, t)$ where Γ^0 is a set of *vertices*, Γ^1 is a set of

edges, and s and t are the *source* and *target* functions from Γ^1 to Γ^0 . We say Γ is finite if both the vertex and edge sets are finite, and say Γ is locally finite if each of $s^{-1}(v)$ and $t^{-1}(v)$ are finite, for all $v \in \Gamma^0$.

We say that a *path* in a directed graph Γ is either a vertex, or a sequence $e_1 e_2 \dots e_n$ of edges with $s(e_i) = t(e_{i+1})$ for $1 \leq i \leq n - 1$ (so our paths are read right-to-left, like morphisms being composed). An example is shown in Figure 1.1. We write $\ell(p)$ for the length of the path p , which is 0 if p is a vertex and n if $p = e_1 e_2 \dots e_n$. Finally we extend the source and target maps to paths: we say that the path $p = e_1 e_2 \dots e_n$ has source $s(p) = s(e_n)$ and target $t(p) = t(e_1)$, and if p is a vertex v , then $s(p) = t(p) = v$. Later we'll also consider infinite paths, which are infinite sequences of edges $e_1 e_2 e_3 \dots$ obeying the same condition $s(e_i) = t(e_{i+1})$. The target of an infinite path is defined as $t(e_1)$, but the source is not defined.



Figure 1.1: An example of a path

We will write Γ^* for the set of all (finite) paths, and Γ^n for the set of all paths $e_1 e_2 \dots e_n$, of length n . We write Γ^ω for the set of infinite paths. If K is a field, the path algebra $K\Gamma$ is then the K -algebra with Γ^* as basis, and with multiplication defined on the basis by concatenation of paths, wherever defined (i.e. for paths p and q , $p \cdot q$ is pq whenever pq is a path and 0 otherwise, when $s(p) \neq t(q)$.)

In this thesis, we will always specify when a graph is directed. The word ‘graph’ alone will mean a graph in the sense of Serre (which is a directed graph with some additional structure), as given in the next section.

1.1 Bass-Serre theory

We begin by studying automorphism groups of trees. The most important work in this area was developed by Serre and explained clearly in his book [51], and later developed by Bass in [8]. Our exposition is a combination of [51] and [13].

1.1.1 Basic definitions

Here we use Serre’s notion of a graph. It is a particular kind of directed graph:

Definition 1.1.1. A graph Γ is a directed graph $(\Gamma^0, \Gamma^1, s, t)$ equipped with a function $- : \Gamma^1 \rightarrow \Gamma^1$, written $e \mapsto \bar{e}$, such that $s(\bar{e}) = t(e), t(\bar{e}) = s(e)$, and $e \neq \bar{e}$ but $e = \bar{\bar{e}}$.

This means that every edge e of a graph is directed, and comes with a reverse \bar{e} ; the edges can be partitioned into disjoint pairs $\{e, \bar{e}\}$. We allow multiple edges between the same pair of vertices, and we allow loops whose source and target is the same edge. It may help to think of Serre's graphs as undirected graphs in the usual sense, but where every edge has been replaced by a pair of edges, oppositely directed.

For $v \in \Gamma^0$, we define the *degree* (or *valence*) of v to be $|s^{-1}(v)|$, the cardinality of the set of edges with source v . We say the graph Γ is locally finite if all valencies are finite, and will usually work with locally finite graphs in this thesis. Since $s(e) = t(\bar{e})$ for all edges e , we have $|t^{-1}(v)| = |s^{-1}(v)|$, so locally finite graphs also have finitely many edges with target v for any vertex v .

When we draw graphs, we will normally draw only one direction of each edge, and we may or may not specify which direction the edge e goes in if we don't include \bar{e} : for example, we might draw the same graph in each of the three ways shown in Figure 1.2. We say that an *orientation* of graph Γ is a choice of one edge from each pair $\{e, \bar{e}\}$, and write an orientation as $E_+ \subset \Gamma^1$. For example, the upper-right diagram of Figure 1.2 shows a particular orientation.

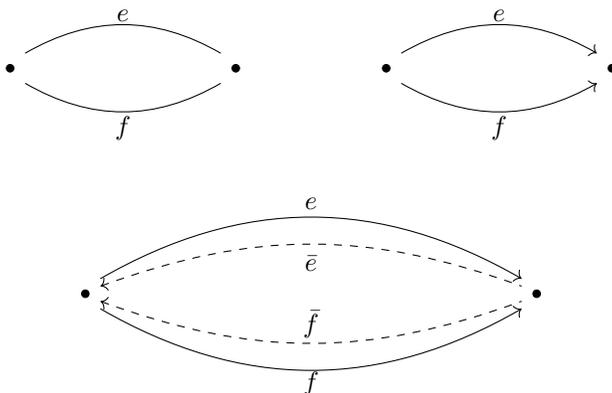


Figure 1.2: A graph with two pairs of edges

The graph Γ can be topologized as a *CW*-complex: there is a 0-cell (i.e. a point) for each vertex $v \in \Gamma^0$, and for each pair $\{e, \bar{e}\}$ of an edge and its reverse, we add in a copy of the unit interval $[0, 1]$ by identifying 0 with $s(e)$ and 1 with

$t(e)$. We call this the realization of Γ . It lets us define topological notions on graphs: for example, we say a graph is connected if its realization is a connected topological space, and we define the fundamental group of a connected graph as the fundamental group of its realization (at any point). The topology described is metrizable, by extending the usual metric from $[0, 1]$.

Paths: Recall the definition of paths and infinite paths in a directed graph Γ as sequences $e_1e_2e_3\dots$ of edges of Γ . We say that a path (finite or infinite) is *without backtracking* if $e_i \neq \bar{e}_{i+1}$ for any i . In contrast with the directed case, we write Γ^n for the set of all paths of length n *without backtracking* (this agrees with the notation Γ^0 for vertices, which are paths of length 0, and Γ^1 for edges, which are paths of length 1). Also, we define $\Gamma^* = \cup_{n=0}^{\infty} \Gamma^n$, the set of all (finite) paths without backtracking, and we write Γ^ω for the set of infinite paths without backtracking.

Trees: Let Γ be a connected graph. We define a *closed path* in Γ to be a path ρ of Γ without backtracking such that $s(\rho) = t(\rho)$, and we say the closed path $\rho = e_1e_2\dots e_n$ is a *cycle* if $s(e_i) \neq s(e_j)$ for any $1 \leq i, j \leq n$, with $i \neq j$. That is, a cycle is a closed path that returns to its source only once; other closed paths are formed by concatenating cycles. A *tree* is a connected graph without cycles. In a tree, any two vertices are connected by a unique path without backtracking.

It is clear that connected subgraphs of a tree T are trees; we call them subtrees of T . In particular, if ρ is a path, then the set of edges of ρ and their source and target vertices forms a subtree. If S_1 and S_2 are subtrees of T , and v, w are vertices of both S_1 and S_2 , then the unique path from v to w in T must lie in both S_1 and S_2 . Thus, $S_1 \cap S_2$ contains a path from v to w , and so is connected. This proves that a non-empty intersection of two subtrees is a subtree.

If T is an infinite tree, we define an equivalence relation on the infinite paths of T by saying they are equivalent if their intersection is also an infinite path. In other words, infinite paths $e_1e_2e_3\dots$ and $f_1f_2f_3\dots$ are equivalent if there exist $N \in \mathbb{N}, k \in \mathbb{Z}$ such that $e_n = f_{n+k}$ for all $n \geq N$. We write ∂T for the equivalence classes under this relation, and call the equivalence classes the *ends* of T .

1.1.2 Defining and classifying automorphisms

We will be interested in automorphisms of trees. A morphism of graphs, from graph Γ to graph Δ , is a pair of functions $f_0 : \Gamma^0 \rightarrow \Delta^0, f_1 : \Gamma^1 \rightarrow \Delta^1$, which commute with the source, target and reversal maps. An automorphism is a morphism from a graph to itself which is bijective on both vertices and edges. For example, let \mathcal{C}_n be the graph consisting of a single cycle of length n , with vertices and edges as shown in figure 1.3. Then the automorphism group of \mathcal{C}_n is the dihedral group D_{2n} - there are n rotations which send e_i to e_{i+k} for some $k \bmod n$, and n reflections which fix a vertex v_i and send e_k to \bar{e}_{i-k} , indices taken mod n .

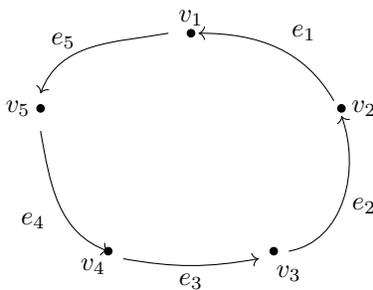


Figure 1.3: The cycle \mathcal{C}_5

Now we can state Tits' classification of tree automorphisms.

Proposition 1.1.2 (Tits, [53]). *Let α be an automorphism of a tree T . Then precisely one of the following holds:*

1. α fixes a vertex v of T (it might fix more than one vertex).
2. α interchanges the two ends $s(e)$ and $t(e)$ of an edge e .
3. There exists a unique doubly-infinite sequence of edges $\dots e_{-1}e_0e_1e_2\dots$, such that $s(e_i) = t(e_{i+1})$ and $e_{i+1} \neq \bar{e}_i$, on which α acts by a translation by some non-zero integer k , so that $e_i \mapsto e_{i+k}$. In particular, α fixes two ends, namely $e_1e_2\dots$ and $\bar{e}_{-1}\bar{e}_{-2}\dots$.

In particular, the identity automorphism α satisfies the first condition.

We remark that the first and third of these cases preserve some orientation of T : indeed, if α fixes v , then we orient each edge towards v , and if α translates a doubly-infinite path, we orient the edges of the path towards one end of it, and

all other edges towards the path. The second case cannot preserve an orientation (since it sends e to \bar{e}). So we make the definition that α acts *without inversion* if it preserves some orientation of T .

Sometimes we will want to restrict to the case of automorphisms without inversion. This can be done without loss of generality, by the following construction. We form the *barycentric division* of Γ , written Γ' , as follows: the vertex set of Γ' is Γ^0 with an extra vertex v_e added for each pair of edges $\{e, \bar{e}\}$. For each edge e of Γ , Γ' has an edge from $s(e)$ to v_e and an edge from $t(e)$ to v_e , and their reverses. An example is in Figure 1.4 (drawn with one edge of each pair $\{e, \bar{e}\}$). It's easy to see that this construction just adds a vertex in the middle of each edge.

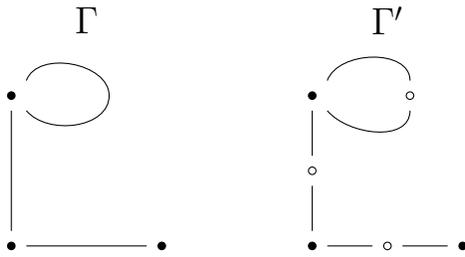


Figure 1.4: A barycentric subdivision

If α acts on T with inversion, by swapping $s(e)$ and $t(e)$, then it acts on T' without inversion, stabilizing the barycentre v_e . So we can always work with automorphism groups acting without inversion so long as we're willing to pass to a barycentric subdivision.

If every element of a group G acts without inversion on T , then there's an orientation of T preserved by G : indeed, no G -orbit of edges of T can contain both an edge and its reverse, so we can just choose a consistent orientation for each orbit independently.

Finally, suppose G is a group acting on a graph X by automorphisms. Then we can define a quotient graph $G \backslash X$ in the obvious manner: $(G \backslash X)^0$ is the quotient of X^0 by the G -action (the set of orbits of vertices), and $(G \backslash X)^1$ is the quotient of X^1 ; the source, target and reversal maps are defined in the quotient, because they are preserved by G .

1.1.3 Groups acting on trees

We now give Serre's general theory of groups acting on trees. Again we're following [51] and [13] for this explanation. All actions will be by tree automorphisms, without inversion.

Free actions

The first important result on groups acting on trees is the following theorem of Serre.

Theorem 1.1.3 ([51] I.3 Theorem 4). *Let G be a group that acts on a tree T . We say that G acts freely if it acts without inversion and no non-identity element of G fixes any vertex of T (so every element of G is of the third type in Tits' classification). If G acts freely, then G is a free group. Conversely, if G is free on the set S , then the Cayley graph of G with respect to the generators S is a tree on which G acts freely.*

We remark that Schreier's theorem is an immediate corollary:

Corollary 1.1.4 (Nielsen-Schreier). *A subgroup of a free group is free.*

The proof of Theorem 1.1.3 also finds a subset of G which generates G freely. So the case of groups acting with trivial vertex stabilizer is completely understood. Going forward, we'll understand other groups acting on trees in terms of their vertex stabilizers.

Graphs of groups

The general case of a group acting on a tree is studied using graphs of groups.

Definition 1.1.5. *A graph of groups \mathcal{G} is a tuple $\mathcal{G} = (\Gamma^0, \Gamma^1, G, \alpha)$ where:*

- Γ^0 and Γ^1 are the vertex and edge sets of a graph Γ , in the sense of Definition 1.1.1. In this thesis, we assume that the graph Γ is connected.
- G is a function assigning to each vertex $v \in \Gamma^0$ a group G_v , and to each edge $e \in \Gamma^1$ a group $G_e = G_{\bar{e}}$.
- α is a set of injective group homomorphisms $\alpha_e : G_e \hookrightarrow G_{t(e)}$, for each edge e .

We say Γ is the underlying graph of \mathcal{G} ; we do not include the source, target and reversal maps of Γ in the notation. For $v \in \Gamma^0$, we will write 1_v for the identity of G_v , or more commonly just 1 when v is clear from context.

The two most important examples occur when Γ has just one pair of edges $\{e, \bar{e}\}$. They are as follows (see also [13]:)

1. An *edge of groups* is a graph of groups as shown in Figure 1.5, with two vertices and one pair of edges.



Figure 1.5: An edge of groups.

The underlying graph has two vertices v and w and one pair of edges $\{e, \bar{e}\}$, where $s(e) = y, r(e) = x$. The embeddings are $\alpha_e : G_e \hookrightarrow G_w$ and $\alpha_{\bar{e}} : G_e \hookrightarrow G_v$.

2. A *loop of groups* is a graph of groups as shown in Figure 1.6, with one vertex and one pair of edges.

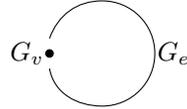


Figure 1.6: A loop of groups

The underlying graph has one vertex v and one pair of edges $\{e, \bar{e}\}$. There are two homomorphisms α_e and $\alpha_{\bar{e}}$ embedding G_e into G_v .

We now return to the general case. For each edge $e \in \Gamma^1$, choose a set Σ_e of left coset representatives for the subgroup $\alpha_e(G_e)$ of $G_{t(e)}$. Say that the graph of groups \mathcal{G} is *locally finite* if its underlying graph is locally finite, and if each Σ_e is finite. We'll work with locally finite graphs of groups from now on.

Suppose T is a locally finite tree on which a group G acts without inversion. We describe the situation with a graph of groups, $\mathcal{G}(G, T)$, defined as follows. The underlying graph Γ of $\mathcal{G}(G, T)$ will be the quotient graph $G \backslash T$, where p is the projection $p : T \rightarrow G \backslash T$. For each $v \in \Gamma^0$, let G_v be the stabilizer of any vertex $\hat{v} \in p^{-1}(v)$. Similarly for $e \in \Gamma^1$, let G_e be the stabilizer of any edge $\hat{e} \in p^{-1}(e)$. Notice that as abstract groups, G_v and G_e don't depend on the

choice of preimage, since all the possible stabilizers are conjugate subgroups of G .

Finally we must define the embeddings α_e for $e \in \Gamma^1$. Let $e \in \Gamma^1$ have target v . We seek an embedding from $\text{Stab}(\hat{e})$ to $\text{Stab}(\hat{v})$. Clearly the stabilizer of an edge is the intersection of the stabilizers of its source and target vertices. By definition of the quotient graph, $t(\hat{e})$ and \hat{v} lie in a common G -orbit, so there exists $g \in G$ such that $gt(\hat{e}) = \hat{v}$, and then:

$$gG_e g^{-1} = g\text{Stab}(\hat{e})g^{-1} \leq g\text{Stab}_G(t(\hat{e}))g^{-1} = \text{Stab}(\hat{v}) = G_v.$$

So we define (cf [13]) $\alpha_e : G_e \rightarrow G_{t(e)}$ by $h \mapsto ghg^{-1}$. This is well-defined, and completes the definition of a graph of groups $\mathcal{G}(G, T)$.

1.1.4 Example: $SL_2(\mathbb{Z})$

A good example of this theory is provided by the group $G = SL_2(\mathbb{Z})$. Recall

$$SL_2(\mathbb{Z}) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} : a, b, c, d \in \mathbb{Z}, ad - bc = 1 \right\}.$$

As a subgroup of $SL_2(\mathbb{C})$, G acts on the extended complex plane $\mathbb{C} \cup \{\infty\}$ by Möbius maps, where the matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ corresponds to the transformation $z \mapsto \frac{az+b}{cz+d}$. The subgroup G preserves the upper half plane, on which it acts with fundamental domain:

$$\Delta = \left\{ z \in \mathbb{C} : |z| \geq 1, |\Re(z)| \leq \frac{1}{2}, \Im(z) > 0 \right\}.$$

This set can also be defined as

$$\Delta = \{z \in \mathbb{C} : d_H(z, 2i) \leq d_H(z, \theta(2i)) \text{ for all } \theta \in G\};$$

here, d_H is the hyperbolic distance, and the definition says that Δ consists of points z where $2i$ is the closest point in its G -orbit to z . See Figure 1.7 for a picture that summarizes this section and see [23] for details, which gives a good short account of the theory in an online expository paper.

Now consider the action of $SL_2(\mathbb{Z})$ on the circular arc between i and $\omega = \frac{1+\sqrt{-3}}{2}$. The orbit of this arc consists of other arcs of circles or line segments, and is shown in figure 1.7, along with the fundamental domain Δ . In particular, one can verify that the orbit forms (the realization of) a tree. Every vertex of

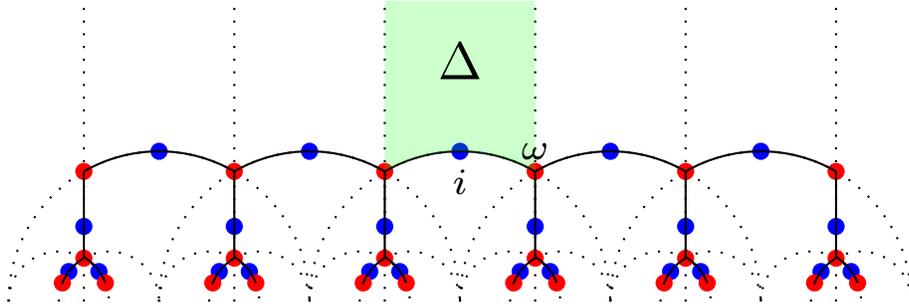


Figure 1.7: A tree that $SL_2(\mathbb{Z})$ acts on.

the tree has valence 2 or 3 - the valence 2 vertices form the orbit of i and the valence 3 vertices form the orbit of ω . This gives a tree T on which G acts without inversion.

Finally we identify the quotient $G \backslash T$. By construction of T , this is a graph with a single edge (and its reverse). It's easy to verify that the stabilizer in G of i is cyclic of order 4, generated by $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$; whereas the stabilizer of ω is cyclic of order 6, generated by $\begin{pmatrix} 0 & 1 \\ -1 & 1 \end{pmatrix}$. Their intersection is $\pm \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$. So the quotient graph of groups is as shown in Figure 1.8.

$$C_4 \bullet \xrightarrow{C_2} \bullet C_6$$

Figure 1.8: The quotient graph of groups for $SL_2(\mathbb{Z})$

We don't need to specify the embeddings because there is only one embedding possible.

1.1.5 The fundamental group

The aim here is to reverse the construction above and construct a group acting on a tree from a graph of groups. The group will be called the fundamental group, and will generalize the topological notion of the fundamental group of a graph. The tree will be called the Bass-Serre tree and will be analogous to a universal cover. We take most of the definitions from [13], beginning with definitions of paths in a graph of groups that generalize the notion of a path in a graph.

Definition 1.1.6. [13] *Let \mathcal{G} be a graph of groups, with underlying graph Γ , and transversals Σ_e . Assume that each Σ_e includes the identity $1_{t(e)}$.*

- A \mathcal{G} -word ρ is either an element g_1 of a vertex group G_{v_1} , or a sequence

$$g_1 e_1 g_2 e_2 \dots g_n e_n \text{ or } g_1 e_1 \dots g_n e_n g_{n+1},$$

where $s(e_i) = t(e_{i+1})$, and $g_i \in G_{t(e_i)}$, $g_{n+1} \in G_{s(e_n)}$ wherever these make sense. (It may help to imagine a \mathcal{G} -word as a walk around the graph, during which you record what edges you walk along, and add in a group element at each vertex.) The source of ρ is $s(e_n)$ and its target is $t(e_1)$ (or if $\rho = g_1 \in G_{v_1}$, then the source and target are both v_1). The length of ρ is the number n of edges in its expression (so the length of a vertex group element is 0). We write the length $\ell(\rho)$.

- A \mathcal{G} -word is reduced if for all $1 \leq i \leq n$, then $g_i \in \Sigma_{e_i}$, and if $e_i = \bar{e}_{i+1}$, then $g_{i+1} \neq 1_{t(e_{i+1})}$. The final group element g_{n+1} is free to be anything in a reduced word; in particular, any word of length 0 is reduced.
- A \mathcal{G} -path is a reduced \mathcal{G} -word of the form $g_1 = 1$ or $g_1 e_1 \dots g_n e_n$. We write \mathcal{G}^n for the set of all length n \mathcal{G} -paths, and \mathcal{G}^* for the set of all \mathcal{G} -paths of any length. For $v \in \Gamma^0$ we write $v\mathcal{G}^*$ (similarly, $v\mathcal{G}^n$) for the set of all \mathcal{G} -paths with target v (or of length n with target v).

Following [13] and [51] we define one auxiliary group before the fundamental group.

Definition 1.1.7. Let \mathcal{G} be a graph of groups with underlying graph Γ . The path group of \mathcal{G} , written $F(\mathcal{G}, \Gamma)$ is the group generated by all vertex groups G_v and symbols for each edge $e \in \Gamma^1$, with the relations:

- Reversal: $e^{-1} = \bar{e}$
- Conjugation: $\alpha_e(g)e = e\alpha_{\bar{e}}(g)$ for all $g \in G_e$.

We remark that elements of the path group can consist of arbitrary products of edges and vertex group elements; they don't all correspond to \mathcal{G} -words.

As an example, let \mathcal{G} be a loop of groups, with vertex group G ; let $\alpha_e(G_e) = H$ and $\alpha_{\bar{e}}(G_e) = H'$ be the two subgroups of G that are images of the embeddings of G_e into G . Then $F(\mathcal{G}, \Gamma)$ is generated by G and a single extra element e that conjugates H to H' ; in other words, it is a HNN extension of G (cf Chapter 1 of [51]).

Importantly, \mathcal{G} -words have representatives in the path group. Indeed, if we represent a \mathcal{G} -path by g_1 , by $g_1 e_1 \dots g_n e_n$, or by $g_1 e_1 \dots g_n e_n g_{n+1}$, then

we identify this with the element of the path group written the same way. Different \mathcal{G} -words can give the same element of the path group (such as $1e1\bar{e}$ giving the identity), but it can be shown that every reduced \mathcal{G} -word gives a different element. This is Theorem 5.2 of [51], and it has the most technical and lengthy proof in this theory.

We use these representatives to define the fundamental group. Write $\pi[v, w]$ for the image in the path group of the set of all reduced \mathcal{G} -words with source w and target v . There are two definitions, which are equivalent.

Definition 1.1.8. *Let \mathcal{G} be a graph of groups with underlying graph Γ . Then the fundamental group, $\pi_1(\mathcal{G})$, may be defined in either of the following ways:*

1. *Fix any vertex v of Γ . Observe that $\pi[v, v]$ forms a subgroup of $F(\mathcal{G}, \Gamma)$. We define the fundamental group of \mathcal{G} rooted at v to be $\pi[v, v]$, and write it $\pi_1(\mathcal{G}, v)$.*
2. *Choose a maximal subtree T of Γ (such subtrees exist, and meet every vertex, in any connected graph). We define the fundamental group of \mathcal{G} relative to T , written $\pi_1(\mathcal{G}, T)$, to be $F(\mathcal{G}, \Gamma)$ quotiented out by the extra relations $e = 1$ for all edges $e \in T^1$.*

The abstract groups $\pi_1(\mathcal{G}, v)$ and $\pi_1(\mathcal{G}, T)$ are isomorphic, and independent of the choice of v and T . Moreover, the projection from $F(\mathcal{G}, \Gamma)$ onto $\pi_1(\mathcal{G}, T)$ restricts to an isomorphism on $\pi[v, v]$ (this is [51] Proposition 4.20). We give three examples:

- **Trivial graphs of groups:** If all the vertex and edge groups of \mathcal{G} are trivial, then $\pi_1(\mathcal{G})$ is the usual fundamental group (of the realization of the graph Γ), since it is generated by loops at a vertex with the relations $e\bar{e} = 1$.
- **Edges of groups:** Let \mathcal{G} be an edge of groups. The underlying graph is already a tree, and so $T = \Gamma$ is a maximal tree. Thus $\pi_1(\mathcal{G}, T)$ is generated by the two vertex groups G_v and G_w and the edge e , with the relations $\alpha_e(g)e = e\alpha_{\bar{e}}(g)$ and $e = 1$. In other words, it is an amalgamated free product $G_v *_H G_w$, where the common subgroup H is $\alpha_e(G_e) \cong \alpha_{\bar{e}}(G_e)$.
- **Loops of groups:** Here there is only one vertex and one edge (and its reverse). So $\pi[v, v]$ is all of $F(\mathcal{G}, \Gamma)$ (and a maximal tree consists of the single vertex v). Thus the fundamental group and the path group are the same in this case, and we've seen that this is a HNN extension.

More complicated graphs of groups can be built up one edge at a time. Adding an edge and a new vertex corresponds to doing an amalgamated free product, and joining two existing vertices corresponds to an HNN extension.

1.1.6 The universal cover

Finally we construct the *Bass-Serre tree*, or *universal cover*, for a graph of groups \mathcal{G} . This is a tree $T_{\mathcal{G},v}$ with an action of $\pi_1(\mathcal{G}, v)$, such that the quotient by $\pi_1(\mathcal{G}, v)$ recovers the graph Γ , and the vertex and edge groups of \mathcal{G} are corresponding vertex and edge stabilizers in $T_{\mathcal{G},v}$. This is done as follows ([13] Definition 2.13):

Definition 1.1.9. Fix vertex v of graph of groups \mathcal{G} . The vertex set of the universal cover is:

$$T_{\mathcal{G},v}^0 = \{\gamma G_x : \gamma \in \pi[v, x], x \in \Gamma^0\}.$$

There is an edge $f \in T_{\mathcal{G},v}^1$ with $s(f) = \gamma' G_{x'}$ and $t(f) = \gamma G_x$ if and only if $\gamma^{-1}\gamma' \in G_{x'e}G'_x$. The fundamental group $\pi_1(\mathcal{G}, v)$ acts on $T_{\mathcal{G},v}$ via multiplication in $F(\mathcal{G}, \Gamma)$, since if $\alpha \in \pi[v, v]$ and $\gamma \in \pi[v, x]$ then $\alpha\gamma \in \pi[v, x]$.

We summarize the point of this construction in the Fundamental Theorem of Bass-Serre theory, as stated in [13].

Theorem 1.1.10. Let \mathcal{G} be a graph of groups, with underlying graph Γ ; let $v \in \Gamma^0$. Then $T_{\mathcal{G},v}$ is a tree, $\pi_1(\mathcal{G}, v)$ acts on it without inversion, and the graph of groups $\mathcal{G}(\pi_1(\mathcal{G}, v), T_{\mathcal{G},v})$ is isomorphic to \mathcal{G} . Conversely, suppose the group G acts without inversion on a tree T , and let the associated graph of groups be $\mathcal{G}(G, T)$. Then for any $v \in \Gamma^0$ we have that $\pi_1(\mathcal{G}(G, T), v) \cong G$ as groups, and $T_{\mathcal{G}(G, T), v} \cong T$ as graphs with G -action.

*SL*₂(\mathbb{Z}), **again:** We return to the example of *SL*₂(\mathbb{Z}) to show how this works in practice. We have seen that *SL*₂(\mathbb{Z}) acts on a (2,3)-regular tree T , with quotient graph of groups:

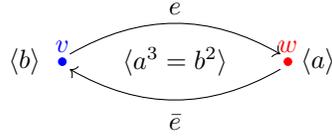
$$C_4 \bullet \xrightarrow{C_2} \bullet C_6$$

We have also seen that the fundamental group of an edge of groups is an

amalgamated free product, so we recover the well-known isomorphism:

$$SL_2(\mathbb{Z}) \cong C_4 *_{C_2} C_6.$$

Conversely, we will show how to rebuild the tree T from the finite graph of groups. Say that C_6 is generated by a and C_4 by b (so that C_2 is generated by $a^3 = b^2$). We add labels to the graph of groups to help make the notation clear:



Now we can construct the Bass-Serre tree $T_{\mathcal{G},v}$. The definition given tells us everything we need to know to construct it, so we just draw (a finite portion of) the result in Figure 1.9.

We see that this is isomorphic to the $(2, 3)$ -regular tree that we found earlier as a subset of \mathbb{C} . The action of $\pi_1(\mathcal{G}, v) = \langle b, \bar{e}ae \rangle$ is easy to work out: for example, b effects a reflection across G_v , whilst $\bar{e}ae$ fixes the vertex $\bar{e}G_w$. The subgroup C_2 (generated by b^2) acts trivially on the tree; this also holds for the action on \mathbb{C} , because $SL_2(\mathbb{Z})$ acts on its tree only after factoring through $PSL_2(\mathbb{Z})$.

Free actions, again: Finally, we remark that Theorem 1.1.3 is a corollary of the fundamental theorem. Indeed, suppose that G acts freely on a graph T . In the graph of groups $\mathcal{G}(G, T)$, all vertex and edge groups must be trivial, because

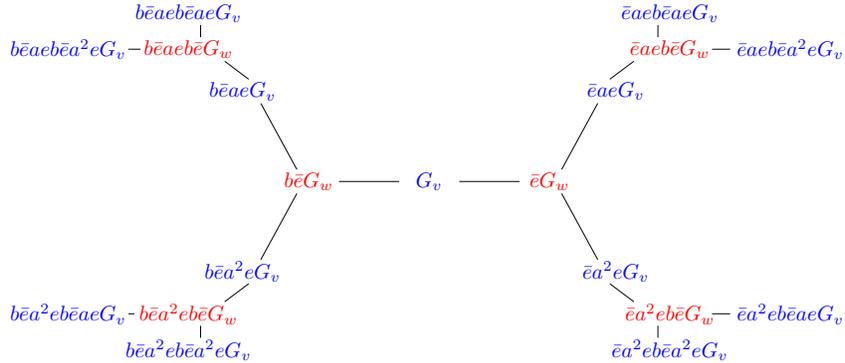


Figure 1.9: The Bass-Serre tree for $SL_2(\mathbb{Z})$, again

no non-trivial subgroup of G fixes any vertex or edge of T . Thus $\mathcal{G}(G, T)$ is a graph of trivial groups. Its fundamental group is then just the fundamental group of the underlying graph, and it's a standard result in topology that this is a free group.

1.2 Leavitt path algebras

1.2.1 Overview

In the rest of the introduction, we will introduce various closely related constructions, all of which can be defined by having them act on paths in a directed graph. Figure 1.10 shows the objects we will be defining and the relations between them. A good recent summary of these constructions is [21], which is particularly good for étale groupoids and Steinberg algebras.

We will first explain this theory for Leavitt path algebras, which relate directly to paths in a directed graph, and then give the more general theory of topological groupoids. One thing we will do in the thesis is to study a particular case of the topological groupoid construction in some detail. We will begin with a C^* -algebra for graphs of groups defined in [13] and show how to define all the other objects in Figure 1.10 for this C^* -algebra.

We begin with an overview of the theory of Leavitt path algebras. There are many aspects to this theory, and we can only give some of them. Much fuller versions of the theory are given in the overview paper [1] and especially in the book [3]. We will be largely following [1] in this introduction. However, our paths will be written right-to-left, whereas left-to-right is more common for Leavitt path algebra conventions. This is done in order to agree with [13], because we will be adapting constructions from that paper most closely. We will briefly describe the history of Leavitt path algebras and sketch their links to C^* -algebras. We will then quote two key theorems about homomorphisms (or ideals) of Leavitt path algebras, which have C^* analogues, and which we will later state generalizations of to related algebras.

1.2.2 History of Leavitt path algebras

Invariant basis number: This history follows the exposition in [54], but we spell out some results in more detail. We begin with some classical ring theory. Say that unital ring R has *invariant basis number* if, whenever the

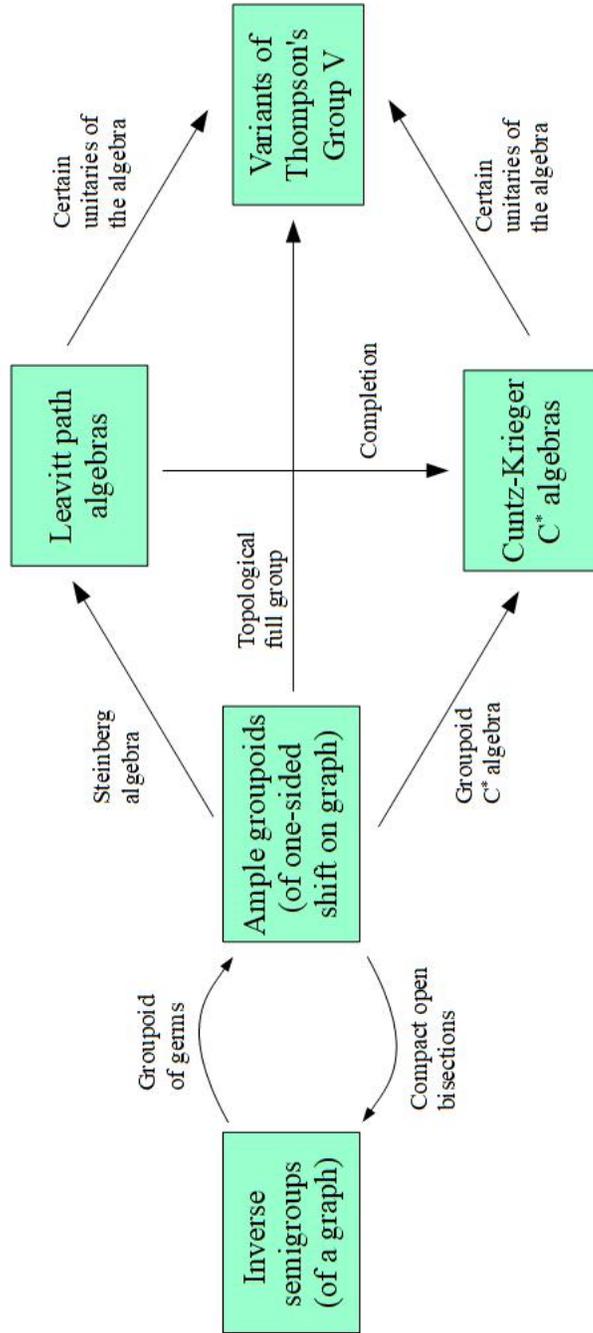


Figure 1.10: Relations between various objects acting on paths

left R -modules R^i and R^j are isomorphic (for $i, j \in \mathbb{N}$) then $i = j$. Most familiar examples of rings, such as all commutative or Noetherian rings, have this property; we want to describe the ways this property can fail.

Suppose R does not have invariant basis number, and that R^m and R^n are isomorphic (for $m \neq n$). Then there exist matrices $X : R^m \rightarrow R^n$ and $Y : R^n \rightarrow R^m$ such that $XY = I_n$ and $YX = I_m$ are identity matrices. In particular, if $n = 1$, we find $X = (x_1, x_2, \dots, x_n)^T$, $Y = (y_1, y_2, \dots, y_n)$, such that:

$$y_i x_i = 1; \quad y_j x_i = 0 \text{ for } i \neq j; \quad \sum_{i=1}^n x_i y_i = 1. \quad (1.1)$$

Now assume R does not have invariant basis number; we ask for what subsets $S \subset \mathbb{N} \times \mathbb{N}$ is it possible that $R^m \cong R^n$ iff $(m, n) \in S$. We notice that $\{R, R^2, R^3, \dots\}$ forms a monogenic semigroup (under direct product). There is a standard classification of these: namely, a semigroup generated by a single element a is either isomorphic to \mathbb{N} , or to the semigroup $S(N, k)$ for some $N, k \in \mathbb{N}$, where $a^m \cong a^n$ if and only if $m, n \geq N$ and $m \cong n \pmod k$. In particular, knowledge of the least m, k such that $R^m \cong R^{m+k}$ is enough to determine all pairs (m, n) for which $R^m \cong R^n$. Thus, we make the definition that a ring R has *module type* (m, n) if $m < n$ are minimal such that $R^m \cong R^n$.

Leavitt, in [36], showed that rings of all module types (m, n) exist. The proof is constructive and provides, for any field K and pair (m, n) with $m < n$, a K -algebra $L_K(m, n)$ of module type (m, n) . The algebra $L_K(m, n)$ is called the *Leavitt algebra*, and it has the universal property that it has a non-zero homomorphism to any other K -algebra of the same module type. Moreover, $L_K(m, n)$ is given explicitly in terms of generators and relations, essentially by defining the generators to be coefficients of matrices X, Y with $XY = I_m, YX = I_n$. In particular, $L_K(1, n)$ has a presentation with generators $x_1, y_1, x_2, y_2, \dots, x_n, y_n$ satisfying the relations in equation 1.1.

C^* -algebras and Cuntz algebras: Now we recall some parallel C^* -algebra theory. It is usually possible to complete Leavitt algebras and their variants with respect to a suitable norm. The completed algebra is a C^* -algebra, which often has analytic properties that are very similar to the algebraic properties of the Leavitt algebra. The proofs of these results tend to be different (one analytic, one algebraic) so this is genuinely surprising. Our work will be algebraic rather

than analytic, but it will adapt the C^* -algebras described in [13], so it's worth giving a short overview of C^* theory.

Recall that a C^* -algebra is a subalgebra of the space $B(\mathcal{H})$ of continuous linear operators on a complex Hilbert space \mathcal{H} that is closed under taking adjoints and closed with respect to the norm topology. Gelfand and Naimark introduced an abstract classification of C^* -algebras, summarized as follows. We say a $*$ -algebra is a \mathbb{C} -algebra A with a conjugate-linear involution $*$, satisfying $(ab)^* = b^*a^*$ for all $a, b \in A$. A C^* -algebra is then a $*$ -algebra A which also has a norm $\|\cdot\|$, with respect to which it is a Banach space with $\|ab\| \leq \|a\|\|b\|$, and where A satisfies the *C^* -condition* $\|a^*a\| = \|a\|^2$. However, since any abstract C^* -algebra can be represented by linear maps on a Hilbert space, we will mostly work with the operator definition.

We say that a *projection* in a C^* -algebra A is a self-adjoint idempotent $p \in A$ (that is $p^2 = p = p^*$); we say an *isometry* is an element $x \in A$ with $x^*x = 1$, and a *partial isometry* is an element $x \in A$ such that x^*x is a projection. We say a C^* -algebra is simple if it contains no closed two-sided ideals, and separable if it has a countable dense subset (as is usual in topology). Finally, we say that the simple, unital C^* -algebra A is *infinite* if it contains an element $x \in A$ where $xx^* = 1$, but $x^*x \neq 1$. In particular, no finite-dimensional C^* -algebra can be infinite.

Cuntz ([25]) constructed the first example of an infinite, separable, simple C^* -algebra as follows: suppose that $\{S_i\}_{i=1}^n$ is a set of isometries on a Hilbert space \mathcal{H} satisfying $\sum_{i=1}^n S_i S_i^* = 1$. For any such choices of \mathcal{H} and $\{S_i\}$, the C^* -algebras generated by the S_i are isomorphic. This C^* -algebra is called the *Cuntz algebra* \mathcal{O}_n , and it is infinite, separable and simple. We remark that the relations defining \mathcal{O}_n are the same as the relations defining the Leavitt algebra $L_{\mathbb{C}}(1, n)$, and both algebras are simple. Moreover, \mathcal{O}_n is a completion of $L_{\mathbb{C}}(1, n)$ - see eg [1] for the proof.

Cuntz algebras were generalized in many ways as the interest in them grew. Most important in this thesis are the Cuntz-Krieger algebras associated to directed graphs, discussed in [34]:

Definition 1.2.1. *Let Γ be a directed graph. Then the Cuntz-Krieger algebra $C^*(\Gamma)$ is the C^* algebra given by the following generators and relations: the generators are mutually orthogonal projections $\{P_v : v \in \Gamma^0\}$ and partial isometries $\{S_e : e \in \Gamma^1\}$, and the relations are*

- $S_e^* S_e = P_{s(e)}$ for all $e \in \Gamma^1$

- If v is any vertex with $0 < |t^{-1}(v)| < \infty$, then

$$P_v = \sum_{e \in \Gamma^1: t(e)=v} S_e S_e^*.$$

This definition has been given with sources and targets the opposite way round to many references, because we are following the conventions of [13]. It is worth pointing out that C^* -algebra presentations are more complicated than algebra presentations. Indeed, a family of generators and relations only presents a C^* -algebra if the generators can then be realized as bounded operators on a Hilbert space. This is not possible in general. However, all the examples we give here will actually have a C^* -algebra where the given relations hold. In particular, \mathcal{O}_n is a Cuntz-Krieger algebra, for the directed graph with one vertex and n edges.

Leavitt path algebras: The algebraic equivalent of the Cuntz-Krieger C^* algebra is the Leavitt path algebra. The history of these objects is somewhat complicated and is explained in the first chapter of [1]: to summarize, they were defined independently (and for different reasons) by Abrams and Aranda Pino in [2], and by Ara, Moreno and Pardo in [6], and the notation has developed since then. We give the definition for any graph Γ , although we will mostly work with the case when Γ is finite.

Definition 1.2.2. Let Γ be a directed graph and let K be a field. Take $(\Gamma^1)^*$ to be a set of symbols e^* , one for each edge $e \in \Gamma^1$. We define the Leavitt path algebra of Γ , written $L_K(\Gamma)$, to be the K -algebra generated by the set $\Gamma^0 \cup \Gamma^1 \cup (\Gamma^1)^*$, satisfying relations:

1. $v^2 = v$ for all $v \in \Gamma^0$; $vw = 0$ for $v, w \in \Gamma^0, v \neq w$.
2. $t(e)e = es(e) = e$ for all $e \in \Gamma^1$; $s(e)e^* = e^*t(e) = e^*$ for all e^* where $e \in \Gamma^1$.
3. $e^*e = s(e)$ for all $e \in \Gamma^1$; $f^*e = 0$ for all $e, f \in \Gamma^1$ with $e \neq f$.
4. Whenever $v \in \Gamma^0$ and $0 < |t^{-1}(v)| < \infty$, then

$$v = \sum_{e \in \Gamma^1: t(e)=v} ee^*,$$

The operators P_v, S_e of the Cuntz-Krieger algebra $C^*(\Gamma)$ satisfy the same relations that v, e satisfy in the Leavitt path algebra. In fact the assignment $v \mapsto P_v, e \mapsto S_e, e^* \mapsto S_e^*$ embeds $L_{\mathbb{C}}(\Gamma)$ as a dense subalgebra of $C^*(\Gamma)$. This gives one connection in the diagram of Figure 1.10. Sometimes we will write the generators of the Leavitt path algebra as P_v, S_e, S_e^* instead of v, e, e^* . This makes sense because of the embedding described, and it will help us distinguish the edge e from the element S_e of the Leavitt path algebra.

There is a second definition of the Leavitt path algebra, as the quotient of a path algebra. Given a directed graph Γ , we define the extended graph of Γ , written $\hat{\Gamma}$, to be the directed graph with vertex set Γ^0 and edge set $\Gamma^1 \cup (\Gamma^1)^*$. For an edge e , we define $s(e^*) = t(e)$ and $t(e^*) = s(e)$. Then we could alternatively define $L_K(\Gamma)$ as the quotient of $K\hat{\Gamma}$ by the relations (3) and (4) of Definition 1.2.2.

It's sometimes helpful to think of the Leavitt path algebra $L_K(\Gamma)$ as an algebra acting on infinite paths in Γ . This description is cleanest when $|t^{-1}(v)| > 0$ for all $v \in V$. If $p \in \Gamma^\omega$, we think of e as the operator sending p to ep whenever this is a path, and to 0 otherwise. We think of e^* as removing e from p (to form p' , where $p = ep'$) where possible. If these operations are impossible, then they send the path to zero. Finally, $v \in \Gamma^0$ fixes paths p with $t(p) = v$ and sends other paths to zero. This provides a way to think about $L_K(\Gamma)$ acting on the set $K\Gamma^\omega$ of K -linear combinations of paths.

We describe why some of the relations hold for this action. Consider the third relation of Definition 1.2.2. It says that if you add e , then remove e on the left of a path p , that leaves p unchanged so long as its target is $s(e)$. If instead you add e then try to remove f , this will never be possible, so you always get zero.

Next, consider the fourth relation. Assume $v \in \Gamma^0$ has $0 < |t^{-1}(v)| < \infty$. One can think of ee^* as checking whether a path begins with e (by removing e then replacing it). The sum of terms ee^* in relation (4) then fixes $p \in \Gamma^\omega$ if and only if the leftmost edge of p has target v . This is precisely the set of infinite paths fixed by the element v of $L_K(\Gamma)$. We remark that this relation would fail if we tried to define an action on finite paths, because $v \in L_K(\Gamma)$ fixes the path v , whereas $\sum ee^*$ sends v to zero. So we do need to define the action on $K\Gamma^\omega$ rather than $K\Gamma^*$.

We conclude the introduction with the remark that $L_K(\Gamma)$ is unital if and only if Γ^0 is finite. Indeed, if Γ^0 is finite, then the sum of all vertices is a multiplicative identity (since by properties 1 and 2 of the definition, it behaves

as such when multiplying any generator). If instead Γ^0 is infinite, then for any $p \in L_K(\Gamma)$, there exists $v \in \Gamma^0$ such that $vp = 0$. But $L_K(\Gamma)$ still has local units: that is, there exists a set $E \subset L_K(\Gamma)$ of commuting idempotents, such that for every $p \in L_K(\Gamma)$ there exists $u \in E$ with $up = p$. A set of local units is given by finite sums of vertices.

1.2.3 Structure, ideals and uniqueness theorems

Here we study images of Leavitt path algebras under algebra homomorphisms. We will follow Mark Tomforde's work in [54] in this section, and the results quoted will be results from that paper. Later, when we define a new family of algebras from graphs of groups, we'll adapt these theorems to the new algebras.

Since Leavitt path algebras are defined by a presentation, they have a universal property which lets us define homomorphisms out of them. Indeed, if A is any K -algebra containing elements $\{a_v : v \in \Gamma^0\}$ and $\{b_e, b_e^* : e \in \Gamma^1\}$ which satisfy the relations of Definition 1.2.2, then there is a homomorphism $\phi : L_K(\Gamma) \rightarrow A$ sending v to a_v and e, e^* to b_e, b_e^* . We want to understand when these homomorphisms are injective. [54] provides a theory similar to that of closed ideals of Cuntz-Krieger algebras. First we describe the structure of elements of Leavitt path algebras:

Definition 1.2.3. *Let Γ be a directed graph. Suppose $p = e_1 e_2 \dots e_n$ is a path in Γ (with each $e_i \in \Gamma^1$). Then we will also write p for the product $e_1 e_2 \dots e_n$ in $L_K(\Gamma)$, and write p^* for $e_n^* e_{n-1}^* \dots e_1^* \in L_K(\Gamma)$.*

Proposition 1.2.4. *Let Γ be a directed graph and let K be a field. Then every element of $L_K(\Gamma)$ is a K -linear combination of monomials pq^* , where p and q are paths in Γ (and either p or q may be a single vertex v).*

This is proved by using relation (3) to eliminate any appearances of f^*e from an expression for an element of $L_K(\Gamma)$. We also record the fact that the algebra $L_K(\Gamma)$ comes with a \mathbb{Z} -grading.

Proposition 1.2.5. *In the usual notation, the Leavitt path algebra $L_K(\Gamma)$ is \mathbb{Z} -graded. The subspace $L_K(\Gamma)_n$ is spanned by monomials pq^* , where p and q are paths with $\ell(p) - \ell(q) = n$.*

Proof. With this definition of the graded components, it's easy to verify that $L_K(\Gamma)$ is the sum of the subspaces $L_K(\Gamma)_n$, but perhaps less obvious that the sum is direct. Instead, we use the fact that the path algebra of the extended

graph, $K\hat{\Gamma}$, is graded, where each vertex has degree 0, each edge e has degree 1 and each edge e^* has degree -1, and we extend the grading to paths multiplicatively (in fact any assignment of degree to each edge would give a \mathbb{Z} -grading on the path algebra). Now notice that the third and fourth relation of definition 1.2.2 consist of homogeneous elements of $K\hat{\Gamma}$ of degree 0. This means that the grading passes to the quotient by these relations, which is $L_K(\Gamma)$. \square

With this grading in place, it is natural to consider graded homomorphisms and graded ideals. These were classified in generality by Tomforde in [54].

Theorem 1.2.6 (Graded Uniqueness Theorem). *Let Γ be a directed graph with Leavitt path algebra $L_K(\Gamma)$, \mathbb{Z} -graded as above. Suppose that A is a \mathbb{Z} -graded K -algebra and $\pi : L_K(\Gamma) \rightarrow A$ is a graded algebra homomorphism. Suppose also that $\pi(v) \neq 0$ for all $v \in \Gamma^0$. Then π is injective.*

A proof is given in [54], Theorem 4.8. Using this result, it's possible to determine the graded ideals of Leavitt path algebras. The following is the finite case of Theorem 5.7 of Tomforde's paper [54]. We remark that our paths are directed in the opposite way to his, so the following definition might seem a bit unusual.

Definition 1.2.7. *Let Γ be a finite directed graph. A subset $H \subset \Gamma^0$ is hereditary if for any $e \in \Gamma^1$, then $t(e) \in H$ implies $s(e) \in H$ (in other words, H is closed upon passing 'upstream' or 'to ancestors'). A hereditary subset H is saturated if whenever $0 < |t^{-1}(v)|$ and $\{s(e) : e \in E^1, t(e) = v\} \subset H$, then $v \in H$. In other words, there is no single vertex v outside H which can be added to H whilst keeping H hereditary.*

Theorem 1.2.8. *Let Γ be a finite directed graph, and let $L_K(\Gamma)$ be its Leavitt path algebra over K . Then graded ideals of $L_K(\Gamma)$ are in bijection with saturated hereditary subsets of Γ^0 . The bijections are as follows: given $H \subset \Gamma^0$ which is hereditary and saturated, define I_H to be the ideal generated by H ; given a graded ideal I , define $H(I)$ to be $I \cap \Gamma^0$, which can be shown to be hereditary and saturated. These maps are mutually inverse.*

In fact, more can be proved, and $L_K(\Gamma)/I_H$ can be identified as the Leavitt path algebra of a graph formed by removing H from Γ . We instead go straight on to summarize the results of [54] about general homomorphisms, that are not required to be graded.

Definition 1.2.9. Let $p = e_1 e_2 \dots e_n$ be a path in directed graph Γ . Recall that p is said to be a closed path (based at v) if $s(e_n) = t(e_1) (= v)$. We say that an edge $f \in \Gamma^1$ is an entrance for the closed path p if $t(f) = t(e_i)$, but $f \neq e_i$, for some $i = 1, 2, \dots, n$.

As usual, the definition of entrance is opposite from [54] (which defines exits), because our paths are read right-to-left. We impose the key condition that every closed path has an entrance, which will let us study all homomorphisms out of $L_K(\Gamma)$.

Theorem 1.2.10 (Cuntz-Krieger Uniqueness Theorem). *Suppose Γ is a directed graph in which every closed path has an entrance, with Leavitt path algebra $L_K(\Gamma)$. If $\pi : L_K(\Gamma) \rightarrow A$ is a K -algebra homomorphism where $\pi(v) \neq 0$ for all $v \in \Gamma^0$, then π is injective.*

Notice that this theorem has the same conclusion as the graded uniqueness theorem, but that by assuming more conditions on the graph Γ , we have been able to drop the requirement that π is a graded homomorphism. As an example of why the condition on Γ is necessary, consider the graph Γ with one vertex v and one edge e . Its Leavitt path algebra is isomorphic to $K[X, X^{-1}]$, under an isomorphism sending v to 1, e to X and e^* to X^{-1} . So every non-zero homomorphism out of $L_K(\Gamma)$ has the property that $\pi(v) \neq 0$, whether or not it is injective.

As a corollary, [54] is able to achieve the following condition on when Leavitt path algebras are simple:

Theorem 1.2.11. *Let Γ be a directed graph. The Leavitt path algebra $L_K(\Gamma)$ is simple if and only if every closed path has an entrance, and the only saturated hereditary subsets of Γ^0 are \emptyset and Γ^0 itself.*

Finally, we state C^* -algebraic results which provide inspiration, or parallels, for these uniqueness theorems. First, we state the C^* version of the graded uniqueness theorem. Instead of having a grading, though, we use a gauge action. We quote these results from the overview given in [48], Chapter 2.

Definition 1.2.12. *Let A be a C^* -algebra and let G be a locally compact group. An action of G on A is a homomorphism $G \rightarrow \text{Aut}(A)$, mapping $s \in G$ to the automorphism α_s , such that for each $a \in A$ the map $s \mapsto \alpha_s(a)$ is a continuous function from G to A .*

Proposition 1.2.13 ([48] Theorem 2.1). *Let Γ be a locally finite directed graph. Let \mathbb{T} be the circle group $\{z \in \mathbb{C} : |z| = 1\}$ (under multiplication, with the subspace topology from \mathbb{C}). Then there is an action γ of \mathbb{T} on the Cuntz-Krieger algebra $C^*(\Gamma)$ such that $\gamma_z(P_v) = P_v$ and $\gamma_z(S_e) = zS_e$ for all $v \in \Gamma^0, e \in \Gamma^1$.*

The existence of the action follows from the universal property of Cuntz-Krieger algebras, since P_v and zS_e satisfy the same defining relations as P_v and S_e do; it just remains to check that the continuity property of the definition of actions holds.

Theorem 1.2.14 ([48] Theorem 2.2, Gauge-Invariant Uniqueness). *Let Γ be a locally finite directed graph, with $C^*(\Gamma)$ its Cuntz-Krieger algebra. Suppose that $\phi : C^*(\Gamma) \rightarrow B$ is a homomorphism of C^* algebras, such that $\phi(P_v) = Q_v$ and $\phi(S_e) = T_e$. Suppose that \mathbb{T} acts on B by action β , such that $\beta_z(T_e) = zT_e$ and $\beta_z(Q_v) = Q_v$ (for every $e \in \Gamma^1, v \in \Gamma^0$). Suppose also that each $Q_v \neq 0$. Then ϕ is injective.*

Just like the graded invariant theorem, this lets us tell that morphisms are injective if they preserve some extra structure and don't vanish on the vertices. Similarly, we have a Cuntz-Krieger uniqueness theorem, where we can drop the gauge action if we assume that every closed path in Γ has an entrance:

Theorem 1.2.15 ([48] Theorem 2.4). *Let Γ be a locally finite directed graph, with $C^*(\Gamma)$ its Cuntz-Krieger algebra, and where every closed path of Γ has an entrance. Suppose that $\phi : C^*(\Gamma) \rightarrow B$ is a homomorphism of C^* algebras, such that $\phi(P_v) = Q_v$ and $\phi(S_e) = T_e$. Suppose also that each $Q_v \neq 0$. Then ϕ is injective.*

We also get the same consequences concerning simplicity:

Corollary 1.2.16 ([1] Theorem 1.11). *Let Γ be a finite graph. Then $C^*(\Gamma)$ is simple if and only if every closed path in Γ has an entrance, and the only saturated hereditary subsets of Γ are \emptyset and Γ itself.*

Interestingly, the Leavitt path algebra results and Cuntz-Krieger results are not deduced from each other, but are typically proved by very different methods. Later in this thesis, we will define a Leavitt-type algebra from a C^* -algebra associated to a graph of groups. We will show how similar invariant theorems can be deduced for the new Leavitt-type algebra as well.

1.3 Groupoids and their Leavitt algebras

We have discussed the theory of Leavitt path algebras and the similarities they share with graph C^* -algebras. We continue the explanation of Figure 1.10 by describing how a more general family of C^* -algebras arises from groupoids. This theory has proved interesting to C^* -algebraists because it reproduces many interesting known families of C^* -algebras. It's interesting to us, because this theory generalizes both Leavitt path algebras (on the algebra side) and Cuntz-Krieger algebras (on the C^* side). We describe groupoids quite fully and then sketch the C^* theory briefly; the full C^* picture is described in Paterson's book [46]. Our groupoid theory is taken from various sources including [46], [12] and [22].

1.3.1 Elementary theory of groupoids:

A groupoid can be defined as a small category with inverses. More fully, a groupoid is a set G with a partially defined product $\cdot : G^2 \rightarrow G$ (for $G^2 \subset G \times G$) and an inverse map $a \mapsto a^{-1}$ from G to itself, such that:

- Associativity: if $(a, b), (b, c) \in G^2$, then $(a \cdot b, c), (a, b \cdot c) \in G^2$, with

$$a \cdot (b \cdot c) = (a \cdot b) \cdot c.$$

- Inverses: $(a^{-1})^{-1} = a$ for all $a \in G$; $(a, a^{-1}) \in G^2$ for all $a \in G$, and if $(a, b) \in G^2$, then

$$a^{-1} \cdot (a \cdot b) = b; (a \cdot b) \cdot b^{-1} = a.$$

These axioms informally say a groupoid is 'like a group, but with the multiplication only partially defined'; alternatively, $a, b, c, \dots \in G$ can be thought of as invertible morphisms of some category. As with groups, we normally drop the multiplication symbol and brackets.

For G a groupoid and $x \in G$, we define the domain of x to be $d(x) = x^{-1}x$ and the range of x to be $r(x) = xx^{-1}$. From the category point of view, this makes sense as $d(x), r(x)$ are then the identity maps on the domain and range objects of the morphism x . Then $(x, y) \in G^2$ if and only if $d(x) = r(y)$. The *unit space* of G is written $G^{(0)}$ and defined as:

$$G^{(0)} = \{xx^{-1} : x \in G\} = \{x^{-1}x : x \in G\}.$$

If we're thinking of G as a category, we can identify the set of objects of G with $G^{(0)}$, so that x becomes a morphism from $d(x)$ to $r(x)$. For $v \in G^{(0)}$, we write G_v for the set of morphisms of G with domain and range v : it is a group, called the *isotropy group* at v .

We give a few examples of groupoids (see [12], [47] for more, including proofs).

1. Any group G is a groupoid, where the multiplication is defined on every pair of elements (in the language of categories, G is a groupoid with one object).
2. Let \sim be an equivalence relation on the set X . Then pairs (x, y) where $x \sim y$ form a groupoid, with multiplication defined by $(x, y)(y, z) = (x, z)$ and undefined elsewhere. As a category, X is the set of objects and there is a unique morphism from x to y whenever $x \sim y$; moreover, any groupoid with trivial isotropy groups is an equivalence relation.
3. Perhaps the best example for intuition is a *transformation groupoid*: suppose H is a group that acts on the set X on the left. Then the set $G = H \times X$ is given a groupoid structure where $G^2 = \{((h, kx), (k, x))\}$ (for $x \in X, h, k \in H$); the product is

$$(h, kx) \cdot (k, x) = (hk, x)$$

and inversion is given by:

$$(h, x)^{-1} = (h^{-1}, hx).$$

As a category, this groupoid has X as the set of objects, and a morphism (h, x) from x to hx for each $h \in H, x \in X$.

4. If G_1 and G_2 are groupoids, then the direct product $G_1 \times G_2$ is a groupoid (where $(g_1, g_2)(g'_1, g'_2)$ is defined precisely when $g_1g'_1, g_2g'_2$ are both defined, and the product is $(g_1g'_1, g_2g'_2)$). Moreover the disjoint union $G_1 \sqcup G_2$ is a groupoid, where the multiplication is not defined between elements of different groupoids. In fact, every groupoid is a disjoint union of groupoids that are direct products of groups and equivalence relations.

Finally we introduce a topology. Say that a *topological groupoid* is a groupoid G equipped with a topology such that G^2 is a closed subset of $G \times G$ (in the

product topology), and such that multiplication and inversion are continuous. These conditions imply that the range and domain maps are also continuous. For example, suppose that H is a topological group acting continuously on the left of topological space X . Then the transformation groupoid $H \times X$ is a topological groupoid with the product topology.

The fundamental groupoid of a graph of groups We pause at this point to give an important example of a groupoid that will be very important in our study of graphs of groups. We take these definitions from [13], which quotes the definitions from [31], which introduced the construction.

Let \mathcal{G} be a graph of groups with underlying graph Γ . The fundamental groupoid $F(\mathcal{G})$ is defined as followed: its set of objects is Γ^0 . The morphisms of $F(\mathcal{G})$ are generated by $e \in \Gamma^1$ and by $g \in G_v$ for each $v \in \Gamma^0$, just as for the path group. For $v \in \Gamma^0$ and $g \in G_v$, then the source and domain of g are both v . For $e \in \Gamma^1$, the domain of the morphism e is $s(e)$ and its range is $t(e)$. The inverse of e is \bar{e} , and we take the relation $e\alpha_{\bar{e}}(g)e^{-1} = \alpha_e(g)$ for all $e \in \Gamma^1, g \in G_e$. Then every element of $F(\mathcal{G})$ can be written as a reduced \mathcal{G} -word (proved in [31]). Moreover, for $v \in \Gamma^0$, the isotropy group of $F(\mathcal{G})$ at v recovers the fundamental group $\pi_1(\mathcal{G}, v)$.

We'll often use the multiplication in the fundamental groupoid to define more complicated algebraic objects later on.

1.3.2 C^* -algebras from groups and groupoids

We shall use groupoids to study Leavitt path algebras, as shown in Figure 1.10. First, we give the older C^* theory which relates groupoids to Cuntz-Krieger algebras. We introduce the C^* -algebra of a groupoid below, summarizing the notes in [47].

Group C^* -algebras First let G be a discrete group (rather than a groupoid), and let $\mathbb{C}G$ be its group algebra over the complex numbers, which is a $*$ -algebra. We write elements a of $\mathbb{C}G$ as sums $a = \sum_{g \in G} a_g g$, where $a_g \in \mathbb{C}$ and all but finitely many of the a_g are zero, and then define $a^* = \sum_{g \in G} a_g^* g^{-1}$ (where a_g^* is the complex conjugate of a_g). This map $*$ is then an involution, making $\mathbb{C}G$ into a $*$ -algebra. Recall that a unitary representation of $\mathbb{C}G$ is a homomorphism of $*$ -algebras from $\mathbb{C}G$ to the algebra $\mathcal{B}(\mathcal{H})$ of bounded linear operators on some Hilbert space \mathcal{H} . In particular, the left regular representation π_λ of G is defined

on the Hilbert space $\ell^2(G)$ by:

$$\left(\pi_\lambda \left(\sum_{g \in G} a_g g \right) \cdot \xi \right) (h) = \sum_{g \in G} a_g \xi(g^{-1}h),$$

for all $h \in G$, $\xi \in \ell^2(G)$, and $\sum_{g \in G} a_g g \in \mathbb{C}G$.

We use representations of $\mathbb{C}G$ to form two C^* -algebras from G . First, define the *reduced norm* on $\mathbb{C}G$ by

$$\|a\|_r = \|\pi_\lambda(a)\|,$$

for any $a \in \mathbb{C}G$. We define the reduced C^* -algebra of G to be the completion $C_r^*(G)$ of $\mathbb{C}G$ with respect to this norm. It can be shown that this completion is a $*$ -algebra. Similarly, we define the *full norm* on $\mathbb{C}G$ by

$$\|a\| = \sup\{\|\pi_u(a)\| : \pi_u \text{ a representation of } G\}.$$

It can be shown that this supremum is always finite. The completion of $\mathbb{C}G$ in the full norm is a C^* -algebra, which we write $C^*(G)$ and call the group C^* -algebra of G .

The reduced norm and full norm are in general different, with $\|a\|_r \leq \|a\|$. The condition for them to be equal is that G is *amenable*. Amenability is an important property in group theory, but all definitions are quite technical, and we won't need this property in the rest of this thesis so we don't elaborate further.

More generally, let G be a topological group (so we no longer assume that G is discrete). We will assume G is Hausdorff and locally compact. The theory of Haar measure on topological groups is well known: there is (up to constant multiples) a unique left-translation invariant measure μ on G that is finite on compact sets and satisfies certain regularity conditions. We define the left regular representation as before, but on the space $L^2(G, \mu)$ (of L^2 integrable functions on G) rather than $\ell^2(G)$. The full and reduced norms can be defined as before, but on the space $C_c(G)$ of continuous functions on G of compact support, rather than on $\mathbb{C}G$. In these norms, $C_c(G)$ completes to C^* -algebras $C^*(G)$ and $C_r^*(G)$.

Groupoid C^* -algebras Now we describe a generalization to the case of a topological groupoid G . First we'll generalize the case of a discrete group. Let G be a topological groupoid. We say that G is *étale* if the range and domain maps r, d are local homeomorphisms. In an étale groupoid, we define an *open bisection* to be an open set $U \subset G$ such that r and d are homeomorphisms on U (open bisections are called open G -sets in [47], but open bisection seems to be the more common name). The definition of étale groupoids implies that the collection of open bisections forms a basis for the topology on G . Moreover, every étale groupoid is *r -discrete*, meaning that for every $g \in G^{(0)}$, the sets $r^{-1}(g)$ and $d^{-1}(g)$ are discrete. In particular, for each $x \in G^{(0)}$, the isotropy group $G_x = \{g \in G : d(g) = r(g) = x\}$ is discrete. If H is a discrete group acting continuously on the right of topological space X , then $G = X \times H$ is an étale groupoid.

Now let G be a locally compact, Hausdorff, étale groupoid. As in the group case, the space $C_c(G)$ of continuous functions of compact support on G is a $*$ -algebra. Unlike the group case, it has not a single regular representation, but a family of left regular representations π_λ^x , one for each $x \in G^{(0)}$, defined by:

$$(\pi_\lambda^x(a) \cdot \xi)(g) = \sum_{r(h)=r(g)} a(h)\xi(h^{-1}g),$$

for each $a \in C_c(G)$, $\xi \in \ell^2(d^{-1}(x))$, and $g \in d^{-1}(x)$. We define the left regular representation π_λ of G to be the direct sum of all these π_λ^x . The full and reduced C^* algebras $C^*(G)$ and $C_r^*(G)$ are then defined as completions of $C_c(G)$ with norms coming from representations, as before.

We finish with remarks about what happens if G is not étale. The equivalent of Haar measure is a family of measures called a Haar system, with the measures indexed by $G^{(0)}$. Haar systems do not always exist, and if they do exist, are not necessarily unique. However, when a Haar system $\{\mu^u : u \in G^{(0)}\}$ does exist, one can define convolution on the space $C_c(G)$ of continuous functions on G of compact support:

$$f * g(x) = \int_{y \in G} f(xy)g(y^{-1})d\mu^{d(x)}(y),$$

for any f and $g \in C_c(G)$. This operation makes $C_c(G)$ into a $*$ -algebra; it can again be completed to give a C^* -algebra $C^*(G)$ with appropriate norm, the I -norm ([46] Section 2.2), or to a reduced C^* -algebra $C_r^*(G)$ with a reduced

I -norm. We will only study étale groupoids in this thesis, so we don't describe this theory any more.

1.3.3 Groupoids and graph C^* algebras

It was shown in [34] that the graph C^* -algebra of a directed graph Γ arises from a groupoid G_Γ under certain mild hypotheses on Γ . The construction is as follows: let Γ be a locally finite directed graph, with vertex set Γ^0 and edge set Γ^1 , and assume $t^{-1}(v) \neq \emptyset$ for all $v \in \Gamma^0$. For simplicity, we also assume that Γ is connected (so that there is a directed path from any vertex to any other) and not a single cycle. Recall Γ^ω is the space of all infinite paths in Γ , which can be topologized with basic open sets the cylinder sets

$$Z(\alpha) = \alpha\Gamma^\omega = \{x \in \Gamma^\omega : x = \alpha x'\},$$

for each finite path α . The set Γ^ω will be the unit space for G_Γ . General elements of G_Γ will be of the form (x, k, y) , where $k \in \mathbb{Z}$ and $x, y \in \Gamma^\omega$, such that $x_n = y_{n+k}$ for all sufficiently large $n \in \mathbb{N}$. In words, $(x, k, y) \in G_\Gamma$ if and only if it is possible to replace the leftmost N edges of path x with $N + k$ other edges and get y . The multiplication is then:

$$(y, l, z) \cdot (x, k, y) = (x, k + l, z),$$

and the inversion is

$$(x, k, y)^{-1} = (y, -k, x).$$

G_Γ is topologized by taking basic open sets

$$Z(\mu, \nu) = \{(\mu x, \ell(\mu) - \ell(\nu), \nu x) : x \in \Gamma^\omega, t(x) = s(\mu)\},$$

for each pair of finite paths μ, ν with $s(\mu) = s(\nu)$. $Z(\mu, \nu)$ is then a finite disjoint union of sets $Z(\mu e, \nu e)$ where e is an edge with $s(\mu) = t(e)$. From this, one can deduce that $Z(\mu, \nu)$ has the topology of a Cantor set, so is compact Hausdorff, and G is an étale groupoid. This construction is sometimes called the groupoid of the one-sided shift, by people studying the dynamics of the shift map that sends a path $e_1 e_2 e_3 \dots$ to $e_2 e_3 e_4 \dots$.

The paper [34] then studies the C^* -algebra $C^*(G_\Gamma)$, in particular showing that it is isomorphic to the graph C^* -algebra generated by the Cuntz-Krieger

family $\{S_e : e \in \Gamma^1\}$, satisfying relations:

$$S_e^* S_e = \sum_{f:s(f)=t(e)} S_f S_f^*.$$

The authors use Renault's theory ([49]) of groupoid C^* -algebras to study the ideals of $C^*(G_\Gamma)$, in particular finding a lattice isomorphism between ideals of $C^*(G_\Gamma)$ and hereditary saturated subsets of Γ .

1.3.4 The Steinberg algebra of an ample groupoid

The algebraic version of this theory began independently with Steinberg ([52]) and with Clark, Farthing, Sims and Tomforde ([20]), both of whom defined algebras associated to topological groupoids, generalizing Leavitt path algebras. Although the constructions are different, [20] establishes that the resulting algebras are the same. These algebras became known as Steinberg algebras.

The construction is as follows (we use the exposition in [22]). Say that an étale groupoid G is *ample* if there is a base for its topology consisting of *compact* open bisections. For an ample groupoid G , the Steinberg algebra of G over field K is written $A_K(G)$ and is defined to be the space of *locally constant* functions $G \rightarrow K$ of compact support, under convolution. For example, the groupoid G_Γ associated to a directed graph is ample. A Steinberg algebra $A_K(G)$ is spanned by characteristic functions 1_B , where B is a compact open bisection of G , and the convolution on them is given by:

$$1_B * 1_{B'} = 1_{BB'},$$

where BB' is the product $BB' = \{xx' : x \in B, x' \in B', (x, x') \in G^2\}$. Since locally constant functions of compact support are dense in the space of all continuous functions of compact support, it follows that $A_K(G)$ is dense in $C_c(G)$ and so is dense in $C^*(G)$.

The Leavitt path algebra of a directed graph is a Steinberg algebra, in the same way that we saw graph C^* -algebras were a special case of groupoid C^* -algebras. Indeed, we've observed that G_Γ is an ample groupoid, from which one can construct the Steinberg algebra $A_K(G_\Gamma)$. Then $A_K(G_\Gamma)$ is isomorphic to the Leavitt path algebra, $L_K(\Gamma)$.

1.3.5 Uniqueness theorems for Steinberg algebras

Versions of the graded and Cuntz-Krieger uniqueness theorems exist even in the full generality of Steinberg algebras. We will be interested in a particular family of Steinberg algebras related to graphs of groups, and will be able to prove stronger versions of these results for that particular family. These results were proved in [19], and we take the statements from the summary paper [21]. The results hold over any ring R (that is commutative, with 1) so we state them in full generality, allowing the Steinberg algebra to be an algebra over R .

Theorem 1.3.1 (Uniqueness for Steinberg algebras). *Let G be a second-countable, ample, Hausdorff groupoid, and let R be a ring (commutative, with 1). Suppose that A is an R -algebra and that $\pi : A_R(G) \rightarrow A$ is a ring homomorphism. Let H be the subgroupoid of G given by the interior of the isotropy bundle on G (where the isotropy bundle is the disjoint union of the isotropy groups of G). Suppose that π is injective on the subalgebra $A_R(H)$ of $A_R(G)$. Then π is injective.*

There is also a graded uniqueness theorem. We state it for \mathbb{Z} -graded groupoids.

Theorem 1.3.2 (Graded uniqueness for Steinberg algebras). *Let G be a \mathbb{Z} -graded Hausdorff ample groupoid, such that the interior of the isotropy bundle of G is $G^{(0)}$, and let R be a ring (commutative, with 1). Suppose that A is a \mathbb{Z} -graded R -algebra and that $\pi : A_R(G) \rightarrow A$ is a graded ring homomorphism. Suppose that π does not vanish on $r1_K$ for any compact open $K \subset G^{(0)}$ and $r \in R$. Then π is injective.*

As usual, one of these theorems requires the stronger assumption that π is a graded homomorphism, but also gives a stronger statement, since it asks for π to be injective on a smaller set.

1.3.6 Higher-rank graphs

Finally we record the existence of a particular generalization of Leavitt path algebras that can also be understood in terms of groupoids. Kumjian and Pask in [33] introduced k -graphs, a kind of higher rank graph. Essentially, a k -graph is a directed graph whose edges come in k different colours, with commutative squares added to relate the different colours. The case $k = 1$ is just a directed graph. They defined C^* -algebras for these higher rank graphs, which expanded the range of C^* -algebras that could be studied using this kind of graph theory. The analogous algebraic construction was carried out in [5] by Aranda Pino,

Clark, an Huef and Raeburn. As usual, they define an algebra to satisfy the same Cuntz-Krieger type relations as the C^* -algebra, and prove graded uniqueness theorems for it. In what follows, we will be interested in a different groupoid variation of the Leavitt path algebra, coming from graphs of groups, and we will follow a similar programme.

1.4 Inverse semigroups

Steinberg's motivation in defining his algebras in [52] was being able to imitate inverse semigroup theory in the C^* -algebra case. We recap this theory from [46], continuing our exposition of Figure 1.10.

1.4.1 Defining inverse semigroups

An inverse semigroup is an algebraic structure that models closed sets of partial bijections of a set X , in the same way that a group models bijections. Formally, an inverse semigroup is a set S closed under an associative multiplication, where for each $s \in S$ there exists a unique $t \in S$ such that $sts = s$ and $tst = t$. We usually write s^{-1} for this t .

The most important example is as follows: if X is any set, we write $\mathcal{I}(X)$ for the set of all bijections between two subsets of X (i.e. partial bijections of X). We define the multiplication on $\mathcal{I}(X)$ by function composition wherever this makes sense (that is, $f \circ g$ is defined on $\text{dom}(g) \cap g^{-1}(\text{dom}(f))$, where $\text{dom}(f)$ is the domain of f). This makes $\mathcal{I}(X)$ into an inverse semigroup, where the inverse is given by function inversion. In fact, every inverse semigroup can be written as a sub-inverse-semigroup of some $\mathcal{I}(X)$. This result is known as the Vagner-Preston theorem, after the two independent inventors of inverse semigroups. Its proof is similar in spirit to Cayley's theorem for groups, although a lot more technical. Moreover, many basic facts about inverse semigroups are easier to understand by working inside some $\mathcal{I}(X)$. For example, it's not obvious algebraically that the set $E(S)$ of idempotents of any inverse semigroup is commutative; however, it's easy to see that the idempotents of $\mathcal{I}(X)$ are the identity functions on subsets of X , which clearly commute.

The inverse semigroup of a graph: We will find the following definition of a *graph inverse semigroup* very useful. It was first given in [4], and was related to groupoids and to path algebras in [45].

Definition 1.4.1. Let Γ be a directed graph. The graph inverse semigroup S_Γ is the semigroup given by the following presentation. The generators are Γ^0 , Γ^1 and $(\Gamma^1)^*$, a set $\{e^* : e \in \Gamma^1\}$ of formal symbols that will act as inverses to the edges. The relations are:

- $uv = 0$ whenever $u, v \in E^0$ are distinct; $v^2 = v$ for all $v \in E^0$.
- $t(e)e = es(e) = e$ for all $e \in E^1$; $s(e)e^* = e^*t(e) = e^*$ for all $e \in E^1$.
- $e^*f = 0$ whenever $e, f \in E^1$ are distinct; $e^*e = s(e)$ for all $e \in E^1$.

Here 0 is a zero element of S_Γ , so $0x = x0 = 0$ for all $x \in S_\Gamma$. In general, inverse semigroups can have a unique zero element, which corresponds to an empty partial bijection.

As usual, our definition uses the opposite source and targets to most references because we're following [13]. This definition of S_Γ is given by a semigroup presentation: that is, as a quotient of a free semigroup by the smallest congruence containing certain relations. So one has to check that S_Γ is an inverse semigroup. In fact, every non-zero element can be uniquely written as pq^* , where p and q are paths with $s(p) = s(q)$, and pq^* has inverse qp^* .

If S is any inverse semigroup, then one can form the inverse semigroup algebra kS over a field k . It is defined as the k -algebra with basis S (or basis $S \setminus \{0\}$, if S has a zero) and with multiplication extended k -linearly. For a graph inverse semigroup S_Γ , it's easy to see that its inverse semigroup algebra is the Leavitt path algebra $L_k(\Gamma)$, because this algebra has basis pq^* for p, q paths with the same source. Moreover, one can define a semigroup C^* -algebra, and the semigroup C^* -algebra for Γ is also the graph C^* -algebra of Γ .

1.4.2 Relating groupoids and inverse semigroups

We now can describe the leftmost section of the master diagram of Figure 1.10.

The inverse semigroup G^a : First let G be an ample groupoid, and let G^a be the set of its compact open bisections. It's easy to check that G^a is closed under (groupoid) products and inverses, so that G^a becomes an inverse semigroup. There is a natural morphism of inverse semigroups $\pi : G^a \mapsto \mathcal{I}(G^0)$, since any bisection gives a bijection from its domain to its range, which are two subsets of $G^{(0)}$. The morphism π is injective if and only if G is *effective*: that is, the interior of the isotropy bundle of G is $G^{(0)}$ (see eg [21], Section 3.4).

The groupoid C^* -algebra $C^*(G)$ and inverse semigroup G^a are linked via ([46] Theorem 3.3.1):

$$C^*(G) \cong C_0(G^{(0)}) \times_{\beta} G^a,$$

where $C_0(G^{(0)})$ is the algebra of continuous functions on $G^{(0)}$ vanishing at infinity, and the notation \times_{β} describes the crossed product C^* -algebra for the natural action β of G^a on $C_0(G^{(0)})$. The definition of the crossed product doesn't matter here, just that it is a C^* -algebra formed from $G^{(0)}$ and G^a .

The groupoid of germs of an inverse semigroup: In the other direction, section 4 of [45] constructs an ample groupoid G from a (suitable) inverse semigroup S such that $C^*(G)$ is isomorphic to $C^*(S)$, the completed semigroup algebra of S . We follow the exposition in [52].

Let S be an inverse semigroup which acts on a locally compact Hausdorff space X (that is, there is a homomorphism $\phi : S \rightarrow \mathcal{I}_X$, whose image consists of partial homeomorphisms). Assume that the action is *non-degenerate*, meaning that every $x \in X$ is in the domain of some $s \in S$. We will form a groupoid written $S \times_{\phi} X$ from this action. As a set, $S \times_{\phi} X$ is $\{(s, x) \in S \times X : x \in d(\phi(s))\}$, quotiented by the equivalence relation where $(s, x) \sim (t, y)$ precisely when $x = y$ and there exists $u \in S$ such that $\phi(u)$ is defined on x , and is a restriction of both $\phi(s)$ and $\phi(t)$. We write $[s, x]$ for the equivalence class of (s, x) and call it the *germ* of s at x . It can be thought of as the action of s in small neighbourhoods of x . The multiplication $[s, x] \cdot [t, y]$ is defined if and only if $ty = x$, and in this case the product is $[st, y]$. $S \times_{\phi} X$ can be topologized, with basic open sets $[s, U] = \{[s, x] : x \in U\}$, whenever U is an open subset of X on which $\phi(s)$ is defined. This topology makes $S \times_{\phi} X$ into an étale groupoid, and $[s, U]$ an open bisection. If X has a basis of compact open sets then $S \times_{\phi} X$ is ample.

Finally, suppose that G is a Hausdorff ample groupoid and $S = G^a \subset \mathcal{I}(G^{(0)})$. We claim that $S \times G^{(0)}$ recovers G (where we drop the action ϕ from the notation, because there is a unique action on $G^{(0)}$ given by S being a subset of $\mathcal{I}(G^{(0)})$). Indeed, $S \times G^{(0)}$ is a groupoid with unit space $G^{(0)}$. For $g \in G$, let U be a compact open bisection containing g , so that $U \in S$ and $[d(g), U] \in S \times G^{(0)}$. We will show that $g \mapsto [d(g), U]$ is an isomorphism. It is well-defined, because if U, V are two compact open bisections containing $d(g)$, then $U \cap V$ is also a compact open bisection and $[d(g), U] = [d(g), U \cap V] = [d(g), V]$. It is injective, since G is Hausdorff, and surjective, by definition of $S \times G^{(0)}$. Finally it's easy

to check this is a homomorphism. So this relates our two constructions.

We now complete the overview of Figure 1.10 by introducing Thompson's group V . We first define it as a group of permutations of the ends of a tree, before relating it to Leavitt path algebras.

1.5 Thompson's group V

In 1965, Richard Thompson defined three infinite groups now called F, T and V . They were originally used to answer questions about groups with solvable word problem. Since then, they have been found to have many further interesting properties, such as T and V being infinite simple groups that are finitely presented (in fact, FP_∞). Moreover, it has become popular to define generalizations of Thompson's groups, and study how properties like simplicity pass to the generalizations.

In this work, we shall mostly be concerned with V , and we will introduce it by its action on (the ends of) an infinite binary tree. We shall see how V has many similarities to the symmetric groups \mathfrak{S}_n , and has better finiteness properties than the group \mathfrak{S}_∞ of permutations of an infinite set. Our chief reference will be the exposition given in [17], which is an excellent summary of the basic results on all three of Thompson's groups. We shall also describe some more modern generalizations of V .

1.5.1 Defining V

Preliminaries on trees: First we fix some notation on binary trees; this introduction follows [40]. Let $X = \{a, b\}$ and let X^* be the set of all finite (possibly empty) words $x_1x_2\dots x_k$ over the alphabet X (which can also be thought of as the free monoid generated by X). Write ℓ for the length function on X^* , and let X^n be the subset of words of length n , so that X^* is a union of the sets X^n over $n \in \mathbb{N}_0$. Let X^ω be the set of infinite words $x_1x_2x_3\dots$ over X , which we can identify with a Cartesian product of \mathbb{N} copies of X . For $v \in X^*$ and $w \in X^*$ or $w \in X^\omega$ the product vw is defined by the obvious concatenation. We give X^ω the topology of the Cartesian product - its basic open sets are cylinder sets $Z(v) = vX^\omega = \{vw : w \in X^\omega\}$, for each $v \in X^*$. The cylinder sets are both open and closed, so X^* is totally disconnected; it is compact, as a product of compact sets (in fact, finite sets), and has no isolated points, so is homeomorphic to a Cantor set.

There is an infinite binary tree $\mathcal{T} = \mathcal{T}_X$ associated to X , whose vertices are labelled by X^* . This tree will be a graph in the sense of Serre, so every edge is directed and has a reverse. For every $w \in X^*$, there will be an edge $e_{w,a}$ with source w and target wa , and an edge $e_{w,b}$ with source w and target wb (as well as their reverses). This means that every vertex lies on three (pairs of) edges, except for the root \emptyset , which lies on two. This tree is drawn in Figure 1.11, although the edges are unlabelled. Finally we consider the space $\partial\mathcal{T}$ of ends of \mathcal{T} : since every end has a unique representative with target \emptyset , we see that $\partial\mathcal{T}$ is in bijection with X^ω , and we give $\partial\mathcal{T}$ the topology from X^ω .

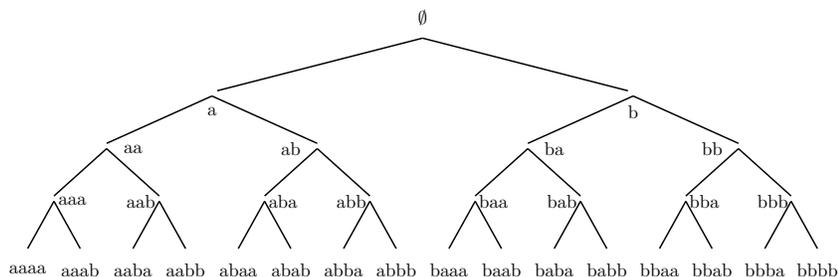


Figure 1.11: The top of the tree \mathcal{T}

Finally we define some orders on X^* which formalize the geometric arrangement of vertices in our pictures of \mathcal{T} . If $v \in X^*$ and $v' \in X^*$ or X^ω , we say that v' lies *below* v (or v lies *above* v') if one can write $v = x_1x_2 \dots x_n$ and $v' = x_1x_2 \dots x_{n+k}$ or $v' = x_1x_2 \dots$ as words over X . If $\{v_i : i \in I\}$ is a set of vertices such that none lies below another, we say the vertices v_i are *incomparable*. Suppose now that $v = x_1x_2 \dots x_m$ and $v' = x'_1x'_2 \dots x'_n$ are incomparable (where each $x_i, x'_i \in X$). Then there is a least i such that $x_i \neq x'_i$; if $x_i = a$ and $x'_i = b$, we say v lies *to the left* of v' , and otherwise say v lies *to the right* of v' .

Defining the group Here, we will define V as a group of permutations (in fact, homeomorphisms) of the end space $\partial\mathcal{T}$. One can instead define V as a group of functions from $[0, 1]$ to itself, which are bijective and linear except at finitely many dyadic rationals, and whose slopes are powers of 2 on any linear section. We will prefer to work with trees.

Say that a finite subtree S of \mathcal{T} is *full* if \emptyset is a vertex of S , and whenever v is a vertex of S , then either both va and wb are vertices of S , or neither of them are. Define a *leaf set* to be a finite subset L of X^* such that every element of X^ω lies below precisely one element of L . Equivalently, a leaf set is a finite set

L of incomparable vertices that is maximal, in the sense that no vertex of \mathcal{T} is incomparable to all of L . It's easy to see that if S is a full subtree of \mathcal{T} , then the leaves (valence 1 vertices) of S form a leaf set; conversely, given any leaf set L , there's a unique finite subtree T_L of \mathcal{T} with leaves L .

Now we can define V .

Definition 1.5.1. V is the group of all permutations of X^ω which can be written in the following form: let L_1 and L_2 be leaf sets, and let $\phi : L_1 \rightarrow L_2$ be a bijection. We define a permutation $\bar{\phi}$ of X^ω as follows: given $\rho \in X^\omega$, let $\rho = w\rho'$, where $w \in L_1$ (which exists and is unique, by definition of leaf sets). Then we define $\bar{\phi}(\rho) = \phi(w)\rho'$.

We often write a bijection ϕ between leaf sets L_1, L_2 by drawing the full subtrees T_{L_1}, T_{L_2} of which they are the leaves. We call T_{L_1} the domain tree and T_{L_2} the range tree of ϕ . An example is drawn in figure 1.12. Here we've used colours and shading to specify the bijection but it's more common to use numerical labels. For example, to define $\bar{\phi}(bbaabbaa\dots)$, we write $\rho = bbaabbaa\dots$ as $bb\rho'$, and then since $\phi(bb) = aab$, we get that $\bar{\phi}(\rho) = aab\rho' = aabaabbaa\dots$

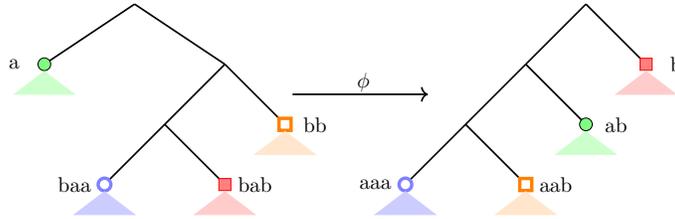


Figure 1.12: A bijection of leaf sets

It's important to note that V is a group of permutations of X^ω (which gives the composition); different pairs of trees can represent the same element of V . In fact, the following facts are easy to show:

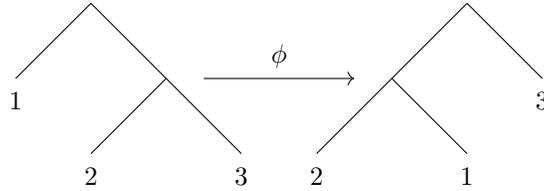
Proposition 1.5.2. Let L be a leaf set and let $v \in L$; we define a simple expansion of L at v to be the set L' formed by replacing the word v with the two words va and vb . Then L' is a leaf set. Define an expansion of L to be any leaf set formed by a sequence of simple expansions.

If $\phi : L_1 \rightarrow L_2$ is a bijection between leaf sets, suppose L'_1 is a simple expansion of L_1 at v , and L'_2 is a simple expansion of L_2 at $\phi(v)$. Define

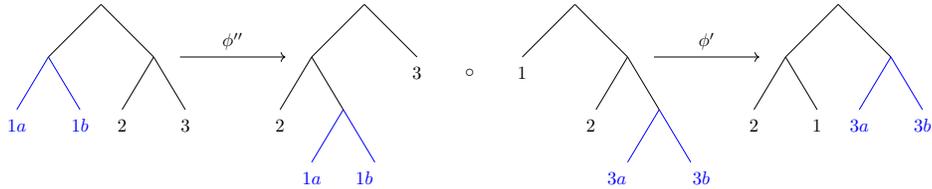
$\phi' : L'_1 \rightarrow L'_2$ to be the bijection equalling ϕ on $L_1 \setminus \{v\}$, and with $\phi'(va) = \phi(v)a, \phi'(vb) = \phi(v)b$. Then we say ϕ' is a simple expansion of ϕ (at v), and we have that $\bar{\phi}' = \bar{\phi}$. We say an expansion of ϕ is any bijection between leaf sets formed by a sequence of simple expansions. For ϕ and ψ bijections between leaf sets, we have that $\bar{\phi} = \bar{\psi}$ if and only if ϕ and ψ have a common expansion. Finally, any two leaf sets L and L' have a common expansion.

Expansions are useful for calculating products. If $\phi : L_1 \rightarrow L_2$ and $\psi : L_3 \rightarrow L_4$ are bijections between leaf sets, we compose $\bar{\psi} \circ \bar{\phi}$ by finding expansions $\phi' : L'_1 \rightarrow L'_2$ and $\psi' : L'_3 \rightarrow L'_4$ of ϕ and ψ respectively, with $L'_2 = L'_3$. Then we can compose $\psi' \circ \phi' : L'_1 \rightarrow L'_4$, and this bijection defines the composition in V .

An example is drawn out below; let $L_1 = \{a, ba, bb\}$ and $L_2 = \{aa, ab, b\}$, with bijection ϕ sending a to ab , ba to aa and bb to b . We label the leaf sets of the corresponding full subtrees to represent this bijection.

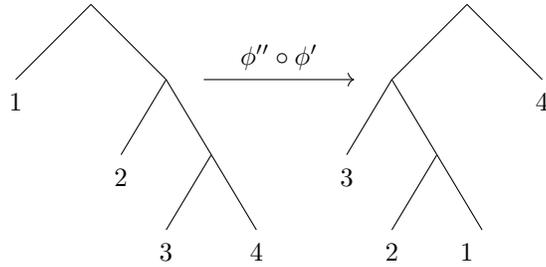


We want to calculate $\bar{\phi}^2$. To do this, we find expansions of ϕ such that the range tree of one expansion is the domain tree of the next. This is shown below. You can see how in each case, a simple expansion of ϕ has been formed by adding a caret (an inverted V-shape) below two corresponding leaves in the domain and range trees.



The composition $\phi'' \circ \phi'$ is now easy to write down, and $\bar{\phi}^2 = \bar{\phi}'' \circ \bar{\phi}'$, as shown below.

From now on, we will often not bother to distinguish ϕ and $\bar{\phi}$, and just write ϕ for both, remembering that different bijections between leaf sets may represent the same element of V .



1.5.2 Properties of V

The following is a very brief overview of important properties of V . Proofs and more details are given in [17].

If L is any leaf set, let \mathfrak{S}_L be the set of permutations of L . This symmetric group \mathfrak{S}_L embeds into V (permuting L): call this embedding θ_L . So V contains every finite symmetric group, and hence every finite group, as a subgroup. If L is a leaf set, let L^+ be the leaf set formed by doing $|L|$ simple expansions, one at each vertex of L (so it contains precisely the vertices va, vb where $v \in L$). If σ is a permutation of L , then we can expand σ at every vertex of L to get a permutation of L^+ . Thus the symmetric groups \mathfrak{S}_L inside V come with natural embeddings $\mathfrak{S}_L \hookrightarrow \mathfrak{S}_{L^+}$. If $|L| = n$, this gives an embedding $\mathfrak{S}_n \hookrightarrow \mathfrak{S}_{2n}$, that sends transpositions to double transpositions. We shall study these embeddings more when we think about Hecke algebras.

We define a permutation in V to be the image of an element of some \mathfrak{S}_L under the embedding θ_L , and define a transposition of V to be the image of a transposition in some \mathfrak{S}_L under θ_L . [17] proves that V is generated by its permutations. Since every symmetric group is generated by its transpositions, V is also generated by its transpositions. In fact, V is finitely generated, and [17] gives a finite presentation. Moreover, V is F_∞ , a finiteness property that is stronger than being finitely presented. We won't study F_∞ properties here, but we will discuss presentations for V and related groups later. We will quote important presentations as they come up.

Finally, V is a simple group, and so in particular is equal to its own commutator subgroup. Thus V provides an example of a finitely presented infinite simple group. Most groups we define as variants of V will also have infinite simple commutator subgroups.

1.5.3 The group F and the Higman-Thompson groups

F is an important subgroup of V consisting of all elements of V where the bijection ϕ is left-to-right order preserving. F is another finitely presented group; however, it is not simple. Indeed, suppose that $L_1 = \{x_1, x_2, \dots, x_n\}$ and $L_2 = \{y_1, y_2, \dots, y_n\}$ are two leaf sets of size n , where $x_i, y_j \in X^*$, and the leaf sets as written are given in left-to-right order. Then $\phi : x_i \mapsto y_i$ defines an element of F (in fact, it's the only element of F with domain leaf set L_1 and range leaf set L_2). Consider now the pair $(s_l(\phi), s_r(\phi))$ where $s_l(\phi) = \ell(y_1) - \ell(x_1)$, $s_r(\phi) = \ell(y_n) - \ell(x_n)$ (here ℓ is the length function on the set X^*). Expanding ϕ to ϕ' does not change the pair $(s_l(\phi'), s_r(\phi'))$, and in fact (s_l, s_r) gives a homomorphism from F onto \mathbb{Z}^2 , so F is not simple. One can show that the kernel of this homomorphism is the commutator subgroup F' , and that F' instead is simple.

There has been much interest in defining variants of F and V . One of the commonest and simplest is due to Higman. For $n, d \in \mathbb{N}$ with $n > 1$, the Higman-Thompson group $V_{n,d}$ is defined in the same way as V , except instead of working with a single 2-regular tree, one starts with a forest of d trees, each of which is n -regular. To be precise: if X is a set of size n , and D is a set of size d , then vertices of this family of trees can be represented by the set $DX^* = \{vx_1x_2\dots x_k : x_i \in X, v \in D\}$. There is an edge (and its reverse) between two vertices w and w' precisely when w and w' differ by the addition of a single character on the right. The ends are parametrized by $DX^\omega = \{vx_1x_2\dots : x_i \in X, v \in D\}$. Leaf sets are then defined as before, as finite subsets L of DX^* such that every element of DX^ω lies below a unique element of L . The *Higman-Thompson group* $V_{n,d}$ is then defined as all permutations $\bar{\phi}$ of DX^ω extended from a bijection ϕ between leaf sets, in the same way as for V .

Thompson's group V is then $V_{2,1}$. In general, each $V_{n,d}$ is finitely presented (in fact, F_∞), just like V . However, $V_{n,d}$ is not simple in general. If n is odd: there is an index 2 normal subgroup whose elements are products of an even number of transpositions. This subgroup is simple, and analogous to the alternating group inside the symmetric group. For n even, a single transposition is also the product of an even number of transpositions (on expanding L to L^+). So the parity of a permutation is not well-defined, and in that case $V_{n,d}$ is indeed simple. In either case, the commutator subgroup of $V_{n,d}$ is simple.

As well as F and V , Richard Thompson defined a third group T , with

$F \leq T \leq V$. Its elements are defined by bijections between leaf sets preserving cyclic order. This is covered in [17], but we don't need it in this thesis.

1.5.4 Relating V to Leavitt algebras

Here we describe a way to form V from a Leavitt algebra (in the sense of 1.2.2), which Nekrashevych describes in [40] for a Cuntz algebra and some of its generalizations. The construction of [40] never uses the C^* -norm or analytic structure, just the generators of the algebra, so we can explain it using the dense Leavitt subalgebra of the Cuntz algebra.

Let L be the Leavitt algebra $L_{\mathbb{C}}(1, 2)$, which is the Leavitt path algebra for the directed graph Γ with one vertex v and two edges e, f (both of which have source and target v). Recall that this Leavitt algebra is spanned by elements pq^* , where p, q are paths in Γ . We define the group V_{Γ} to be the set of all elements x of L of the form $x = \sum_{i=1}^n p_i q_i^*$ which are invertible, with inverse $\sum_{i=1}^n q_i p_i^*$ (that is, they are unitaries in the Cuntz algebra: $x^{-1} = x^*$). We claim that V_{Γ} is isomorphic to Thompson's group V .

Indeed, recall that L acts faithfully on the \mathbb{C} -span of infinite paths in Γ . But the infinite paths are also labelled by the set $\{e, f\}^*$ of ends of the infinite binary tree, so V acts on this set also. An element x of V can be specified by a bijection $\phi : q_i \mapsto p_i$ between two leaf sets of the tree. The element x then maps the end $q_i \rho$ to $p_i \rho$ (for any $\rho \in \Gamma^{\omega}$), so it acts on the ends identically to $\sum p_i q_i^*$. Moreover, the inverse of x maps $p_i \rho$ to $q_i \rho$, so is $\sum q_i p_i^*$, and this means $\sum p_i q_i^*$ is a unitary element of L . Conversely, it can be shown that if $\sum_{i=1}^n p_i q_i^*$ is unitary, then $\{p_i : 1 \leq i \leq n\}$ and $\{q_i : 1 \leq i \leq n\}$ both form leaf sets, giving the reverse map.

This construction, done in more generality, is used to construct Nekrashevych-Röver groups. These are one of two families of variants of V that we discuss in the next section.

1.5.5 More variants on Thompson groups

In this section we define two other generalizations of the Higman-Thompson groups that will become important later.

Colour-preserving Thompson groups

First we describe a family of Thompson-like groups that can be defined by restricting which bijections between leaf sets are permitted. These groups were defined by Matui in [39] where he described them as topological full groups of groupoids from certain shifts. Here we will prefer to define them as permutation groups of the ends of trees, but we will explain the groupoid theory later. The description of these groups given here appears to be original, although Lederle made a similar construction in [37] to identify them as subgroups of Neretin group variants (of which more later). The name ‘colour-preserving Thompson groups’ is my own.

The information to define a colour-preserving Thompson group is as follows. Let \mathbf{C} be a finite set. We call the elements of \mathbf{C} colours, and will usually write them with letters A, B, \dots . We write χ for a colouring function that will assign to vertices v of a graph colours $\chi(v) \in \mathbf{C}$. We equip \mathbf{C} with a production rule p , that assigns to each colour c of \mathbf{C} a tuple $p(c) = (c_1, c_2, \dots, c_r)$ of colours. Let $\mathbf{S} = (s_1, s_2, \dots, s_m)$ be another tuple of colours, the *starting set*. Fix for each colour c a set $\{x_{c,1}, x_{c,2}, \dots, x_{c,r}\}$ of size $|p(c)|$.

Now form a family $\mathcal{T}_{\mathbf{C},\mathbf{S}} = (\mathcal{T}_1, \dots, \mathcal{T}_m)$ of m trees inductively, as follows. We start by taking a set $S_v = \{v_1, v_2, \dots, v_m\}$ of vertices at depth 0, and give each v_i the colour $\chi(v_i) = s_i$. Inductively, whenever x is a depth n vertex of colour c , take $|p(c)|$ new vertices of depth $n+1$, coloured by the tuple $p(c)$, and connected to x . We name these vertices $xx_{c,i}$, for $1 \leq i \leq |p(c)|$ (where $xx_{c,i}$ is coloured with the colour c_i). This defines a family of m trees, with coloured vertices, and we define \mathcal{T}_m to be the tree in this family containing the depth 0 vertex v_m . All the edges will be directed and reversible, in the sense of Serre. Vertices are labelled by sequences $vx_{c_1,i_1}x_{c_2,i_2} \dots x_{c_n,i_n}$, where $v \in S_v$ and c_{k+1} is the i_k th element of $p(c_k)$. We write \mathbf{SC}^* for the set of all such sequences, which is in bijection with the set of vertices.

Definition 1.5.3. *We say that a family of subtrees $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ generated by the above process is self-similarly coloured.*

If $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ is self-similarly coloured, then whenever v, w are vertices of $\mathbf{T}_{\mathbf{C},\mathbf{S}}$ of the same colour, then the subtrees below v and w are isomorphic as coloured trees.

An example is in Figure 1.13, where $\mathbf{C} = \{A, B\}$, with production rule $p(A) = (A, B, A)$ and $p(B) = (A, B)$, and where $\mathbf{S} = (A, B)$. We just show

the colours of the vertices, and don't write the elements of \mathbf{SC}^* which label them. To translate between the formal description and the picture, it is helpful to think of $x_{c,i}$ as meaning that we take the i th edge of the edges below a vertex with colour c .

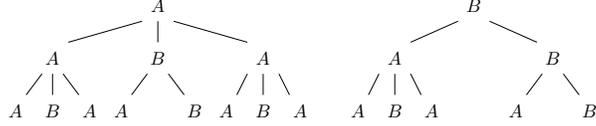


Figure 1.13: An example of coloured trees

As before, we define an end of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ as an equivalence class of infinite paths $e_1e_2e_3\dots$ under the relation where $e_1e_2e_3\dots \sim f_1f_2f_3\dots$ if $e_{n+k} = f_n$ for some $k \in \mathbb{N}$, all large enough n . We write $\partial\mathcal{T}_{\mathbf{C},\mathbf{S}}$ for the set of ends.

The colour-preserving Thompson group $V_{\mathbf{C},\mathbf{S}}$ will be defined as a group of permutations of $\partial\mathcal{T}_{\mathbf{C},\mathbf{S}}$. We observe that $\partial\mathcal{T}_{\mathbf{C},\mathbf{S}}$ is in bijection with the set:

$$\mathbf{SC}^\omega = \{vx_{c_1,i_1}x_{c_2,i_2}\dots\},$$

where $v \in S_v$, c_1 is the colour of v , and c_{k+1} is the i_k th component of $p(c_k)$ for all $k \geq 1$. This is true because an end can be specified by giving one vertex of each depth, each connected to the next. We will define $V_{\mathbf{C},\mathbf{S}}$ as a set of permutations of \mathbf{SC}^ω .

Define a full subforest of $\mathcal{T}_{\mathbf{S},\mathbf{C}}$ to be a tuple (S_1, S_2, \dots, S_m) , where S_i is a subtree of the tree \mathcal{T}_i such that S_i contains the depth 0 vertex, and where if v is a depth n vertex of S_i , then either all or none of the depth $n+1$ vertices connected to v in \mathcal{T}_i lie in S . We will say that a leaf set is the collection of degree 1 vertices in a full subforest. Equivalently, a leaf set is a subset L of \mathbf{SC}^* such that every element of \mathbf{SC}^ω lies below a unique element of L . This is all analogous to the construction of V .

We define $V_{\mathbf{C},\mathbf{S}}$ as a group of permutations of \mathbf{SC}^ω . Suppose L_1, L_2 are two leaf sets and $\phi : L_1 \rightarrow L_2$ is a *colour-preserving* bijection between them. Then there exists a permutation $\bar{\phi}$ of \mathbf{SC}^ω defined as follows: given $\rho \in \mathbf{SC}^\omega$, write $\rho = w\rho'$ for $w \in L_1$ (w exists, and is unique, by definition of leaf sets). Then define $\bar{\phi}(\rho) = \phi(w)\rho'$, just as for V . Notice that this is still an element of \mathbf{SC}^ω because ϕ is colour preserving. We define $V_{\mathbf{C},\mathbf{S}}$ to be the group of all permutations of \mathbf{SC}^ω which can be written as $\bar{\phi}$. Notice that a simple expansion of a colour-preserving bijection is still colour-preserving, so that we can compose

two bijections by expanding until the range tree of one equals the domain tree of the other. We give some examples:

1. Let $\mathbf{C} = \{A\}$, let $\mathbf{S} = (A)$, and let $p : A \mapsto (A, A)$. Then $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ is the infinite binary tree where every vertex is labelled A . So there are no restrictions on the bijection, and $V_{\mathbf{C},\mathbf{S}}$ is the usual Thompson's group V .
2. Let $\mathbf{C} = \{A\}$ again, but now let S be a d -tuple of A s and let p map A to an n -tuple of A s. There are again no restrictions on the bijection, and we get the Higman-Thompson group $V_{n,d}$.
3. Let $\mathbf{C} = \{A, B\}$ and take $\mathbf{S} = \{A\}$, and let p map A to (A, B, A) and B to (B, A, B) . Consider the element of $V_{\mathbf{C},\mathbf{S}}$ defined in Figure 1.14, where the numbers define the (order-preserving) bijection:

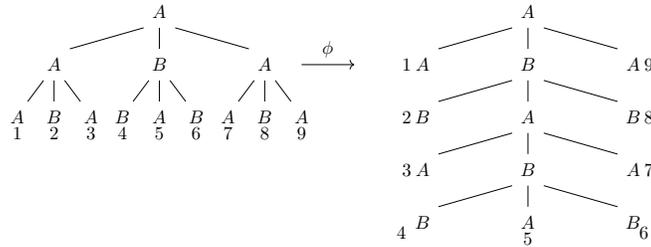


Figure 1.14: A colour-preserving Thompson element with an interesting quotient

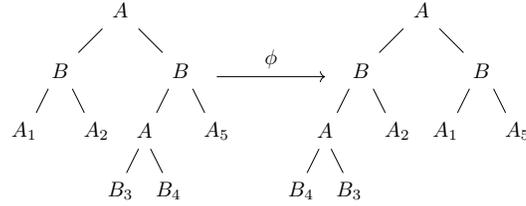
Observe that the domain tree of ϕ has 3 A s and one B as colours of its interior vertices, whilst the range tree has 2 A s and 2 B s (and they both have 5 A s and 4 B s on the leaves). Moreover, doing a simple expansion of ϕ adds either one A or one B to both trees.

Generalizing, we can define a function δ from $V_{\mathbf{C},\mathbf{S}}$ to \mathbb{Z}^2 which sends $\psi \in V_{\mathbf{C},\mathbf{S}}$ to the pair $(\delta_A(\psi), \delta_B(\psi))$, where δ_A is the difference in the number of A s between domain and range tree, and δ_B is the difference in the number of B s; for example, $\delta(\phi) = (-1, 1)$. It's clear that this is a homomorphism. So we get a non-trivial homomorphism from $V_{\mathbf{C},\mathbf{S}}$ to \mathbb{Z}^2 , which is interesting as most variants of V are either simple or have an index 2 simple subgroup, so cannot have such homomorphisms.

Isomorphisms of colour-preserving Thompson groups: Finally we make some remarks about when some of these colour-preserving Thompson groups are

isomorphic. Suppose c is a colour in a set \mathbf{C} , such that the tuple $p(c)$ does not contain c . Then if $\phi : L_1 \rightarrow L_2$ is a bijection of leaf sets defining an element of the colour-preserving Thompson group $V_{\mathbf{C},\mathbf{S}}$, we can perform a simple expansion at each leaf in L_1 of colour c . This gives a bijection ψ between two larger leaf sets, with $\bar{\phi} = \bar{\psi}$, and where the larger leaf sets have no elements of colour c . Thus, we can replace \mathbf{C} with a smaller set of colours \mathbf{C}' , by removing c from \mathbf{C} , and replacing c with $p(c)$ whenever it occurs in a production rule. We also replace c with $p(c)$ wherever it occurs in \mathbf{S} to form a new starting set \mathbf{S}' . The ends of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ and $\mathcal{T}_{\mathbf{C}',\mathbf{S}'}$ will then be in bijection, and the groups $V_{\mathbf{C},\mathbf{S}}$ and $V_{\mathbf{C}',\mathbf{S}'}$ will be isomorphic.

As an example, consider Figure 1.15. Let $\mathbf{C} = \{A, B\}$ with production rule $A \mapsto (B, B), B \mapsto (A, A)$ and $\mathbf{S} = (A)$. We give an example of a typical element of $V_{\mathbf{C},\mathbf{S}}$ and rewrite it in terms of $\mathbf{C}' = \{A\}$ with production rule $A \mapsto (A, A, A, A)$. This yields an isomorphism between $V_{\mathbf{C},\mathbf{S}}$ and the Higman-Thompson group $V_{1,4}$.



corresponds to the element

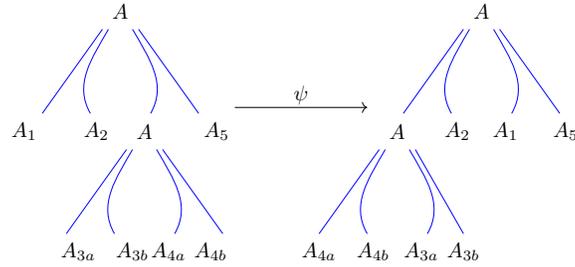


Figure 1.15: Corresponding elements of two isomorphic colour-preserving Thompson groups

Thus, we say that \mathbf{C} is *minimal* if c appears in the tuple $p(c)$ for all colours $c \in \mathbf{C}$. We will often restrict ourselves to minimal sets of colours, with no loss of generality. We will also not consider the situation $\mathbf{C} = \{A\}, p(A) = (A)$,

when the tree $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ just consists of a single end at each element of \mathbf{S} .

Finally, we remark that \mathbf{S} can be replaced by any leaf set of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ to give a new colour-preserving Thompson group which is isomorphic to $V_{\mathbf{C},\mathbf{S}}$ (as an abstract group). Indeed, this just removes finitely many vertices from the top of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$; the end space remains the same, and any bijection between leaf sets has an expansion not involving the removed vertices. So replacing \mathbf{S} by an expansion of \mathbf{S} gives an isomorphic colour-preserving Thompson group. For Higman-Thompson groups $V_{n,r}$, a simple expansion replaces one of the r root vertices with n extra vertices. Thus we get a (known) isomorphism $V_{n,r} \cong V_{n,r+k(n-1)}$, for any $k \in \mathbb{N}$. Enrique Pardo proved the converse of this result in [44], showing that $V_{n,r} \cong V_{n',r'}$ if and only if $n = n'$ and $r \equiv r' \pmod{n-1}$. This solves the isomorphism problem for Higman-Thompson groups.

Nekrashevych-Röver groups

Here we describe a family of variants of Thompson's group V that include a group G of tree automorphisms as a subgroup. These groups were first studied systematically by Nekrashevych in [40], where he produced them from C^* -algebras as groups of particular unitaries (in the manner of Section 1.5.4). The first such group was defined by Röver in [50], for the particular case where G is the Grigorchuk group, a particular group of automorphisms of the infinite rooted binary regular tree. Hence, this family of groups have become known as Nekrashevych-Röver groups. We follow the exposition in [40].

In this section, let X be a finite set, of size d . As before, let X^* be the set of finite words $x_1x_2 \dots x_n$ over the alphabet X (so that $x_i \in X$) and let X^ω be the set of infinite words $x_1x_2 \dots$. The set X^* then labels the vertices of a (rooted, d -regular) tree $\mathcal{T}_X = \mathcal{T}_d$ whose ends are labelled by X^ω . We drew out the 2-regular case in Figure 1.11.

We'll consider groups of automorphisms of \mathcal{T}_X . Automorphisms of \mathcal{T}_X induce permutations of X^ω . Conversely, a permutation σ of X^ω arises from a tree automorphism if and only if for any $w \in X^*$ there exist $\sigma_n(w) \in X^*$ and a permutation $\sigma|_w$ of X^ω , such that:

$$\sigma(w\rho) = \sigma_n(w) \circ \sigma|_w(\rho),$$

for all $\rho \in X^\omega$. Here σ_n is a permutation of the set X^n of length n words, and $\sigma|_w$ describes the action of σ on the tree below w . Notice that (for any fixed n), σ is determined by the permutation σ_n and the tree automorphisms

$\{\sigma|_w : w \in X^n\}$. This motivates the following definition:

Definition 1.5.4. *Let G be a group of automorphisms of \mathcal{T}_X . We say that the action of G on \mathcal{T}_X is self-similar if for any $\sigma \in G$ and $x \in X = X^1$, then $\sigma|_x \in G$ also.*

The definition implies that $\sigma|_w \in G$ for all $w \in X^*$, since if $w = x_1x_2 \dots x_n$, then

$$\sigma|_w = ((\sigma|_{x_1})|_{x_2}) \dots |_{x_n}.$$

If $\sigma \in G$, for G a self-similar group, we will write

$$\sigma = (\sigma_1; \sigma|_{x_1}, \sigma|_{x_2}, \dots, \sigma|_{x_d}).$$

We now give some examples.

1. **The adding machine:** We define an automorphism a of \mathcal{T}_2 by the formula

$$a = (\tau; 1, a),$$

where τ is the unique non-identity element of the symmetric group \mathfrak{S}_2 (and 1 is the identity permutation). Despite first appearances, this definition is not circular, and does define images of $x \in X^\omega$ under a . For example, let $X = \{0, 1\}$. Then

$$\begin{aligned} a(11001100\dots) &= \tau(1)a(1001100\dots) \\ &= 0\tau(1)a(001100\dots) \\ &= 00\tau(0)01100\dots \\ &= 00101100\dots \end{aligned}$$

The group generated by a is an infinite cycle group with self-similar action on \mathcal{T}_2 . This example is called the adding machine, because it can be thought of as adding 1 to an infinite binary number (written in the opposite direction to normal).

2. **Products:** Suppose that σ, σ' are automorphisms of \mathcal{T}_X such that

$$\sigma = (\sigma_1; \sigma|_{x_1}, \sigma|_{x_2}, \dots, \sigma|_{x_d})$$

and

$$\sigma' = (\sigma'_1; \sigma'|_{x_1}, \sigma'|_{x_2}, \dots, \sigma'|_{x_d}).$$

Then observe, for any $1 \leq i \leq d$ and $\rho \in X^\omega$,

$$\sigma\sigma'(x_i\rho) = \sigma(\sigma'_1(x_i) \circ \sigma'|_{x_i}(\rho)) = (\sigma_1\sigma'_1(x_i)) \circ \sigma|_{\sigma'_1(x_i)}\sigma'|_{x_i}(\rho),$$

where \circ denotes concatenation. This tells us that $\sigma\sigma'$ can be written:

$$\sigma\sigma' = (\sigma_1\sigma'_1; \sigma|_{\sigma'_1(x_1)}\sigma'|_{x_1}, \dots, \sigma|_{\sigma'_1(x_d)}\sigma'|_{x_d}).$$

One can check similarly that:

$$\sigma^{-1} = (\sigma_1^{-1}; (\sigma|_{\sigma^{-1}(x_1)})^{-1}, \dots, (\sigma|_{\sigma^{-1}(x_n)})^{-1}).$$

In particular, to check that a group G is self-similar, it suffices to check that each element of a generating set has the self-similarity property.

3. **The Grigorchuk group:** This is a group of automorphisms of \mathcal{T}_2 generated by the following four elements, defined recursively:

$$a = (\tau; 1, 1)$$

$$b = (1; a, c)$$

$$c = (1; a, d)$$

$$d = (1; 1, b)$$

As before, τ is the non-identity element of \mathfrak{S}_2 . The Grigorchuk group is a finitely generated infinite torsion group (Burnside famously asked whether such groups could exist, and this is one of the simplest examples). It was also the first known group of intermediate growth, among other interesting properties. Groups with properties like Grigorchuk's are now called branch groups and are studied in their own right (see eg [10]).

Now let Γ_d be the directed graph with one vertex, v , and d edges with source and target v , with the edges labelled by X . Notice that the infinite paths of Γ_d are labelled by X^ω . This means that both the Leavitt path algebra $L_K(\Gamma_d)$ and the self-similar group G act on the set KX^ω of finitely supported K -valued functions on the ends (indeed $L_K(\Gamma_d)$ can be defined as an algebra of linear maps on KX^ω .) We define the algebra $L_K(G)$ to be the algebra of linear maps of KX^ω generated by $L_K(\Gamma_d)$ and G , and call $L_K(G)$ the Leavitt algebra of G .

We now study elements of $L_K(G)$. We write the generators of the Leavitt path algebra $L_K(\Gamma_d)$ as P_v, S_e, S_e^* rather than v, e, e^* . Importantly, the self-

similarity of G tells us that, for $x \in X$ and $\sigma \in G$, there exist $y = \sigma_1(x) \in X$, $\sigma|_x \in G$, such that

$$\sigma(x \circ \rho) = y \circ (\sigma|_x \rho),$$

for all $\rho \in X^\omega$. This implies the equation:

$$\sigma S_x = S_y \sigma|_x,$$

which holds in $L_K(G)$. Similarly, for any $\rho \in X^\omega$, and $y = \sigma(x)$ as before,

$$S_y^* \sigma(x\rho) = S_y^*(y\sigma|_x \rho) = \sigma|_x \rho = \sigma|_x S_x^*(x\rho),$$

and $S_y^* \sigma$ and $\sigma|_x S_x^*$ vanish on all ends not of the form $x\rho$. This implies:

$$S_y^* \sigma = \sigma|_x S_x^*.$$

These two results together imply that any element of $L_K(G)$ can be written as a linear combination of terms $S_p \sigma S_q^*$, where $p, q \in X^*$. As usual, this representation is not unique, because $\sum_{x \in X} S_x S_x^* = 1$ in the Leavitt algebra. This implies:

$$S_p \sigma S_q^* = \sum_{x \in X} S_p \sigma S_x S_x^* S_q^* = \sum_{x \in X} S_p S_{\sigma_1(x)} \sigma|_x S_x^* S_q^*.$$

As usual, we describe this process as a simple expansion of $S_p \sigma S_q^*$, and we use expansions to calculate products.

In [40], this construction is done with a completed group algebra \mathcal{A}_ϕ , and so a C^* algebra is produced rather than just a K -algebra. We won't take completions, and will give more details in the next section where we generalize to families of groups.

Finally we introduce Nekrashevych-Röver groups. This follows Section 9 in [40].

Definition 1.5.5. *Let G be a self-similar group of automorphisms of X^* (where X is a set of size d). Let $L_{\mathbb{C}}(G)$ be the Leavitt algebra of G over \mathbb{C} . Let V_G be the set of elements x of $L_{\mathbb{C}}(G)$ of the form $x = \sum_{i=1}^n S_{p_i} \sigma_i S_{q_i}^*$, where $p_i, q_i \in X^n$ and $\sigma_i \in g$, such that x is invertible, with $x^{-1} = \sum_{i=1}^n S_{q_i} \sigma_i^{-1} S_{p_i}^*$ (we say x is unitary). Then V_G is a group. The elements of V_G can be characterized as all elements of $L_{\mathbb{C}}(G)$ of the form $\sum_{i=1}^n S_{p_i} \sigma_i S_{q_i}^*$, where $\{p_i\}_{i=1}^n, \{q_i\}_{i=1}^n$ form leaf sets. Multiplication can be calculated by forming expansions. V_G is called the Nekrashevych-Röver group associated with G .*

Clearly, V_G contains a group isomorphic to the Higman-Thompson group $V_{1,d}$, as the set of its elements $\sum S_{p_i} S_{q_i}^*$. Like the Higman-Thompson group, V_G is a group of permutations of the set X^ω . Indeed, suppose $x \in V_G$ and $x = \sum S_{p_i} \sigma_i S_{q_i}^*$ as in the definition, and that $\rho \in X^\omega$. Then ρ can be uniquely written as $q_i \rho'$ for some i , and $x \cdot \rho = p_i \sigma_i(\rho')$.

We will draw elements of the Nekrashevych-Röver group with a bijection between finite trees (as for the usual Thompson group), but we will also record the element σ_i on the leaf p_i . An example is shown in Figure 1.16, which is a modification of Figure 1.12 to include elements of the adding machine group, generated by a as in Example 1. The bijection between leaf sets is shown by coloured triangles, which also are intended to visually resemble the subtree below that leaf, on which the adding machine acts.

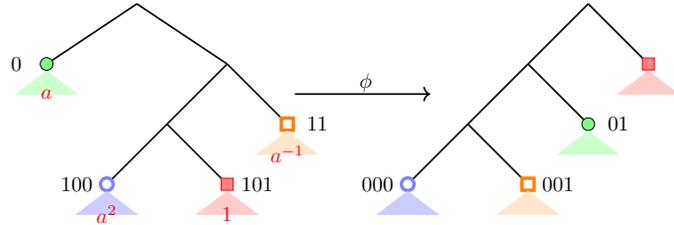


Figure 1.16: An element of a Nekrashevych-Röver group

As an example (with $X = \{0, 1\}$), we work out the image of $\rho = 11001100\dots$ under the group element shown. The Thompson bijection sends the initial segment 11 to 001 , whilst the remainder of ρ , which is $001100\dots$ is acted on by a^{-1} , and $a^{-1}(001100\dots) = 110100\dots$. So ϕ of Figure 1.16 maps $110011001100\dots$ to the end $001110100110011\dots$

In [40], elements of Nekrashevych-Röver groups V_G are often written with a different notation, as 3-row tables where if $S_{p_i} \sigma_i S_{q_i}^*$ is a summand of $X \in V_G$, then one column of a table for X contains p_i, σ_i and q_i in that order. For example, we could draw a table for ϕ as in Figure 1.16 as

$$\phi : \begin{pmatrix} 0 & 100 & 101 & 11 \\ a & a^2 & 1 & a^{-1} \\ 000 & 001 & 01 & 1 \end{pmatrix}$$

This notation is compact, but not as visual as drawing out trees. It can also be used just for Thompson's group V , where don't need to include the middle row

of the table.

Later in the thesis, groups will naturally arise that combine features of both colour-preserving Thompson groups and of Nekrashevych-Röver groups.

Neretin's group

It's worth pointing out the connection between Nekrashevych-Röver groups and a group introduced by Neretin in [42]. More modern introductions are given in [30] and [37]. Let \mathcal{T}_q be a q -regular tree (so that every vertex, even the root if we choose one, has degree q). Let $\text{Aut}(T)$ be the group of its tree automorphisms, which acts on $\partial\mathcal{T}_q$. Then the group N_q is most easily defined as the topological full group of $\text{Aut}(\mathcal{T}_q)$ acting on the boundary (see the next section for definitions). Informally, this is the group of all homeomorphisms of the boundary locally given by a tree automorphism.

Alternatively, choose a root of \mathcal{T}_q , and suppose that T_1, T_2 are full subtrees of \mathcal{T}_q . Then $F_1 = \mathcal{T}_q \setminus T_1$ and $F_2 = \mathcal{T}_q \setminus T_2$ are forests of rooted $q - 1$ -regular trees (where each vertex of depth k is a neighbour of $q - 1$ vertices of depth $k + 1$). Suppose that these two forests have the same number of trees each: say $F_1 = \{s_1, s_2, \dots, s_k\}$ and $F_2 = \{t_1, t_2, \dots, t_k\}$ where the s_i, t_i are the connected components of F_1, F_2 . Let ϕ be a forest isomorphism between them, which yields a homeomorphism $\bar{\phi}$ of $\partial\mathcal{T}_q$. Then the group N_q is also equal to the group of all such homeomorphisms $\bar{\phi}$ as T_1, T_2 and the isomorphism ϕ vary. It's easy to see that this is a Nekrashevych-Röver group, where the self-similar group G consists of all automorphisms of the rooted $(q - 1)$ -regular tree: indeed, ϕ must map each s_i to some t_j , but can apply a tree automorphism to it.

Neretin's group is an infinite simple group, and can be topologized to be a totally disconnected locally compact group with a compact presentation. There are of course many variants of Neretin's group that have been studied, including some isomorphic to Higman-Thompson groups. An interesting case comes from fixing a subgroup A of $\text{Aut}(T)$ and insisting that each restriction $\phi|_{s_i} : s_i \rightarrow t_j$ is given by an element of the subgroup A . This is a topological full group corresponding to A , and is studied more in [37] where it is related to Matui's colour-preserving Thompson groups. We explain the construction of topological full groups in the next section.

1.6 Topological full groups

Topological full groups of groupoids provide another way to define Thompson-like groups. In particular, Matui studied groupoids acting on Cantor sets in [39], which was how he defined the groups that we have been calling colour-preserving Thompson groups, $V_{\mathbb{C}}$. He constructed a topological full group isomorphic to $V_{\mathbb{C}}$, and showed that this group is of type F_{∞} and that its derived subgroup is simple. We will explain Matui's results and demonstrate how his topological full groups are isomorphic to the colour-preserving Thompson groups described earlier. We will give an overview of [39] before explaining which groupoids we are interested in. This completes the connections of Figure 1.10.

1.6.1 Ample groupoids and the topological full group

Let G be a topological groupoid with unit space $G^{(0)}$ homeomorphic to a Cantor set. We take our definitions from Nekrshevych's work in [41]; more general definitions are given in Matui's paper [39], where it is not assumed that $G^{(0)}$ is homeomorphic to a Cantor set. Recall that an open bisection is an open subset F of G on which the domain and range maps are both homeomorphisms (to subsets of $G^{(0)}$); F then defines a homeomorphism $\pi_F : d(F) \rightarrow r(F)$ by $x \mapsto r(d^{-1}|_F(x))$.

Recall also that we say G is *ample* if there is a basis for its topology consisting of *compact* open bisections. We will assume G is ample in this section. There is a product and an inversion on compact open bisections: if F_1, F_2 are compact open bisections, then so are:

$$F_1^{-1} = \{g^{-1} : g \in F_1\}$$

and

$$F_1 F_2 = \{g_1 g_2 : g_1 \in F_1, g_2 \in F_2, r(g_2) = d(g_1)\}.$$

This is enough to define the topological full group:

Definition 1.6.1. *Let G be an ample groupoid with unit space homeomorphic to a Cantor set X . The topological full group of G is the subgroup $[[G]]$ of the group of homeomorphisms of $G^{(0)}$, whose elements are of the form π_F for some compact open bisection F with $d(F) = r(F) = G^{(0)}$.*

If Y is a clopen subset of $G^{(0)}$, then we write $G|Y$ for the restriction of G

to Y (whose elements are $\{g \in G : d(g), r(g) \in Y\}$). This is also an ample groupoid, so the topological full group $[[G|Y]]$ is also defined.

We give an important example immediately.

Thompson's group V as a topological full group: Let $X = \{a, b\}$. We saw, when defining Thompson's group V , that X^ω labels the ends of an infinite binary tree whose vertices are labelled by the set X^* of finite words over X . We also saw that X^ω can be topologized with basic open sets $Z(w) = wX^\omega = \{w\rho : \rho \in X^\omega\}$, for each $w \in X^*$, so that X^ω becomes a Cantor space. If $\rho \in X^\omega$, we will write ρ_n for the n th character in the infinite word ρ (so that each $\rho_n \in X$). If Γ is the directed graph with one vertex and two edges a, b , then X^* is the same as Γ^* , the set of paths in Γ , and X^ω is the set of infinite paths in Γ . Take G to be the groupoid G_Γ as defined in Section 1.3.3, an ample groupoid. We claim $[[G]]$ is isomorphic to Thompson's group V .

Indeed, suppose $\phi \subset G$ is a compact open bisection with domain and range $G^{(0)} = X^\omega$. This set ϕ defines a permutation (in fact a homeomorphism) of $G^{(0)} = X^\omega$; we need to explain why this permutation is in Thompson's group. Indeed, since ϕ is open in the topology on G , then for each $x \in X^\omega$, there exist $\mu, \nu \in X^*$ such that $x \in \mu X^\omega$ and $Z(\mu, \nu) \subset \phi$. By compactness, a finite number of these $Z(\mu, \nu)$ cover ϕ . By repeatedly replacing $Z(\mu, \nu)$ with the two sets $Z(\mu a, \nu a)$ and $Z(\mu b, \nu b)$ and then deleting repetitions, we can assume the $Z(\mu, \nu)$ are disjoint, so

$$F = Z(\mu_1, \nu_1) \sqcup Z(\mu_2, \nu_2) \sqcup \dots \sqcup Z(\mu_n, \nu_n).$$

Then X^ω is the disjoint union of the domains $Z(\mu_1), \dots, Z(\mu_n)$, and is also the disjoint union of the ranges $Z(\nu_1), \dots, Z(\nu_n)$. In other words, $\{\mu_1, \dots, \mu_n\}$ and $\{\nu_1, \dots, \nu_n\}$ are leaf sets, and ϕ describes the permutation on X^ω where the end $\mu_i \rho$ is sent to $\nu_i \rho$ for each i . This is precisely what we need for an element of Thompson's group. Conversely, the same description gives a compact open bisection from any element of Thompson's group, so the two groups are isomorphic.

1.6.2 Properties of ample groupoids and topological full groups

Here we quote some theorems from [39] that tell us properties of ample groupoids. [39] states these results for étale groupoids rather than for ample groupoids, but since the results assume $G^{(0)}$ is a Cantor set, ample and étale groupoids are the same.

Definition 1.6.2. *Let G be a groupoid. The isotropy bundle G' of G is the subset*

$$G' = \{g \in G : d(g) = r(g)\}.$$

This definition implies that G' is a disjoint union of groups, one at each unit $x \in G^{(0)}$. We say that G is principal if $G' = G^{(0)}$, and (for G a topological groupoid) we say that G is essentially principal if the interior of G' equals $G^{(0)}$.

We now assume that G is essentially principal and that $G^{(0)}$ is a Cantor set.

Definition 1.6.3 ([39], Definition 4.9). *1. A clopen subset A of $G^{(0)}$ is properly infinite if there exist compact open bisections U, V , such that $d(U) = d(V) = A$, $r(U) \cup r(V) \subset A$, and $r(U) \cap r(V) = \emptyset$.*

2. G is purely infinite if every clopen subset A of $G^{(0)}$ is properly infinite. In particular, $G^{(0)}$ must be properly infinite.

Say that the topological groupoid G is *minimal* if for every $x \in G^{(0)}$, the G -orbit $Gx = \{r(g \cdot x) : g \in G\}$ is dense in $G^{(0)}$. The following is a summary of the main results on purely infinite groupoids from [39].

Theorem 1.6.4 ([39], Proposition 4.10, Theorem 4.16). *Let G be an essentially principal ample groupoid whose unit space is a Cantor set. If $G^{(0)}$ is properly infinite, then $[[G]]$ contains a subgroup isomorphic to the free product $\mathbb{Z}_2 * \mathbb{Z}_3$. In particular, $[[G]]$ is not amenable. If G is purely infinite and minimal, then the derived subgroup $[[G]]'$ is simple. Moreover, any subgroup of $[[G]]$ normalized by $[[G]]'$ contains $[[G]]'$.*

These results apply to many families of Thompson groups.

1.6.3 Groupoids from shifts and colour-preserving Thompson groups

We've been following [39], which studies groupoids of one-sided shifts in particular detail. We will restate the construction of [39] and then explain why this

gives the same groups as our colour-preserving Thompson groups.

Let Γ be a finite directed graph, with vertex set Γ^0 and edge set Γ^1 , and with source and target maps s and t . We will assume that for every two vertices v, w of Γ^0 , there is a path with source v and range w , and also that Γ is not a single cycle. We recall the set of infinite paths in Γ is:

$$\Gamma^\omega = \{e_1 e_2 e_3 \dots : e_i \in \Gamma^1, s(e_i) = t(e_{i+1}), \text{ all } i \in \mathbb{N}\}.$$

Notice that our paths are read right-to-left, as usual, so that $\rho = e_1 e_2 e_3 \dots$ has target $t(\rho) = t(e_1)$. The étale groupoid G for Γ is then the set:

$$G = \{(\rho, n, \rho') \in \Gamma^\omega \times \mathbb{Z} \times \Gamma^\omega : \rho_N = \rho'_{N+n} \text{ for all sufficiently large } N\}.$$

This works the same way as we saw for Thompson's group V . The multiplication is defined by $(\rho', n, \rho'')(\rho, m, \rho') = (\rho, m + n, \rho'')$ and undefined for other pairs, and the inversion is $(\rho, n, \rho')^{-1} = (\rho', -n, \rho)$. The unit space $G^{(0)}$ is associated with Γ^ω via $x \mapsto (x, 0, x)$. We topologize Γ^ω by taking basic open sets $Z(\mu)$ for each finite path μ in Γ ($\mu = f_1 f_2 \dots f_n$ where $f_i \in \Gamma^1$):

$$Z(\mu) = \{\rho \in \Gamma^\omega : \rho = e_1 e_2 e_3 \dots, e_1 = f_1, \dots, e_n = f_n\}.$$

When μ and ν are both finite paths with $s(\mu) = s(\nu)$, we define a subset $Z(\mu, \nu)$ of the groupoid G by:

$$Z(\mu, \nu) = \{(\mu\rho, n, \nu\rho) : \rho \in \Gamma^\omega, t(\rho) = s(\mu) = s(\nu), n = l(\nu) - l(\mu)\}.$$

As before, the sets $Z(\mu, \nu)$ form a base of open sets for a topology on G , with respect to which G is ample and the $Z(\mu, \nu)$ are compact open bisections. Finally, G is purely infinite and minimal ([39] Lemma 6.1).

Now we describe how the topological full group $[[G]]$ of this groupoid is (isomorphic to) a colour-preserving Thompson group. We've seen the argument before in Section 1.6.1, where we showed Thompson's V was a topological full group; the groupoid we described there comes from the directed graph with one vertex and two edges. By the same compactness argument as before, elements of $[[G]]$ can be written as a disjoint union of basic compact open bisections $Z(\mu, \nu)$, so we consider an element

$$g = Z(\mu_1, \nu_1) \cup Z(\mu_2, \nu_2) \cup \dots \cup Z(\mu_k, \nu_k),$$

and argue that the permutation it gives of G^ω corresponds to a permutation of the ends of an appropriate tree given by an appropriate colour-preserving Thompson group element. As before, we have that Γ^ω is a disjoint union of the domains $Z(\mu_i)$ and a disjoint union of the ranges $Z(\nu_j)$.

We define our set of colours \mathbf{C} to be Γ^0 (remember we assume Γ is finite). Set $\mathbf{S} = \Gamma^0$ also. For $c \in \mathbf{C}$, we take $p(c)$ to be the tuple $(s(e_1), s(e_2), \dots, s(e_m))$ where (e_1, e_2, \dots, e_m) are the edges with target v . We then observe that the vertices of $\mathcal{T}_{\mathbf{C}}$ are in bijection with finite paths in Γ , and the end set $\partial\mathcal{T}_{\mathbf{C}}$ is in bijection with the set Γ^ω of infinite paths in the graph Γ . Under this identification, the compact open bisection $Z(\mu_i, \nu_i)$ maps the ends below μ_i to the corresponding ends below ν_i . We also note that since $g \in [[G]]$, every end $\rho \in \Gamma^\omega$ lies in a unique $Z(\mu_i)$ and a unique $Z(\nu_j)$, which tells us that the μ_i and the ν_j form a leaf set. Thus, g is identified with the Thompson element defined by the bijection on finite trees sending μ_i to ν_i , both giving the same permutation on the ends. So the topological full group of a one sided shift is a colour-preserving Thompson group.

Conversely, suppose that \mathbf{C} is a set of colours with production rule p , and \mathbf{S} is a starting set. Further assume that if c, d are any two colours, then there exist a sequence of colours $c_0 = c, c_1, c_2, \dots, c_n = d$ where $n > 0$, and c_{i+1} appears in the tuple $p(c_i)$ for each i . We informally will say that this means every colour appears somewhere below every other colour. Take Γ as the directed graph with $\Gamma^0 = \mathbf{C}$ and for each colour c , take edges from each colour of $p(c)$ to c (with multiplicity). Assume we're not in the case where Γ is a single cycle (ie $\mathcal{T}_{\mathbf{C}}$ is not just a ray). Then there exist arbitrarily large sets of incomparable words in Γ , and we can find a set X_S of incomparable words with sources given by \mathbf{S} . Then considering the groupoid G for graph Γ restricted to the set X_S , we get that $[[G|X_S]]$ is a group isomorphic to the colour-preserving Thompson group $V_{\mathbf{C}, \mathbf{S}}$.

1.6.4 Theorems about groupoids of finite shifts

We now state some of Matui's results about topological full groups of the one-sided shift and rewrite them in terms of colour-preserving Thompson groups. Let Γ be a finite directed graph with étale groupoid G . Let its adjacency matrix be M ; that is, M is a matrix whose rows and columns are labelled by Γ^0 , and where:

$$M_{v,w} = |\{e \in \Gamma^1 : s(e) = w, t(e) = v\}|.$$

Theorem 1.6.5. *The following summarizes Chapter 6 of [39].*

1. *For any non-empty clopen subset $Y \subset G^{(0)}$, the groupoid $G|Y$ is purely infinite and minimal. Thus the derived subgroup $[[G|Y]]'$ is simple (and any subgroup of $[[G|Y]]$ normalized by $[[G|Y]]'$ contains $[[G|Y]]'$). In Thompson language, suppose that \mathbf{C} is a set of colours whose production rule p means every colour appears somewhere below each other colour, and $(\mathbf{C}, p) \neq (\{A\}, (A))$. Then the derived subgroup of $V_{\mathbf{C}}$ is simple, and any non-trivial subgroup normalized by the derived subgroup contains the derived subgroup.*
2. *In the case above, the abelianization $[[G|Y]]/[[G|Y]]'$ is isomorphic to $(H_0(G) \otimes \mathbb{Z}_2) \times H_1(G)$. The homology groups can be calculated from $H_0(G) = \text{Coker}(\text{Id} - M^T)$ and $H_1(G) = \text{Ker}(\text{Id} - M^T)$. Notice that this does not depend on Y , so in the language of Thompson groups, the abelianization does not depend on the starting set.*
3. *The group $[[G|Y]]$ is of type F_{∞} ; in particular, it is finitely presented, and an explicit presentation is given in [39] Section 6.6.*

Like Matui's previous result in Theorem 1.6.4, we will apply this to a generalization of colour-preserving Thompson groups.

1.6.5 Alternating and symmetric groups in a more general case

We conclude by summarizing the work of Nekrashevych in [41]. Let G be an ample groupoid whose unit space is a Cantor set.

Definition 1.6.6 ([41] Definition 3.1). *Let $d \in \mathbb{N}$. For each $1 \leq i \leq d$, let $F_{i,i}$ be a clopen subset of $G^{(0)}$, and suppose that these d sets are all disjoint. For each $1 \leq i, j \leq d$, let $F_{i,j}$ be a compact open bisection (in the case $i = j$, we remark that $F_{i,i}$ is already defined to be a compact open bisection). Suppose that for all i, j, k we have that $F_{j,k}F_{i,j} = F_{i,k}$. Then the family $F_{i,j}$ is a multisection of degree d . We say that the subsets $F_{i,i}$ of $G^{(0)}$ are the components of the domain of the multisection, and that $\cup_{i=1}^d F_{i,i}$ is the domain.*

In particular, a multisection of degree 2 is just a compact open bisection $F_{1,2}$, together with its inverse $F_{2,1}$, its domain $F_{1,1}$, and its range $F_{2,2}$. A

general multisection is just a collection of clopen subsets of $G^{(0)}$ with a family of compatible bisections between them.

If $\mathcal{F} = \{F_{i,j}\}_{i,j=1}^d$ is a multisection of degree d with domain U , and $\pi \in \mathfrak{S}_d$ is a permutation, then we can define an element \mathcal{F}_π of $[[G]]$ as follows:

$$\mathcal{F}_\pi = \cup_{i=1}^d F_{i,\pi(i)} \cup (G^{(0)} \setminus U).$$

So \mathcal{F}_π permutes the components of the domain of \mathcal{F} according to the permutation π , and is the identity elsewhere.

Definition 1.6.7 ([41] Definition 3.7). *Let \mathcal{F} be a multisection of the groupoid G . Observe that $\pi \mapsto \mathcal{F}_\pi$ is an embedding of \mathfrak{S}_n into $[[G]]$. We define $\mathfrak{S}(\mathcal{F})$ to be the image of this embedding, and $\mathfrak{A}(\mathcal{F})$ to be the image of the alternating group under this embedding. We write $\mathfrak{S}(G)$ (likewise $\mathfrak{A}(G)$) for the subgroup of $[[G]]$ generated by all subgroups $\mathfrak{S}(\mathcal{F})$ (or $\mathfrak{A}(\mathcal{F})$) as \mathcal{F} varies.*

Now recall that an étale groupoid G is *effective* if the interior of the isotropy bundle of G is $G^{(0)}$. Put another way, G is effective if for every non-unit $g \in G$ and every bisection F containing g , then for some $g' \in F$, we have $d(g) \neq r(g')$.

In particular, the groupoid of germs of any semigroup action (of S on a set X) is effective. Indeed, any open bisection has the form $[g, U]$ for some open $U \subset X$, and $g \in S$ defined on U . If $g(x) = x$ for all $x \in U$, then locally about $x \in U$, g is the identity, so the germ $[g, x]$ is a unit $[1, x]$.

Theorem 1.6.8 ([41] Theorem 4.1). *Suppose that G is minimal and effective. Then every non-trivial subgroup of $[[G]]$ normalized by $\mathfrak{A}(G)$ contains $\mathfrak{A}(G)$. In particular, $\mathfrak{A}(G)$ is simple.*

In particular, for Matui's groups of one-sided shifts, it can be verified that $\mathfrak{A}(G)$ is equal to the derived subgroup $[[G]]'$, which we've already seen is simple.

Chapter 2

Constructions of Leavitt-type structures for graphs of groups

We have now finished the introduction and we move to give some original constructions relating graphs of groups and Leavitt path algebras. In [13], the authors define a C^* -algebra associated to a graph of groups. This was first done for a finite graph of finite groups by Rui Okayasu in [43], and extended to a locally finite case in [13]. This C^* -algebra is similar in many ways to a directed graph C^* -algebra, but with a more complicated structure that both includes vertex group elements and forbids paths with backtracking. Theorem 4.1 of [13] describes the resulting C^* -algebra in terms of a crossed product of the fundamental group and the algebra of continuous functions on the boundary of the Bass-Serre tree. This gives what they describe as a C^* -algebraic Bass-Serre theorem, and they go on to study topological properties of the action of the fundamental group on the boundary of the tree.

It is natural to seek an algebraic version of this construction. This is what we accomplish in this section. In fact, we'll find graph-of-groups versions of all the objects in Figure 1.10: a groupoid of shifts, an inverse semigroup, a path algebra, and a variant of Thompson's group V . The Thompson groups we end up with will add tree automorphisms to Matui's colour-preserving Thompson groups (in the manner of Nekrashevych-Röver groups). These groups can be understood as topological full groups, but are perhaps better understood as a

new family of Thompson-like groups with a natural definition.

2.1 The path inverse semigroup of a graph of groups

We begin by defining an inverse semigroup associated to a graph of groups \mathcal{G} . We will define it as a set of partial bijections on the boundary of the Bass-Serre tree $\mathcal{T}_{\mathcal{G}}$ of \mathcal{G} . It will be similar to the graph inverse semigroup S_{Γ} of Definition 1.4.1, which we recall acts on infinite paths in Γ by changing a finite initial segment. Similarly, our inverse semigroup will act on the infinite \mathcal{G} -paths in a graph of groups \mathcal{G} by changing initial segments. However, it isn't possible to change initial segments freely, because this might result in $e\bar{e}$ appearing as a subword of a \mathcal{G} -word, preventing it from being reduced. This means we have to be more careful than in the directed graph case, and explicitly specify domains for the bijections which will prevent this from happening. Also, we allow vertex groups to act on the \mathcal{G} -words, which need not just affect a finite initial segment.

Let $\mathcal{G} = (\Gamma^0, \Gamma^1, G, \alpha)$ be a graph of groups, and for $e \in \Gamma^1$, let Σ_e be a set of coset representatives for $\alpha_e(G_e)$ in $G_{t(e)}$. In what follows, we will only consider graphs of groups \mathcal{G} that are locally finite: that is, the graph Γ is locally finite, and for each edge e , the group $\alpha_e(G_e)$ is of finite index in $G_{t(e)}$. We will also ask for \mathcal{G} to be non-singular: that is, if e is the unique edge with target $t(e)$, then $\alpha_e(G_e)$ is a proper subgroup of $G_{t(e)}$. This implies that every \mathcal{G} -path $g_1e_1 \dots g_n e_n$ can be extended to a \mathcal{G} -path $g_1e_1 \dots g_{n+1}e_{n+1}$ (where $e_i \in \Gamma^1, g_i \in \Sigma_{e_i}$). See [13], Section 2 for more details.

Write $\mathcal{T}_{\mathcal{G},v}$ for the Bass-Serre tree of \mathcal{G} rooted at vertex v . Recall that $\mathcal{T}_{\mathcal{G},v}$ has vertices labelled by the set $v\mathcal{G}^*$ of (finite) \mathcal{G} -paths, and ends labelled by the set $v\mathcal{G}^{\omega}$ of infinite \mathcal{G} -paths. The fundamental group $\pi_1(\mathcal{G}, v)$ acts on the tree, via action on the sets $v\mathcal{G}^*$ and $v\mathcal{G}^{\omega}$. In particular, G_v acts by tree automorphisms fixing the vertex v , so also acts on $v\mathcal{G}^n$ for any $n \in \mathbb{N}$. For $x \in v\mathcal{G}^n$, the action is by $g \cdot x = x'$ if $gx = x'g'$ in the fundamental groupoid of \mathcal{G} (where $x' \in v\mathcal{G}^n$ and $g' \in G_{s(x)}$). This extends to give a formula for the action on $v\mathcal{G}^{\omega}$, where $g \cdot g_1e_1g_2e_2 \dots$ is the unique element of \mathcal{G}^{ω} whose initial length n segment is $g \cdot g_1e_1 \dots g_n e_n$. It's easy to check that $g \cdot g_1e_1g_2e_2 \dots = g'_1e_1g'_2e_2 \dots$ for appropriate $g'_i \in \Sigma_{e_i}$. In other words, the vertex group elements change but the edges do not.

If $x = g_1e_1g_2e_2 \dots$ is a finite or infinite \mathcal{G} -path, we define the function l_1

by $l_1(x) = g_1e_1$, and we define $r_1(x) = e_n$ (we think of these as the left and right length 1 segments of x , although the definitions aren't symmetrical). Now we are ready to define an inverse semigroup. We define it as a sub-inverse-semigroup of $\mathcal{I}(\mathcal{G}^\omega)$, the inverse semigroup of all partial bijections on \mathcal{G}^ω . First we give a rough definition.

Outline: Let \mathcal{G} be a locally finite non-singular graph of groups. In this outline, we'll work with the example of a graph \mathcal{G} of trivial groups with one vertex and two edges, as in Figure 2.1. We will build an inverse semigroup of partial bijections of \mathcal{G}^ω , which consists of infinite words over the set $\{e, \bar{e}, f, \bar{f}\}$.

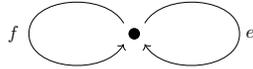


Figure 2.1: An example to motivate our inverse semigroup

Whenever μ and ν are finite \mathcal{G} -paths with the same source, we want a partial bijection that maps μx to νx , whenever μx is a \mathcal{G} -path. However, we're not guaranteed that νx is a valid \mathcal{G} -path, because it may be that $l_1(x) = \overline{r_1(\nu)}$ (and so some $e\bar{e}$ would appear in νx). For example, we may have $\mu = e, \nu = \bar{e}f$, when any x with $l_1(x) = ex'$ has μx a \mathcal{G} -path but νx not. To avoid this, we have to take the partial bijection that maps μx to νx , so long as $\overline{r_1(\nu)} \neq l_1(x)$. Call this partial bijection $B(\mu, \nu)$.

But now there is a new problem, because these partial bijections $B(\mu, \nu)$ are not multiplicatively closed. As an example, consider $B(f, \bar{f}e)B(\bar{f}e, f)$, the partial bijection that sends $\bar{f}ex$ to fx and back to $\bar{f}ex$. This is not the same as the identity on \mathcal{G} -paths $\bar{f}ex$, because it will not be defined if $l_1(x) = \bar{f}$. So instead of just taking bijections $B(\mu, \nu)$, we must take restrictions of $B(\mu, \nu)$ to sets $\{\mu x : l_1(x) \in P\}$, for $P \subset \mathcal{G}^1$. In the example, the product of bijections is only defined if $l_1(x)$ is equal to f or e , so we should take $P = \{e, f\}$.

Finally, our inverse semigroup will be more complicated because we will also have vertex groups. In general instead of bijections mapping $\mu \circ x$ to $\nu \circ x$, we will include bijections mapping it to $\nu \circ (g \cdot x)$. This means elements of the inverse semigroup will be specified by $\mu, \nu \in \mathcal{G}^*$, $g \in G_v$ for vertex v , and $P \subset \mathcal{G}^1$.

Definition 2.1.1. *Let \mathcal{G} be a locally finite non-singular graph of groups. The path inverse semigroup $S_{\mathcal{G}}$ associated to \mathcal{G} is defined as the following subset of $\mathcal{I}(\mathcal{G}^\omega)$: Take a symbol (μ, g, ν, P) whenever μ, ν, g, P satisfy the following conditions:*

(S1) $\mu, \nu \in \mathcal{G}^*$ with $s(\mu) = s(\nu)$.

Write $\mu = g_1 e_1 \dots g_k e_k$ and $\nu = h_1 f_1 \dots h_l f_l$, where as usual, e_i, f_j are edges of \mathcal{G} with $g_i \in \Sigma_{e_i}$ and $h_j \in \Sigma_{f_j}$.

(S2) $g \in G_{s(\mu)}$.

(S3) P is a subset of $\{g_e e : e \in \Gamma^1, t(e) = s(\mu), g_e \in \Sigma_e\} \subset \mathcal{G}^1$.

(S4) Let $f = f_l$, and let k be the element of $\Sigma_{\bar{f}}$ such that $gk \in \alpha_{\bar{f}}(G_f)$ (such a k exists and is unique by definition of $\Sigma_{\bar{f}}$ as a set of coset representatives). Then $k\bar{f} \notin P$.

$S_{\mathcal{G}}$ is the collection of all such symbols together with a zero. We refer to P as the permitted subset (for the element $(\mu, g, \nu, P) \in S_{\mathcal{G}}$).

As a partial bijection of \mathcal{G}^ω , the map (μ, g, ν, P) of $S_{\mathcal{G}}$ will have domain:

$$d((\mu, g, \nu, P)) = \mu P \mathcal{G}^\omega = \{\mu x : x \in \mathcal{G}^\omega, l_1(x) \in P\}.$$

The bijection will be:

$$\mu x \mapsto \nu \circ (g \cdot x),$$

for $x \in \mathcal{G}^\omega$, where \circ represents concatenation. Notice that property S4 tells us that if $r_1(\nu) = e$ and $l_1(x) = g_1 \bar{e}$ then $g \cdot g_1 \notin \alpha_e(G_e)$. This guarantees that the image of (μ, g, ν, P) will consist of infinite \mathcal{G} -paths - the subword $e\bar{e}$, which would stop the path being reduced, does not appear. So each element (μ, g, ν, P) does define a partial bijection of \mathcal{G}^ω .

In summary: (μ, g, ν, P) maps μx to $\nu \circ (g \cdot x)$, and is defined so long as $l_1(x) \in P$. If P is empty, then (μ, g, ν, P) is just another way to write \emptyset , the empty map. First we should show that $S_{\mathcal{G}}$ as defined is an inverse semigroup, so that it is closed under multiplication. In particular, this proof will illustrate the use of P , showing that we don't need to 'look ahead' more than one edge.

Proposition 2.1.2. *The inverse semigroup $S_{\mathcal{G}}$ is well-defined, in that it is a subset of $\mathcal{I}(\mathcal{G}^\omega)$ closed under products and inverses. The multiplication*

$$(\mu, g, \nu, P)(\rho, h, \sigma, Q)$$

is defined in various cases:

1. If $\sigma = \mu$, then the product is $(\rho, gh, \nu, Q \cap h^{-1}P)$. Notice that h^{-1} acts on $s(\sigma)\mathcal{G}^1$, so we can look at the image of P under this action.

2. Suppose $\mu = \sigma\mu'$ for non-empty \mathcal{G} -path μ' . In the fundamental groupoid, write $h^{-1}\mu' = \mu''k^{-1}$, for μ'' a \mathcal{G} -path and $k \in G_{s(\mu')}$. Suppose also that $l_1(\mu'') \in Q$. Then the product is $(\rho\mu'', gk, \nu, k^{-1}P)$.
3. Suppose $\sigma = \mu\sigma'$ for non-empty \mathcal{G} -path σ' . Suppose that $l_1(\sigma') \in P$. Write $g\sigma' = \sigma''k$ in the fundamental groupoid, where σ'' is a \mathcal{G} -path and $k \in G_{s(\sigma')}$. Then the product is $(\rho, kh, \nu\sigma'', Q)$.
4. The product is zero in all other cases. In particular, the product is zero whenever μ and σ are incomparable.

The inverse of (μ, g, ν, P) is $(\nu, g^{-1}, \mu, g \cdot P)$.

The proof is not difficult, but just requires working through the different cases carefully to compose the partial bijections. We use \circ for concatenation and \cdot for the fundamental groupoid action. We will see that the permitted sets P, Q are crucial in this proof, as described in the outline.

Proof. Consider the product $(\mu, g, \nu, P) \cdot (\rho, h, \sigma, Q)$, which is a composition of partial bijections of \mathcal{G}^ω . Suppose the product is non-zero, so that the image $\sigma \circ (h \cdot Q) \circ \mathcal{G}^\omega$ of (ρ, h, σ, Q) meets the domain $\mu \circ P \circ \mathcal{G}^\omega$ of (μ, g, ν, P) . In particular, the \mathcal{G} -paths σ and μ must be comparable. This means we split into three cases from now on, corresponding to 1-3 in the statement of the proposition.

Suppose first that $\mu = \sigma$. This is the case where we need permitted sets. Let $x \in \mathcal{G}^\omega$, and consider the image of $\rho \circ (h^{-1} \cdot x)$. The bijection (ρ, h, σ, Q) maps this \mathcal{G} -path to $\sigma x = \mu x$, and (μ, g, ν, P) then maps it to $\nu \circ (g \cdot x)$, so long as the permitted sets work out. This happens precisely when we have both $l_1(h^{-1} \cdot x) \in Q$ and $l_1(x) \in P$. The latter is equivalent to $l_1(h^{-1} \cdot x) \in h^{-1}P$. So we get the bijection $(\rho, gh, \nu, Q \cap h^{-1}P)$.

Suppose next that $\mu = \sigma\mu'$ for non-empty μ' . Then $\rho \circ (h^{-1} \cdot \mu'x)$ maps under (ρ, h, σ, Q) to $\sigma\mu'x = \mu x$, so long as $l_1(h^{-1}(\mu'x))$ lies in Q . Since $l_1(h^{-1}(\mu'x))$ only depends on μ' and h , not on x , this is either always true or never true as x varies. If it never holds, the product is zero. Otherwise, (μ, g, ν, S) maps $\sigma\mu'x$ to $\nu g \cdot x$, so long as $l_1(x) \in P$ (and so $l_1(k^{-1} \cdot x) \in k^{-1}P$). Overall, the product sends $\rho \circ h^{-1}(\mu'x) = \rho\mu'' \circ (k^{-1} \cdot x)$ to $\nu \circ (g \cdot x)$, and is not defined elsewhere. This gives the bijection $(\rho\mu'', gk, \nu, k^{-1} \cdot P)$ as claimed - we've checked that permitted sets are the same.

The final case is $\sigma = \mu\sigma'$ for non-empty σ' . Consider the image of ρx under the product (it being undefined elsewhere). (ρ, h, σ, Q) maps ρx to

$\sigma \circ (h \cdot x) = \mu\sigma' \circ (h \cdot x)$, so long as $l_1(x) \in Q$. The bijection (μ, g, ν, P) is then defined on this path if and only if $l_1(\sigma') \in P$, and if so, (μ, g, ν, P) maps $\mu\sigma' \circ (h \cdot x)$ to $\nu g \cdot (\sigma' \circ (h \cdot x)) = \nu\sigma'' \circ (kh \cdot x)$. So the product is $(\rho, kh, \nu\sigma'', Q)$ as claimed.

Inversion is easier. We've claimed the inverse of (μ, g, ν, P) is $(\nu, g^{-1}, \mu, g(P))$, and this maps $\nu \circ (g \cdot x)$ to μx so long as $l_1(g \cdot x) \in g \cdot P$, that is, $l_1(x) \in P$. It's easy to see this is an inverse. \square

This shows that $S_{\mathcal{G}}$ as defined is indeed an inverse semigroup. In particular, it's enough for the set P to be a subset of \mathcal{G}^1 - we never have to look ahead further than one edge beyond μ, ν . For example, consider the situation of Figure 2.1, and the product:

$$(\emptyset, 1, e, P_e)(\emptyset, 1, f, P_f)(f, 1, \emptyset, P_f)(e, 1, \emptyset, P_e),$$

where $P_e = \{e, f, \bar{f}\}$ and $P_f = \{e, f, \bar{e}\}$ (so that the domains are as large as possible). This product of elements removes e , then f , from the start of a path, and then replaces them. One might think that you need a permitted set of length 2 vectors to describe this. But in fact the domain of (μ, g, ν, P) is a function of μ as well as P , and the product is $(ef, 1, ef, P_f)$. This represents the identity on all \mathcal{G} -words beginning ef , with the permitted set remaining in \mathcal{G}^1 .

We will define the path groupoid as the groupoid of germs of this inverse semigroup. Since a groupoid of germs is defined in terms of local behaviour, we first study restrictions in $S_{\mathcal{G}}$. Recall that if S is an inverse semigroup, we say that $y \in S$ is a restriction of $z \in S$ if and only if $y = ez$ for e an idempotent of S (equivalently, $y = zf$ for f an idempotent of S). For an inverse semigroup of partial bijections, then y is a restriction of z if and only if y and z are equal everywhere y is defined.

We will see that restrictions can be made to have a particularly nice form.

Definition 2.1.3. *Let $S_{\mathcal{G}}$ be the path inverse semigroup of a graph of groups \mathcal{G} . Let $s = (\mu, g, \nu, P) \in S_{\mathcal{G}}$. We say that the expression (μ, g, ν, P) is full if $r_1(\mu) = r_1(\nu) = e$, say, if $g \in \alpha_{\bar{e}}(G_e)$, and if $P = s(e)\mathcal{G}^1 \setminus \{\bar{e}\}$.*

In particular, full elements have P as big as possible: indeed, P contains every length 1 path whose target is $s(e)$ except for the single path \bar{e} , which is forbidden by (S4), since this would cause $\mu\bar{e}$ not to be reduced. We show that any $s \in S_{\mathcal{G}}$ restricts around any x in its domain to a full element of $S_{\mathcal{G}}$, by first showing products of full elements are full.

Proposition 2.1.4. *Suppose that (μ, g, ν, P) and (ρ, h, σ, Q) are full elements of $S_{\mathcal{G}}$ whose product is non-zero. Then their product, as given in Proposition 2.1.2, is also full, as is their inverse.*

Proof. The case of inverses is easy: the inverse of (μ, g, ν, P) is $(\nu, g^{-1}, \mu, g \cdot P)$ and we certainly have that $r_1(\nu) = r_1(\mu) = e$, that $g^{-1} \in \alpha_e(G_e)$ (since g lies in this group), and then $g \cdot P = P = s(e)\mathcal{G}^1 \setminus \{\bar{e}\}$.

We work through the cases of Proposition 2.1.2 to study products. In the first case, $\sigma = \mu$, we have that $Q = P = h^{-1}P$, and so the product is (ρ, gh, ν, P) which it is easy to check is full. In the second case, where $\mu = \sigma\mu'$, we have that $r_1(\mu'') = r_1(\mu) = r_1(\nu)$, so the edge condition works. Also $g, k \in \alpha_{\bar{e}}(G_e)$ so their product also lies in this group, and moreover $k^{-1}P = P$. This is enough to show the product is full. Finally in the third case, we have that $kh \in G_{s(\sigma')} = G_{s(\rho)}$, and $r_1(\rho) = r_1(\sigma) = r_1(\sigma'')$, so this is full also. \square

We now show restrictions are also full.

Proposition 2.1.5. *Suppose that μ and $\mu\mu_1$ are both \mathcal{G} -paths, where μ_1 is non-empty. Let $z = (\mu, g, \nu, P) \in S_{\mathcal{G}}$. Then either z is defined nowhere on $\mu\mu_1\mathcal{G}^\omega$, or is defined on all of $\mu\mu_1\mathcal{G}^\omega$. In the latter case, the restriction of z to $\mu\mu_1\mathcal{G}^\omega$ is given by $(\mu\mu_1, g', \nu\mu'_1, A)$, where $g\mu_1 = \mu'_1g'$ in the fundamental groupoid, and $A = s(\mu_1)\mathcal{G}^\omega \setminus \{\bar{e}\}$, for $e = r_1(\mu_1)$. Moreover, if $\ell(\mu_1) > 0$, then $g' \in \alpha_e(G_e)$, and this restriction is full.*

Proof. This is straightforward: z is defined on $\mu\mu_1x$ (for $x \in \mathcal{G}^\omega$) if and only if $l_1(\mu_1) \in P$, a condition that doesn't depend on x . This gives the first statement, and also tells us that A should consist of all length 1 paths that can follow μ_1 in a \mathcal{G} -path, which is how A is defined. The restriction is then easy to calculate, since z sends $\mu\mu_1x$ to $\nu \circ (g \cdot \mu_1x) = \nu \circ \mu'_1 \circ (g' \cdot x)$, meaning that z locally acts by $(\mu\mu_1, g', \nu\mu'_1, A)$ as required. Finally $g' \in \alpha_{\bar{e}}(G_e)$ holds by the definition of multiplication in the fundamental groupoid, as does $r_1(\mu_1) = r_1(\mu'_1)$. This gives all the conditions needed for the restriction to be full. \square

Proposition 2.1.5 tells us that locally, any $s \in S_{\mathcal{G}}$ restricts to a full element of $S_{\mathcal{G}}$. Going forwards, it will often be helpful just to look at full elements of $S_{\mathcal{G}}$. This will normally result in no loss of generality, since as a partial function on \mathcal{G}^ω any $s \in S$ can be written as a disjoint union of restrictions which are full.

Now we characterize all restrictions.

Proposition 2.1.6. *Suppose that $z = (\mu, g, \nu, P) \in S_{\mathcal{G}}$ and that $y \in S_{\mathcal{G}}$ is a restriction of z . Then either y can be written in the form $y = (\mu, g, \nu, Q)$ for $Q \subset P$, or y can be written in the form $y = (\mu\mu_1, g', \nu\mu'_1, Q)$, where μ_1 is a non-empty \mathcal{G} -path with $g\mu_1 = \mu'_1g'$ in the fundamental groupoid, with $l_1(\mu_1) \in P$, and Q is unrestricted.*

We remark that we have carefully said y ‘can be written’ in the given form, but it might also have other expressions. For example, (μ, g, ν, P) and (μ, h, ν, P) might define the same element of $S_{\mathcal{G}}$.

Proof. Suppose $y = (\rho, h, \sigma, Q)$. Clearly, μ and ρ must be comparable, because otherwise the domains of y and z are disjoint. If $\rho = \mu$, then y is defined on all μx where $x \in \mathcal{G}^\omega$, $l_1(x) \in Q$, whereas z is defined on the set of all μx where $x \in \mathcal{G}^\omega$, $l_1(x) \in P$. Thus we must have $Q \subset P$, and since y and z are equal on the set where they’re both defined, we must have that $y = (\mu, g, \nu, Q)$.

Next suppose that $\rho = \mu\mu_1$ for non-empty $\mu_1 \in \mathcal{G}^*$. If we restrict z to the set $\rho\mathcal{G}^\omega$, then we get $(\mu\mu_1, g', \nu\mu'_1, A)$ where g', μ'_1, A are as in the previous proposition. Moreover, y must be a restriction of this bijection. This reduces to the $\mu = \rho$ case, so y can be written $(\mu\mu_1, g', \nu\mu'_1, Q)$ for some $Q \subset A$.

Finally suppose that $\mu = \rho\rho_1$ for non-empty $\rho_1 \in \mathcal{G}^*$. This is a slightly strange case, as then the domain of y seems larger than the domain of z , but is possible in some cases for appropriate Q . By the previous proposition, y must be defined on either all or none of $\mu\mathcal{G}^\omega$. Since y is a restriction of z , then z must be defined on all of $\mu\mathcal{G}^\omega$ and y must be zero elsewhere (or it could not be a restriction). Thus $y = z$ in this case. \square

We include one more useful fact about our inverse semigroups. It tells us when it is possible for there to be more than one full representation of an element of $S_{\mathcal{G}}$.

Proposition 2.1.7. *Let \mathcal{G} be a locally finite nonsingular graph of groups. Suppose that $[\mu, g, \nu, P]$ and $[\rho, h, \sigma, Q]$ are both equal to $s \in S$, and are both full, with $\ell(\mu) \leq \ell(\rho)$. If $\ell(\mu) = \ell(\rho)$, then $\mu = \rho, \nu = \sigma$ and $P = Q$. We need not have $g = h$; just that gh^{-1} acts trivially on $P\mathcal{G}^\omega$. If instead $\ell(\rho) > \ell(\mu)$, then $\rho = \mu\mu_1$, and ρ is the unique \mathcal{G} -path of its length extending μ . Moreover, writing $g\mu_1 = \mu'_1g'$ in the fundamental groupoid, then $\sigma = \nu\mu'_1$ and $h^{-1}g'$ acts trivially on $Q\mathcal{G}^\omega$.*

Here we write $P\mathcal{G}^\omega$ for the set of all $\rho \in \mathcal{G}^\omega$ such that $l_1(\rho) \in P$. The important thing to notice here is that full representatives of $s \in S$ are not unique if there are non-identity elements that act trivially on subtrees $P\mathcal{G}^\omega$.

Proof. We first do the case of $\ell(\rho) = \ell(\mu)$. Since (μ, g, ν, P) is defined on all infinite \mathcal{G} -paths extending μ , whereas (ρ, h, σ, Q) is defined on all infinite \mathcal{G} -paths extending ρ , we must have $\mu = \rho$. The two bijections then send μx to $\nu \circ (g \cdot x)$ and $\sigma \circ (h \cdot x)$ (for $x \in \mathcal{G}^\omega$) respectively. In particular, ν and σ must be comparable.

Now we claim that ν must equal σ . Indeed, suppose instead (without loss of generality) that $\ell(\sigma) = \ell(\nu) + k$, for $k > 0$. Consider how one would solve the equation

$$\nu \circ (g \cdot x) = \sigma \circ (h \cdot x).$$

The fact that $\nu \circ (g \cdot x)$ must start in σ determines the initial length k section of x . Then a length $\ell(\sigma) + k$ section of $\sigma \circ (h \cdot x)$ is determined, so we must fix an additional k edges of x for the left hand side to equal this. Repeating this, we determine x uniquely; this is a contradiction since we want the equation to hold for all $x \in \mathcal{G}^\omega$. So $\nu = \sigma$. Then P must equal Q since P is determined by μ for full $[\mu, g, \nu, P]$. We also see that g and h must act identically on all possible x that continue μ to the right; that is, on $P\mathcal{G}^\omega$.

Now suppose $\ell(\rho) > \ell(\mu)$. Again, considering the domains, we see that the sets $\rho\mathcal{G}^\omega$ and $\mu\mathcal{G}^\omega$ must be equal. This tells us ρ must be an extension of μ to the right, $\rho = \mu\mu_1$ (or the domains would not intersect) and that ρ is the unique such extension of its length (or $\rho\mathcal{G}^\omega$ would be a proper subset of $\mu\mathcal{G}^\omega$). Writing $\rho = \mu\mu_1$, we have that (ρ, h, σ, Q) is equal to $(\mu\mu_1, g', \nu\mu'_1, Q)$, the restriction of (μ, g, ν, P) to $\mu\mu_1\mathcal{G}^\omega$ by the description of restrictions in Proposition 2.1.6. Here g', μ'_1 are as described in the statement of the proposition. We've now reduced to the case $\ell(\rho) = \ell(\mu)$. By the case we've already done, we must have $\sigma = \nu\mu'_1$, and $h^{-1}g'$ acting trivially on $Q\mathcal{G}^\omega$, so we're done. \square

2.2 The path groupoid

We can now define a groupoid $G_{\mathcal{G}}$ as the groupoid of germs of $S_{\mathcal{G}}$. We call it the *path groupoid* of \mathcal{G} . Elements of $G_{\mathcal{G}}$ can be written $[s, x]$ for $s \in S_{\mathcal{G}}$ and $x \in d(s) \subseteq \mathcal{G}^\omega$, where we identify $[s, x]$ and $[s', x]$ if they have a common restriction around x . Since we know that restrictions can be chosen to be full,

we can assume s is a full element of $S_{\mathcal{G}}$ whenever this is helpful.

As a groupoid of germs, $G_{\mathcal{G}}$ has a topology. Recall that the topology on $G_{\mathcal{G}}$ has basic open sets

$$[s, U] = \{[s, x] : x \in U\},$$

for each $s \in S_{\mathcal{G}}$ and U an open subset of \mathcal{G}^{ω} contained in the domain of s . The topology on $G_{\mathcal{G}}$ need not be Hausdorff in general, but it's common for groupoids to be non-Hausdorff. It is however ample so long as \mathcal{G} is finite (in the sense that its underlying graph is finite and that its edge groups G_e are of finite index in the vertex groups under the embeddings α_e). Indeed, the open set $[s, U]$ is homeomorphic to U , so since \mathcal{G}^{ω} has a basis of compact open sets, $G_{\mathcal{G}}$ does also.

2.2.1 The grading on $G_{\mathcal{G}}$

We also point out that $S_{\mathcal{G}}$ is a \mathbb{Z} -graded inverse semigroup, meaning there is a map

$$\deg : S_{\mathcal{G}} \setminus \{0\} \rightarrow \mathbb{Z}$$

such that

$$\deg(st) = \deg(s)\deg(t)$$

whenever $s, t \in S, st \neq 0$ (cf [21] Section 2.1). We call $\deg(s)$ the degree of s . The degree is defined by $\deg((\mu, g, \nu, P)) = \ell(\nu) - \ell(\mu)$. One can check that this map is well-defined (so that if $(\mu, g, \nu, P) = (\mu', g', \nu', P')$ then $\ell(\nu) - \ell(\mu) = \ell(\nu') - \ell(\mu')$) and moreover that the degree of a restriction of $s \in S$ is the same as the degree of s . The degree of a product can be calculated via Proposition 2.1.2 and works in all cases. So this map does give a grading.

The groupoid $G_{\mathcal{G}}$ inherits the grading \deg from $S_{\mathcal{G}}$: for $s \in S$ and $x \in \mathcal{G}^{\omega}$, we define $\deg([s, x]) := \deg(s)$. It's clear that $\deg(yz) = \deg(y)\deg(z)$ whenever $(y, z) \in G_{\mathcal{G}}^{(2)}$, which is the condition for a grading on a groupoid.

2.3 The Leavitt graph of groups algebra

Now we define an algebra as the Steinberg algebra of $G_{\mathcal{G}}$. Here we will ask for Γ to be a *finite* graph, so that Γ^{ω} is compact (as a finite union of spaces $v\mathcal{G}^{\omega}$, each of which is the space of ends of a locally finite tree). Since $G_{\mathcal{G}}$ is then ample, we can define its Steinberg algebra.

Definition 2.3.1. *Let \mathcal{G} be a locally finite non-singular graph of groups whose*

underlying graph is finite. Let K be a field. We define the Leavitt graph of groups algebra $L_K(\mathcal{G})$ for K to be the Steinberg algebra of the ample groupoid $G_{\mathcal{G}}$ over K .

We will compare the algebra $L_K(\mathcal{G})$ to the C^* -algebra $C^*(\mathcal{G})$ of [13]. The C^* algebra is defined ([13] Definitions 3.1, 3.3) with generators and relations. The generators are a family of partial isometries $\{S_e : e \in \Gamma^1\}$ and a family of representations $g \mapsto U_{v,g}$ for each $v \in \Gamma^0$, such that $U_{v,g}$ is a partial unitary. The definitions of partial isometry, partial unitary and representation imply some relations between the $U_{v,g}$ and S_e , but we also add some further relations:

$$(G1) \quad U_{v,1}U_{w,1} = 0 \text{ for each } v, w \in \Gamma^0 \text{ with } v \neq w.$$

$$(G2) \quad U_{t(e),\alpha_e(g)}S_e = S_eU_{s(e),\alpha_e(g)} \text{ for each } e \in \Gamma^1 \text{ and } g \in G_e.$$

$$(G3) \quad U_{s(e),1} = S_e^*S_e + S_{\bar{e}}S_{\bar{e}}^* \text{ for each } e \in \Gamma^1.$$

$$(G4)$$

$$S_e^*S_e = \sum_{\substack{t(f)=s(e), h \in \Sigma_f \\ hf \neq 1\bar{e}}} U_{s(e),h}S_fS_f^*U_{s(e),h}^*,$$

for each $e \in \Gamma^1$.

There is indeed a C^* -algebra satisfying these relations (so the operators can be given a C^* -norm). The algebra is described in Remark 3.4 of [13], as a C^* -algebra acting on ℓ^2 -functions on \mathcal{G}^ω . Informally, the generators act on (\mathbb{C} -linear combinations of) \mathcal{G} -paths in the usual way: $U_{v,g}$ sends $x \in \mathcal{G}^\omega$ to $g \cdot x$ whenever $t(x) = v$, whereas $S_e x = ex$ whenever ex is a \mathcal{G} -path, and $S_e^*(ex) = x$. The generators are zero on other \mathcal{G} -paths.

We will produce a similar presentation for the algebra $L_K(\mathcal{G})$. First we remark that $L_K(\mathcal{G})$ has a similar action on K -valued functions on \mathcal{G}^ω , which it inherits from $G_{\mathcal{G}}$. Indeed, $L_K(\mathcal{G})$ is by definition an algebra of locally constant functions on $G_{\mathcal{G}}$ under convolution, so is spanned by indicator functions 1_B of open bisections B . If f is a K -valued function on \mathcal{G}^ω , then we define $1_B f$ by the formula $1_B f(x) = f(B^{-1}x)$ for all $x \in r(B)$, and $1_B f(x) = 0$ otherwise.

Now we state a presentation for $L_K(\mathcal{G})$ similar to the C^* presentation above. There are some extra relations, because we're giving an algebra presentation, not a $*$ -algebra presentation. Verifying that this presentation works takes a while, because there are many relations to check; although the proof is not that difficult, it will occupy the rest of the chapter.

Theorem 2.3.2. *Let $\mathcal{G} = (\Gamma^0, \Gamma^1, G, \alpha)$ be a locally finite non-singular graph of groups whose underlying graph is finite. For each edge e , let Σ_e be a transversal for the subgroup $\alpha_e(G_e)$ of $G_{t(e)}$, with $1 \in \Sigma_e$. Then the Leavitt graph of groups algebra $L_K(\mathcal{G})$ is isomorphic to the K -algebra given by presentation with generators $U_{v,g}, S_e, S_e^*$ (for all $v \in \Gamma^0, g \in G_v, e \in \Gamma^1$) and relations:*

(L1) $U_{v,g}U_{v,h} = U_{v,gh}$ for all $v \in \Gamma^0, g, h \in G_v$. In particular $U_{v,1}$ is idempotent.

(L2) $U_{v,1}U_{w,1} = 0$ whenever $v, w \in \Gamma^0$ and $v \neq w$.

(L3) $U_{t(e),1}S_e = S_e = S_eU_{s(e),1}$ and $U_{t(e),1}S_e^* = S_e^* = S_e^*U_{s(e),1}$, for all $e \in \Gamma^1$.

(L4)

$$U_{t(e),\alpha_e(g)}S_e = S_eU_{s(e),\alpha_{\bar{e}}(g)}$$

for each $e \in \Gamma^1$ and $g \in G_e$; also,

$$U_{s(e),\alpha_{\bar{e}}(g)}S_e^* = S_e^*U_{t(e),\alpha_e(g)}.$$

(L5) For $g \in \Sigma_e$, define $S_{ge} = U_{t(e),g}S_e$ and $S_{ge}^* = S_e^*U_{r(e),g^{-1}}$. Then if $e, f \in \Gamma^1, g \in \Sigma_e, h \in \Sigma_f$

$$S_{ge}S_{ge}^*S_{ge} = S_{ge}, S_{ge}^*S_{ge}S_{ge}^* = S_{ge}^*,$$

and

$$S_{ge}^*S_{hf} = 0$$

whenever $e \neq f$ or $g \neq h$.

(L6) For each $e \in \Gamma^1$,

$$S_e^*S_e = \sum_{h,f} S_{hf}S_{hf}^*,$$

where the sum is taken over all pairs (h, f) such that $t(f) = s(e), h \in \Sigma_f$, and $(h, f) \neq (1, \bar{e})$.

(L7) $U_{s(e),1} = S_e^*S_e + S_{\bar{e}}S_{\bar{e}}^*$ for each $e \in \Gamma^1$. Moreover, $S_eS_{\bar{e}} = S_e^*S_{\bar{e}}^* = 0$.

(L8) Suppose that $e \in \Gamma^1$ and $g \in G_e$ such that $\alpha_e(g)$ fixes $e\mathcal{G}^\omega$ pointwise. Then we add the relation

$$S_eU_{s(e),\alpha_{\bar{e}}(g)}S_e^* = S_eS_e^*.$$

This definition is a bit longer than the equivalent definition of a C^* -algebra in [13], but this is because we've written out explicitly relations that in the C^* -context are covered by saying that elements are partial isometries or partial unitaries. Also, we are giving an algebra presentation rather than a $*$ -algebra presentation, so we have to include $*$ -versions of each relation.

It is easiest to understand this presentation by thinking about the generators $U_{v,g}$, S_e , S_e^* acting on functions on \mathcal{G}^ω . Specifically, consider the image of the indicator function 1_x , for $x \in \mathcal{G}^\omega$. Each of the generators sends 1_x either to 0 or to another indicator function. That is, $U_{v,g}1_x = 1_{g \cdot x}$ whenever $t(x) = v$; $S_e(1_x) = 1_{ex}$ whenever $t(x) = s(e)$ and $l_1(x) \neq \bar{e}$; and S_e^* is a partial inverse to S_e , so $S_e^*(1_{ex}) = 1_x$. The generators send other indicator functions to 0. The mnemonic for this is that $U_{v,g}$ multiplies by group element g , S_e adds the edge e , and S_e^* removes it - whenever these operations are possible. Formally, this gives an action of $L_K(\mathcal{G})$ on the space $K\mathcal{G}^\omega$ of K -valued functions on \mathcal{G}^ω of finite support. All the relations (L1)-(L8) are easy to understand in terms of this action.

The important point of difference between $L_K(\mathcal{G})$ as given by this presentation and $C^*(\mathcal{G})$ is the inclusion of relation (L8). The C^* -algebra doesn't have a relation like this one, but we need it because Steinberg algebras are defined in terms of germs of actions, so if $\alpha_e(g)$ acts locally as the identity, we need to specify this. An earlier version of this work added a condition to the graphs of groups, asking for vertex groups not to act locally as the identity, but it seems better to include (L8) instead.

Before proving Proposition 2.3.2 we find some other presentations of $L_K(\mathcal{G})$. First we give a presentation for $L_K(\mathcal{G})$ using Definition 4.1 of [21], which gives a presentation for Steinberg algebras of a graded groupoid. The generators will be indicators 1_B of homogeneous compact open bisections. We say bisection B is homogeneous if all its elements have the same degree, and refer to this as the degree of the bisection.

Fact 2.3.3. *$L_K(\mathcal{G})$ is isomorphic to the K -algebra given by the following presentation: the generators are symbols t_B , whenever B is a homogeneous compact open bisection of $G_{\mathcal{G}}$, and the relations are:*

$$(R1) : t_\emptyset = 0.$$

$$(R2) : t_B t_D = t_{BD} \text{ for all homogeneous compact open bisections } B, D.$$

$$(R3) : t_B + t_D = t_{B \cup D} \text{ whenever } B, D \text{ are compact open degree zero bisections}$$

such that $B \cup D$ is also a compact open bisection (which necessarily also has degree zero)

The generator t_B in this presentation corresponds to the indicator 1_B in the Leavitt path algebra. We refine this presentation a bit further, and show that we can just take generators corresponding to elements of $S_{\mathcal{G}}$. Notice that for $s \in S$, then $B(s) = [s, d(s)]$ is a compact open bisection. We will find a presentation whose generators correspond to indicators $1_{B(s)}$ where $s \in S$ is full.

Proposition 2.3.4. *Let \mathcal{G} be a finite nonsingular graph of groups. Then the Leavitt graph of groups algebra $L_K(\mathcal{G})$ is isomorphic to the algebra L given by the following presentation: the generators are $\tau_{B(s)}$, for each $s \in S_{\mathcal{G}}$ which is full, and the relations are:*

$$(R1) \quad \tau_{\emptyset} = 0.$$

(R2') *Suppose that $s_1 = (\mu, g, \nu, P)$ and $s_2 = (\rho, h, \sigma, P)$ are both full expressions for elements of S . If σ and μ are incomparable take relation $\tau_{s_1}\tau_{s_2} = 0$. If σ and μ are comparable, then Proposition 2.1.4 tells us that s_1s_2 is also full. So we can take relation*

$$\tau_{s_1}\tau_{s_2} = \tau_{s_1s_2}.$$

(R3') *Let $s = (\mu, g, \nu, P)$ be full and of degree 0. For each $hf \in P$ (so that $f \in \Gamma^1, h \in \Sigma_f$), observe that the restriction s_{hf} of s to $\mu hf \mathcal{G}^\omega$ is full (by Proposition 2.1.5). Then take relation:*

$$\tau_s = \sum_{hf \in P} \tau_{s_{hf}}.$$

Proof. Let L be the algebra presented by relations (R1), (R2'), (R3'). We must find mutually inverse homomorphisms between L and $L_K(\mathcal{G})$. We will look at $L_K(\mathcal{G})$ either as the algebra presented by Fact 2.3.3, or as a Steinberg algebra, whichever is more suitable. Since relations (R1), (R2'), (R3') are special cases of (R1), (R2), (R3), they hold for $t_{B(s)}$ in $L_K(\mathcal{G})$. So there exists a homomorphism $\phi : L \rightarrow L_K(\mathcal{G})$ sending $\tau_{B(s)} \mapsto t_{B(s)}$.

Conversely, let B be a homogeneous compact open bisection, so that t_B is a generator of $L_K(\mathcal{G})$. We know that locally around each $x \in d(B)$, B restricts

to $B(s)$ for some full $s \in S_{\mathcal{G}}$. By compactness, B can be written as a union of finitely many bisections $B(s_i)$, for full $s_i = (\mu_i, g_i, \nu_i, P_i)$. After refining if necessary (replacing $B(s_i)$ with a disjoint union of its restrictions s_{hf} to $\mu_i hf \mathcal{G}^\omega$, for $hf \in P_i$) we can assume the union is disjoint. We will define $\psi(t_B) = \sum \tau_{s_i}$ whenever B is a disjoint union of full bisections B_i , and need to show that this does not depend on the choice of bisections.

Observe that any two expressions for B as a disjoint union of bisections $B(s)$ can be made equal by further refinements. Since (R3') allows us to expand degree zero bisections, we want to be able to relate any two expressions for B by this operation. Now consider a particular $s_i = s = (\mu, g, \nu, P)$. If $\ell(\mu) > \ell(\nu)$, write $\mu = \mu' \mu''$ where $\ell(\mu'') = \ell(\nu)$. Then we can factor:

$$s = s^{(0)} s^{(1)} = (\mu'', g, \nu, P)(\mu', 1, \emptyset, P'),$$

for $P' = \{l_1(\mu'')\}$. The point of this is that the first factor $s^{(0)} = (\mu'', g, \nu, P)$ is then degree zero. Moreover a restriction s_{hf} of s can be calculated as a restriction $s_{hf}^{(0)}$ of $s^{(0)}$ multiplied by $s^{(1)}$:

$$s_{hf} = (\mu hf, g', \nu h' f, Q) = (\mu'' hf, g', \nu h' f, Q)(\mu', 1, \emptyset, P'),$$

for appropriate $Q = s(f) \mathcal{G}^\omega \setminus \{\bar{f}\}$. A similar factorization exists if $\ell(\mu) < \ell(\nu)$.

The point of this is that the definition $\psi(t_B) = \sum \tau_{s_i}$, whenever B is a disjoint union of full bisections B_{s_i} is well-defined. Indeed, the previous paragraph tells us that any two such expressions can be equated by factoring and by expanding degree zero bisections: these operation are permitted by relations (R2'), (R3'). To show that this defines a homomorphism ϕ out of $L_K(\mathcal{G})$ we must check relations (R2), (R3) hold for the image of ϕ .

For (R2), given $\psi(t_B) = \sum \tau_{s_i}$, and $\psi(t_D) = \sum \tau_{s'_i}$, write $s_i = (\mu_i, g_i, \nu_i, P_i)$ and $s'_i = (\mu'_i, g'_i, \nu'_i, P'_i)$. Relation (R2') tells us that $\tau_{s_i s'_j} = \tau_{s_i} \tau_{s'_j}$ for each possible product, giving the result. (R3) is also easy, since partitioning B and D separately into bisections $B(s_i)$ for full s_i gives a partition of $B \cup D$. This establishes that ψ is a homomorphism.

It is clear that ϕ and ψ are inverses, giving the result. \square

Now we return to prove Theorem 2.3.2.

Proof. We find homomorphisms between the Leavitt graph of groups algebra and the algebra of Theorem 2.3.2.

The map from the presentation to the Steinberg algebra: Write $L_K(\mathcal{G})$ for the Leavitt graph of groups algebra, and write L for the K -algebra given by the presentation in Theorem 2.3.2. As before, we find mutually inverse homomorphisms between $L_K(\mathcal{G})$ and L .

First we find a homomorphism ϕ from L to $L_K(\mathcal{G})$. We will define ϕ as follows:

- For $v \in \Gamma^0$ and $g \in G_v$, we map $U_{v,g}$ to the open bisection $\phi(U_{v,g}) = [(\emptyset, g, \emptyset, v\mathcal{G}^1), v\mathcal{G}^\omega]$, which describes the map whose domain is all of $v\mathcal{G}^\omega$, and which sends x to $g \cdot x$.
- For $e \in \Gamma^1$, with $s(e) = v$, we map S_e to the open bisection $[(\emptyset, 1, e, v\mathcal{G}^1 \setminus \{\bar{e}\}), X]$, where $X = \{x \in v\mathcal{G}^\omega : l_1(x) \neq \bar{e}\}$. This open bisection maps the infinite \mathcal{G} -path x to ex , so long as $t(x) = s(e)$ and $l_1(x) \neq \bar{e}$, so that ex is reduced.
- For $e \in \Gamma^1$, we map S_e^* to the open bisection $[(e, 1, \emptyset, s(e)\mathcal{G}^1 \setminus \{\bar{e}\}), e\mathcal{G}^\omega]$. This open bisection is defined on \mathcal{G} -paths ex and maps them to x .

We have to check that these images satisfy relations (L1)-(L8). We will look at $L_K(\mathcal{G})$ as the Steinberg algebra, so talk about open bisections in it as functions on \mathcal{G}^ω . In particular we will be able to check that (L1)-(L8) hold by considering the action of $L_K(\mathcal{G})$ on points $x \in \mathcal{G}^\omega$ (seen as $1_x \in K\mathcal{G}^\omega$).

- (L1) The bisections $\phi(U_{v,g}), \phi(U_{v,h})$ are defined on all of $v\mathcal{G}^\omega$ and send $x \in v\mathcal{G}^\omega$ to $g \cdot x, h \cdot x$ respectively. The fact that $g \mapsto U_{v,g}$ is a representation follows from $x \mapsto g \cdot x$ being an action of G_v .
- (L2) This is immediate, since the domain $v\mathcal{G}^\omega$ of $\phi(U_{v,1})$ is disjoint from the range $w\mathcal{G}^\omega$ of $\phi(U_{w,1})$, so the product of the two bisections is zero.
- (L3) First we show that $\phi(U_{t(e),1})\phi(S_e)$ equals $\phi(S_e)$. This holds, because $\phi(U_{t(e),1})$ is the identity on its domain $t(e)\mathcal{G}^\omega$, which includes the range of $\phi(S_e), e\mathcal{G}^\omega$. The other parts are similar.
- (L4) We compare the open bisections on each side by letting them act on an infinite \mathcal{G} -path x . $\phi(S_e)$ is defined on x if and only if $x \in s(e)\mathcal{G}^\omega$ and $l_1(x) \neq \bar{e}$, in which case it is sent to ex . $\phi(U_{t(e),\alpha_e(g)})$ is then defined on ex , which it maps to $\alpha_e(g) \cdot ex$. In the fundamental groupoid, $\alpha_e(g)e = e\alpha_{\bar{e}}(g)$, so $\phi(U_{t(e),\alpha_e(g)})\phi(S_e)x = e\alpha_{\bar{e}}(g) \cdot x$.

On the other hand, it is clear that $\phi(S_e)\phi(U_{s(e),\alpha_{\bar{e}}(g)})$ also sends x to $e\alpha_{\bar{e}}(g) \cdot x$ wherever it is defined. This set is $\{x \in s(e)\mathcal{G}^\omega : l_1(\alpha_{\bar{e}}(g) \cdot x) \neq \bar{e}\}$

- since the unitary $\phi(U_{s(e),\alpha_{\bar{e}}(g)})$ is defined on all of $s(e)\mathcal{G}^\omega$, and the l_1 restriction comes from $\phi(S_e)$ needing be defined. But if $l_1(x) = \bar{e}$, then since $\alpha_{\bar{e}}(g)\bar{e} = \bar{e}\alpha_e(g)$ in the fundamental groupoid, $l_1(\alpha_{\bar{e}}(g) \cdot x) = \bar{e}$ if and only if $l_1(x) = \bar{e}$. Thus the domain is the same in both cases also. The second part can be deduced from the first by noticing that $\phi(S_e^*)$ is the inverse of $\phi(S_e)$ (as a partial bijection) and taking inverses of the first part with g^{-1} in place of g .

(L5) We study the functions $\phi(S_{ge}S_{ge}^*)$. First we consider $\phi(S_eS_e^*)$. This is defined on the set $e\mathcal{G}^\omega$, and is the identity on that set. Moreover, $\phi(S_{ge}S_{ge}^*)$ (for $g \in \Sigma_e$) is just a conjugate of this map by the action of $g \in G_v$, and so is the open bisection that is identity on the set $ge\mathcal{G}^\omega$ and undefined elsewhere. It is clear that the domain of $\phi(S_{ge}^*)$ is equal to the range of $\phi(S_{ge})$, and that $\phi(S_{ge}S_{ge}^*)$ is the identity on this set and undefined elsewhere, which gives the result. Finally, $\phi(S_{ge}^*)\phi(S_{hf})$ is the product of a bisection with domain $ge\mathcal{G}^\omega$ with a bisection with range $hf\mathcal{G}^\omega$. These sets are disjoint for $ge \neq hf$ so the product is zero.

(L6) The map $\phi(S_e^*S_e)$ adds e then removes it to all paths where this is possible - in other words, it is the identity on the set $\{x \in s(e)\mathcal{G}^\omega : l_1(x) \neq \bar{e}\}$. Recall that $\phi(S_{hf}S_{hf}^*)$ is the identity on the set $hf\mathcal{G}^\omega = \{x \in s(e)\mathcal{G}^\omega : l_1(x) = hf\}$. Since every $x \in s(e)\mathcal{G}^\omega$ has $l_1(x) = hf$ for some unique $f \in \Gamma^1$ with $t(f) = s(e)$, $h \in \Sigma_f$ we get the equality claimed.

(L7) This is easy from what we know: $\phi(U_{s(e),1})$ is the identity on the set $s(e)\mathcal{G}^\omega$ and undefined elsewhere. On the other hand, $\phi(S_e^*S_e)$ is the identity function on paths $x \in s(e)\mathcal{G}^\omega$ with $l_1(x) \neq \bar{e}$, whereas $\phi(S_{\bar{e}}S_{\bar{e}}^*)$ is the identity on paths $x \in s(e)\mathcal{G}^\omega$ where $l_1(x) = \bar{e}$. Taken together, these give the first result. The second part is clear since no \mathcal{G} path contains $e\bar{e}$.

(L8) Suppose that $e \in \Gamma^1$ and $g \in G_e$ such that $\alpha_e(g)$ fixes $e\mathcal{G}^\omega$ pointwise. Since $\alpha_e(g)e = e\alpha_{\bar{e}}(g)$ in the fundamental groupoid, we can say that $\alpha_{\bar{e}}(g)$ fixes pointwise all of the set

$$X = s(e)\mathcal{G}^\omega \setminus \bar{e}\mathcal{G}^\omega.$$

Now notice that $\phi(S_e^*)$ maps $e\mathcal{G}^\omega$ to X , which is fixed pointwise by $\phi(U_{s(e),\alpha_{\bar{e}}(g)})$ by the above. So we immediately get $\phi(S_eS_e^*) = \phi(S_eU_{s(e),\alpha_{\bar{e}}(g)}S_e^*)$.

This completes the proof in one direction.

The map in the other direction: Conversely, we must define a homomorphism ψ from $L_K(\mathcal{G})$ to L . First we fix notation: in L , if $p = g_1 e_1 \dots g_n e_n$ is a \mathcal{G} -path (with $e_i \in \Gamma^1$, and $g_i \in \Sigma_{e_i}$ as normal), where $g_i \in G_{v_i}$, write

$$S_p = U_{v_1, g_1} S_{e_1} \dots U_{v_n, g_n} S_{e_n},$$

and

$$S_p^* = S_{e_n}^* U_{v_n, g_n^{-1}} \dots S_{e_1}^* U_{v_1, g_1^{-1}}.$$

Then we define the homomorphism ψ by setting $\psi(1_B) = S_\nu U_{s(\mu), g} S_\mu^*$, whenever $B = B(s)$ for full $s \in S$, with $s = (\mu, g, \nu, P)$. We're going to use the result of Proposition 2.3.4 which tells us that such $1_{B(s)}$ generate. We must check various properties of this map.

ψ is uniquely defined: Here we show that the definition of ψ does not depend of how s is written, where we have $B = B(s)$. Indeed, suppose that $[\mu, g, \nu, P]$ and $[\rho, h, \sigma, Q]$ are two different full ways of writing the same element of S , with $\ell(\rho) \geq \ell(\mu)$ say. Proposition 2.1.7 tell us that $\rho = \mu \mu_1$, and that if $g \mu_1 = \mu'_1 g'$ in the fundamental groupoid, then $\sigma = \nu \mu'_1$ and $h^{-1} g'$ acts trivially on $Q\mathcal{G}^\omega$; moreover, ρ is the unique \mathcal{G} -path of its length extending μ . Because of this uniqueness, we must have $\mu_1 = \mu'_1$. We need to show that:

$$S_\mu U_{s(\mu), g} S_\nu^* = S_\rho U_{s(\rho), h} S_\sigma^*$$

or written another way

$$S_\mu U_{s(\mu), g} S_\nu^* = S_\mu S_{\mu_1} U_{s(\rho), h} S_{\mu'_1}^* S_\nu^*.$$

Write $r_1(\mu) = r_1(\nu) = e$. We will show:

$$S_e U_{s(\mu), g} S_e^* = S_e S_{\mu_1} U_{s(\rho), h} S_{\mu'_1}^* S_e^*.$$

Let $\mu_1 = g_1 e_1 \dots g_k e_k$ in usual notation. By assumption (g_1, e_1) is the only pair (h, f) satisfying the condition of (L6), so that $S_e^* S_e = S_{g_1 e_1} S_{g_1 e_1}^*$. By (L5), we have $S_e^* = S_e^* S_e S_e^*$, and hence this equals $S_{g_1 e_1} S_{g_1 e_1}^* S_e^*$. Repeat-

ing this argument, we see that

$$\begin{aligned}
S_{g_1 e_1} S_{g_1 e_1}^* &= S_{g_1 e_1} U_{s(e_1), 1} S_{g_1 e_1}^* \\
&= S_{g_1 e_1} (S_{\bar{e}_1} S_{\bar{e}_1}^* + S_{g_2 e_2} S_{g_2 e_2}^*) S_{g_1 e_1}^* \\
&= S_{g_1 e_1 g_2 e_2} S_{g_1 e_1 g_2 e_2}^*.
\end{aligned}$$

Here we've used again that (g_2, e_2) is the only pair satisfying the condition of (L6) for the edge e_1 . Inductively, we get

$$S_e^* = S_{\mu_1} S_{\mu_1}^* S_e^*.$$

So we see that

$$S_e U_{s(e), g} S_e^* = S_e U_{s(e), g} S_{\mu_1} S_{\mu_1}^* S_e^*.$$

Finally, repeatedly applying (L4) repeatedly shows that $U_{s(e), g} S_{\mu_1} = S_{\mu_1'} U_{s(\mu_1'), g'}$. Thus it remains to see that

$$S_{\mu_1} U_{s(\mu_1), g'} S_{\mu_1}^* = S_{\mu_1} U_{s(\mu_1), h} S_{\mu_1}^*,$$

for which it's enough to show that

$$S_f U_{s(f), g'} S_f^* = S_f U_{s(f), h} S_f^*,$$

where $f = r_1(\mu_1)$. Write $g' = \alpha_{\bar{f}}(x)$, $h = \alpha_{\bar{f}}(y)$ (which is possible by assumption that the elements are full). Then we can rearrange the statement to be proved to

$$S_f S_f^* = S_f U_{s(f), \alpha_{\bar{f}}(x^{-1}y)} S_f^*,$$

which is true by relation (L8).

(R1) holds To show that ψ can be defined as a homomorphism, we need that (R1)-(R3) hold on the image of ψ . First, (R1) is immediate if we define ψ to send 0 to 0.

(R2') holds Suppose that (μ, g, ν, P) and (ρ, h, σ, Q) are full elements of S_G . First suppose σ, μ are incomparable. We need to show that

$$S_\nu U_{s(\nu), g} S_\mu^* S_\sigma U_{s(\rho), h} S_\rho^* = 0.$$

It's enough to show $S_\mu^* S_\sigma = 0$. Suppose $\mu = g_1 e_1 \dots g_n e_n$ and $\sigma =$

$g'_1 e'_1 \dots g'_m e'_m$ in usual notation. If $g_1 e_1 \neq g'_1 e'_1$ then $S_\mu^* S_\sigma = 0$ by (L5). If instead $g_1 e_1 = g'_1 e'_1$, then by (L6),

$$S_{g_n e_n}^* \dots S_{g_1 e_1}^* S_{g'_1 e'_1} \dots S_{g'_m e'_m} = \sum_{h,f} S_{g_n e_n}^* \dots S_{g_2 e_2}^* S_{hf} S_{hf}^* S_{g'_2 e'_2} \dots S_{g'_m e'_m},$$

summed over appropriate pairs (h, f) . By (L5) again, the only term that is non-zero is when $(h, f) = (g_2, e_2)$, giving $S_{g_n, e_n}^* \dots S_{g_2, e_2}^* S_{g'_2 e'_2} \dots S_{g'_m, e'_m}$. This gives the result, inductively on n .

Now we consider σ and μ comparable; first suppose that $\sigma = \mu$. We wish to examine

$$S_\nu U_{s(\nu),g} S_\mu^* S_\mu U_{s(\mu),h} S_\rho^*.$$

First we argue $S_\mu^* S_\mu = S_e^* S_e$, where $e = r_1(\mu)$. This will be similar to the argument that $S_\mu^* S_\sigma$ vanishes for μ and σ incomparable. We induct on length of μ , the base case being $S_{g_e}^* S_{g_e} = S_e^* S_e$ which is clear. For the inductive step, write $\mu = g_1 e_1 \mu'$ (where $g_1 e_1 = l_1(\mu)$). We have:

$$\begin{aligned} S_\mu^* S_\mu &= S_{\mu'}^* S_{g_1 e_1}^* S_{g_1 e_1} S_{\mu'} = S_{\mu'}^* S_{e_1}^* S_{e_1} S_{\mu'} \\ &= S_{\mu'}^* (U_{s(e_1),1} - S_{\bar{e}_1} S_{\bar{e}_1}^*) S_{\mu'} \\ &= S_{\mu'}^* U_{s(e_1),1} S_{\mu'} \\ &= S_{\mu'}^* S_{\mu'} \end{aligned}$$

where we have used that $\ell_1(\mu') \neq \bar{e}_1$. This completes the induction.

We notice that since $g \in \alpha_{\bar{e}}(G_e)$, then $U_{s(\nu),g}$ commutes with $S_e^* S_e = S_\mu^* S_\mu$, by (L4). Thus $S_\nu U_{s(\nu),g} S_\mu^* S_\mu U_{s(\mu),h} S_\rho^*$ is equal to $S_\nu S_e^* S_e U_{s(\mu),gh} S_\rho^*$. Finally we use $S_e S_e^* S_e = S_e$ (L5) to replace $S_\nu S_e^* S_e$ with S_ν (noting that $r_1(\nu) = e$ since (μ, g, ν, P) is full). This completes the proof in the case $\sigma = \mu$.

Finally, suppose that $\mu = \sigma \mu'$ (the case $\sigma = \mu \sigma'$ is similar). We need to study

$$S_\nu U_{s(\nu),g} S_{\mu'}^* S_\sigma^* S_\sigma U_{s(\sigma),h} S_\rho^*.$$

Similarly to before,

$$\begin{aligned}
S_{\mu'}^* S_{\sigma}^* S_{\sigma} &= S_{\mu'}^* S_{r_1(\sigma)}^* S_{r_1(\sigma)} \\
&= S_{\mu'}^* \left(U_{s(\mu'),1} - S_{r_1(\sigma)} \frac{S_{r_1(\sigma)}^*}{S_{r_1(\sigma)}} \right) \\
&= S_{\mu'}^*.
\end{aligned}$$

Thus the product becomes

$$S_{\nu} U_{s(\nu),g} S_{\mu'}^* U_{s(\sigma),h} S_{\rho}^*,$$

and we have seen that $S_{\mu'}^* U_{s(\sigma),h} = U_{t(\mu'),h'} S_{\mu''}^*$, where $\mu' h = h' \mu''$ in the fundamental groupoid, by applying (L4) repeatedly. This is enough to give the correct product.

(R3') holds This is a consequence of (L6). Specifically, if $s = (\mu, g, \nu, P)$, then

$$\psi(\tau_s) = S_{\mu} U_{s(\mu),g} S_{\nu}^* = S_{\mu} U_{s(\mu),1} U_{s(\mu),g} S_{\nu}^*.$$

Expanding $U_{s(\mu),1}$ using (L6) and (L7) we get

$$\sum_{h,f} S_{\mu} S_{hf} S_{hf}^* U_{s(\mu),g} S_{\nu}^*,$$

a sum over all length 1 \mathcal{G} -paths hf with $t(f) = s(\mu)$. Write $r_1(\mu) = r_1(\nu) = e$. Then we have that $S_e^* S_{\nu}^* = 0$ by (L7), so the sum can be taken over all such hf with $hf \neq \bar{e}$. That is, the sum is over P . Using (L4), we have once again that

$$S_{hf}^* U_{s(\mu),g} = U_{s(f),g'} S_{h'f}^*,$$

whenever $ghf = h'f'g'$ in the fundamental groupoid (with $h' \in \Sigma_f$). Now if we write out s_{hf} as $(\mu hf, g', \nu h'f, P')$ we realize that

$$\psi(\tau_{s_{hf}}) = S_{\mu} S_{hf} U_{s(f),g'} S_{h'f}^* S_{\nu}^*,$$

which gives that $\psi(\tau_s)$ is a sum of terms $\psi(\tau_{s_{hf}})$, as required. This completes the verification of these identities.

It remains to show that the two maps we've defined between L and $L_K(\mathcal{G})$ are mutually inverse. This is easy to check, since the image of each of $S_e, S_e^*, U_{v,g}$ is

a full open bisection under ϕ , on which ψ is defined directly. We omit the details here. \square

We will often compare $L_K(\mathcal{G})$ to a Leavitt path algebra but with group actions added. The following proposition relates $L_K(\mathcal{G})$ to a Leavitt path algebra more formally. The analogous result in the C^* -algebra case was proved as Theorem 3.6 of [13]. Our result is almost the same.

Proposition 2.3.5. *Let \mathcal{G} be a locally finite non-singular graph of groups whose underlying graph Γ is finite. Let $L_K(\mathcal{G})$ be its Leavitt graph of groups algebra, with generators $U_{v,g}, S_e, S_e^*$ as normal. Let $E_{\mathcal{G}}$ be the directed graph whose vertex set is \mathcal{G}^1 and whose edge set is \mathcal{G}^2 ; for $\mu = g_1 e_1 g_2 e_2 \in \mathcal{G}^2$, we put $s_E(\mu) = g_2 e_2$ and $t_E(\mu) = g_1 e_1$. This is a finite directed graph. There is an embedding*

$$\phi : L_K(E_{\mathcal{G}}) \hookrightarrow L_K(\mathcal{G}),$$

defined on generators by:

$$\phi(\nu) = S_{\nu} S_{\nu}^*, \phi(\mu) = S_{\mu} S_{s(\mu)}^*, \phi(\mu^*) = S_{s(\mu)} S_{\mu}^*,$$

for all $\nu \in E_{\mathcal{G}}^0, \mu \in E_{\mathcal{G}}^1$. Moreover if all edge groups G_e are trivial, then the homomorphism ϕ is an isomorphism.

Proof. We sketch the proof, which follows the argument of Theorem 3.6 of [13], but instead of using the gauge-invariant uniqueness theorem, we use the graded uniqueness theorem of Theorem 1.2.6 for Leavitt path algebras.

To show that ϕ is a well-defined homomorphism, we need to just check that the images of the generators satisfy the defining relations for Leavitt path algebras, given in 1.2.2. These can be seen to follow from (L1)-(L8) in the definition of $L_K(\mathcal{G})$, in a similar way to the proof in [13]. The homomorphism ϕ is graded, so we can use the Graded Uniqueness Theorem to show it is injective. It is enough that it does not vanish on any $\nu \in E_{\mathcal{G}}^0$, so we need to show that $S_{\nu} S_{\nu}^*$ is non-zero in $L_K(\mathcal{G})$. This is clear since $S_{\nu} S_{\nu}^*$ is non-vanishing in its action on $K\mathcal{G}^{\omega}$ (specifically, it fixes $\nu\mathcal{G}^{\omega}$). In the case where edge groups are trivial, the surjectivity argument is the same as in [13], just with a third case added to show $S_{g_e}^*$ is in the image, which works the same as the S_{g_e} case. Namely, $\sum_{hf} \phi(geh f) = S_{g_e}$ (where $geh f \in E_{\mathcal{G}}^1$ and we sum over all possible $hf \neq 1\bar{e}$),

so S_{ge} is in the image of ϕ , and then

$$U_{x,g} = U_{x,g} \sum_{hf} S_{hf} S_{hf}^* = \sum_{hf} S_{ghf} S_{hf}^*,$$

so that $U_{x,g}$ is also in the image. \square

2.4 Colour-preserving Nekrashevych-Röver groups and the group of unitaries

We now produce a Thompson group variant from the algebra $L_K(\mathcal{G})$, which will also equal the topological full group of $G_{\mathcal{G}}$. We first define a family of variants of Thompson's group V , which will generalize both the colour-preserving Thompson groups and the Nekrashevych-Röver groups we discussed in the introduction. We then explain how the topological full group of $G_{\mathcal{G}}$ is a Thompson variant of this type. This completes the picture of Figure 1.10.

2.4.1 Self-similar families of groups

We will work in the situation of Section 1.5.5. Let \mathbf{C} be a finite set of colours equipped with production rule p and starting set \mathbf{S} . Let $\mathcal{T}_{\mathbf{C}}$ be the family of coloured trees associated to \mathbf{C} as in Section 1.5.5, writing $\mathcal{T}_{\mathbf{C}} = (\mathcal{T}_1, \dots, \mathcal{T}_m)$. Let \mathcal{T}_i have root v_i , its unique depth 0 vertex, so that the depth n vertices of \mathcal{T}_i are labelled by sequences $v_i x_{c_1, i_1} \dots x_{c_n, i_n}$. Let $\chi : T^0 \rightarrow \mathbf{C}$ be the function that gives the colour of each vertex.

Now suppose that e is an edge of \mathcal{T}_i , directed away from the root v_i and let T_e be the maximal subtree of \mathcal{T}_i containing $t(e)$ but not $s(e)$. Observe that as coloured trees, T_e is isomorphic to T_f whenever $t(e)$ and $t(f)$ have the same colour. So for $c \in \mathbf{C}$, we can fix a coloured tree T_c such that T_e is isomorphic to T_c whenever $\chi(t(e)) = c$, and we choose an isomorphism $\beta_e : T_c \rightarrow T_e$.

With this in place, let ϕ be an automorphism of T_e that preserves the colouring. Then ϕ permutes the set $\{e_1, e_2, \dots, e_k\}$ of edges whose source is $t(e)$, by a permutation $\sigma_e \in \mathfrak{S}_k$, and so ϕ also permutes the subtrees T_{e_i} . But since ϕ preserves the colouring, then $\chi(t(e_i)) = \chi(t(\phi(e_i)))$. By definition of self-similar colourings, T_{e_i} and $T_{\phi(e_i)}$ are both isomorphic to T_c (where $c = c(e_i)$). So the formula

$$\phi_i := \beta_{\phi(e_i)}^{-1} \phi \beta_{e_i}$$

defines a colour-preserving automorphism of T_c . We write:

$$\phi = (\sigma_\phi; \phi_1, \phi_2, \dots, \phi_k),$$

remarking that ϕ is determined by the permutation σ_ϕ and the automorphisms ϕ_i . This is modeled after Definition 1.5.4 where we defined a self-similar action of a single group.

Definition 2.4.1. *Let $\mathcal{T}_{\mathbf{C}}$ be a self-similar family of coloured trees as above. For each colour $c \in \mathbf{C}$, let G_c be a colour-preserving group of automorphisms of T_c that fix the root. For $g \in G_c$, write:*

$$g = (\sigma_g; g_1, g_2, \dots, g_k),$$

where g_i is an automorphism of $T_{c(i)}$. If it happens that $g_i \in G_{c_i}$ for all choices of c, g and i , then we say that the G_c form a self-similar family of groups.

There are two important examples, in both of which we will take \mathbf{S} of size 1, so that $\mathcal{T}_{\mathbf{C}, \mathbf{S}}$ consists of a single tree:

1. If $|\mathbf{C}| = 1$, then all vertices of $\mathcal{T}_{\mathbf{C}}$ must be given the same colour, c . Assume for simplicity that $|\mathbf{S}| = 1$ also, so that $\mathcal{T}_{\mathbf{C}}$ consists of a single tree T , and assume that the production rule p has $p(c) = (c, c, \dots, c)$ where c appears d times. Then T is the d -regular rooted tree.

Now suppose $\{G\}$ is a self-similar family of groups (since \mathbf{C} is a singleton set, there is only one group in the family). G is then a group of automorphisms of T . It has the property that if $g \in G$, and we write:

$$g = (\sigma_g; g_1, g_2, \dots, g_d),$$

then each $g_i \in G$. This exactly says that G is a self-similar group in the sense of Nekrashevych.

2. Let \mathcal{G} be a locally finite graph of groups whose underlying graph is finite. We draw an example in Figure 2.3. For each $v \in \Gamma^0$, form the Bass-Serre tree $T = \mathcal{T}_{\mathcal{G}, v}$ rooted at v . Let \mathbf{C} be the set Γ^1 of edges of Γ (with each edge and its reverse included separately). Suppose that p is a \mathcal{G} -path with target v , so we can write $p = g_1 e_1 \dots g_n e_n$, where $e_i \in \Gamma^1$ with $s(e_i) = t(e_{i+1})$ wherever this applies, and $g_i \in \Sigma_{e_i}$. Set $p' = g_1 e_1 \dots g_{n-1} e_{n-1}$. Then

there is an edge e in T from $p'G_{s(p')}$ to $pG_{s(p)}$. We give the vertex $pG_{s(p)}$ the colour e_n . Notice that the tree T_e has its vertices labelled by paths pw extending p on the right. So T_p is in bijection with the tree T_e whose vertices are labelled by $\{w \in \mathcal{G}^* : t(w) = s(e), l_1(w) \neq \bar{e}\}$.

This gives a self-similar colouring of $\{\mathcal{T}_v : v \in \Gamma^0\}$ except that the vertices $\{G_v : v \in \Gamma^0\}$ are uncoloured. We can fix this if we want by adding v as a colour to \mathbf{C} for each $v \in \Gamma^0$. Then the vertex G_v will be thought of as the root of the Bass-Serre tree, of colour v . The production rule can be extended to v , just by listing the colours of the neighbours of G_v , and the starting set \mathbf{S} will be Γ^0 . Alternatively, we will often deal with the problem of the vertices G_v being uncoloured by dropping them from the graphs; since we'll be working with Thompson-type groups, which are permutations of the end space $\partial\mathcal{T}_{\mathbf{C},\mathbf{S}}$, we can drop finitely many vertices and it won't change the group constructed.

Now we claim that $\{G_e : e \in \Gamma^1\}$ forms a self-similar family of groups. We consider the action of G_e on T_e . Let $g \in G_e$, and let g_1e_1w be a vertex of T_e , so that $g_1 \in \Sigma_{e_1}$. As usual, we write $\alpha_e(g)g_1 = g'_1\alpha_{e_1}(h_g)$, for unique $g'_1 \in \Sigma_{e_1}$ and $h_g \in G_{e_1}$. In the fundamental groupoid, we can calculate:

$$g \cdot g_1e_1w = g'_1e_1 \circ (h_g \cdot w),$$

where \circ denotes concatenation. Thus, we see that g sends the subtree below g_1e_1 to the subtree below g'_1e_1 (preserving the colour e_1) and acts on it by the automorphism $h_g \in G_{e_1}$. This is exactly what we need to have a self similar family of groups.

We illustrate this with the following example. Let \mathcal{G} be the graph of groups, whose underlying graph has two vertices and three edges shown in figure 2.2.

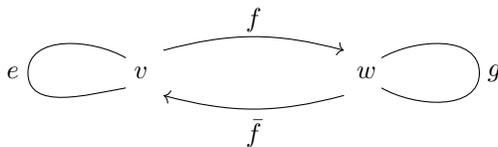


Figure 2.2: An example graph of groups

In our example, both vertex groups will be isomorphic to \mathbb{Z} , with G_v generated by a and G_w generated by b . All edge groups will also be isomorphic to

\mathbb{Z} (with G_e generated by an element e , G_f generated by f , and G_g by g). The embeddings are as follows: $\alpha_f(f) = b^2$; $\alpha_{\bar{f}}(f) = a^2$; $\alpha_e(e) = a$; $\alpha_{\bar{e}}(e) = a^{-1}$; $\alpha_g(g) = b$; $\alpha_{\bar{g}}(g) = b^{-1}$. This gives the Bass-Serre tree shown in Figure 2.3.

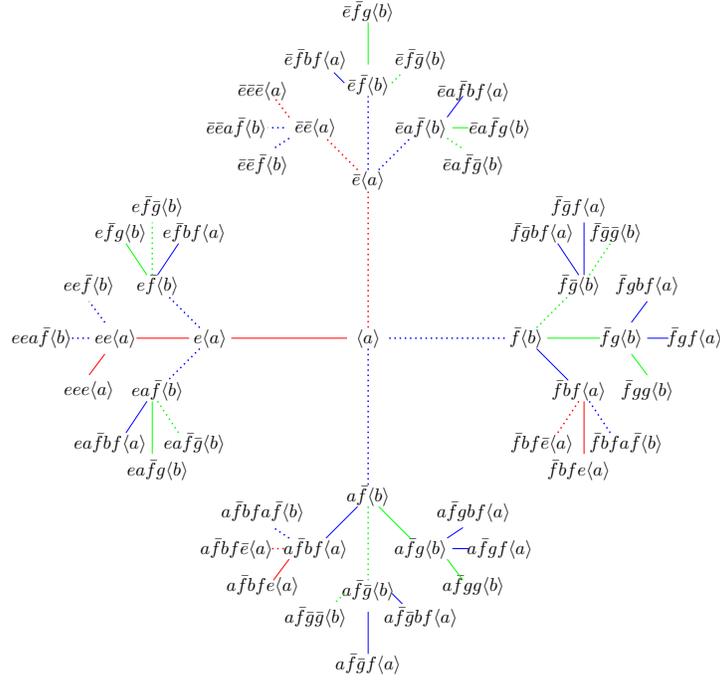


Figure 2.3: The Bass-Serre tree of our example

In this figure, we've coloured e red, f blue and g green, and used dotted lines for the reverse edges. We've drawn the edges coloured, not the vertices, for legibility: you should think of the central vertex as being the root, and then give every other vertex v the colour of the edge with target v . Notice that, for example, the target of each blue edge (oriented outwards) is the source of one red edge, one dotted red edge, and one dotted blue edge. This, and similar facts for each other colour, together show that the colouring is self-similar. We could get a self-similar family of trees using this tree and another Bass-Serre tree, this time rooted at G_w .

We work out one example of the self-similarity of the groups. Consider an edge labelled e (a red edge). Its three children are given by the continuations \bar{f} , $a\bar{f}$ and e - so two of them are blue dotted (from \bar{f}) and the third is red. Consider multiplying these continuations by $a \in \alpha_e(G_e)$. This sends \bar{f} to $a\bar{f}$.

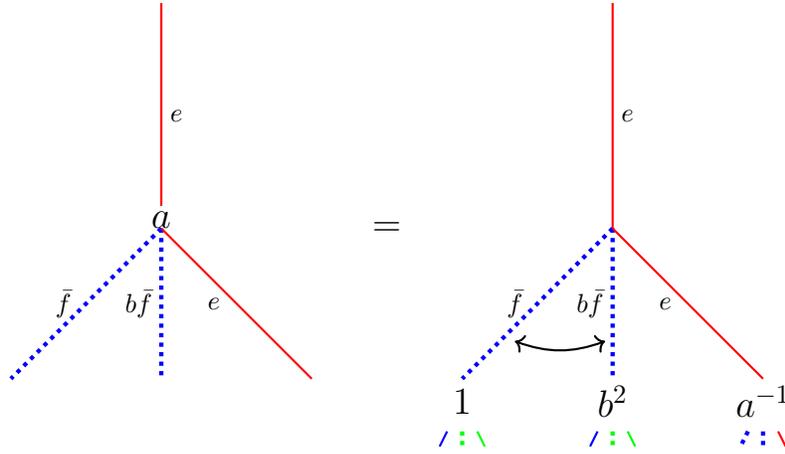


Figure 2.4: The action of $a \in \alpha_e(G_e)$ on an e -labelled subtree

Since $a^2 \bar{f} = \bar{f} b^2$ in the fundamental groupoid, we see that a -multiplication sends $a\bar{f}$ back to \bar{f} , but multiplies by b^2 on the subtree below \bar{f} . Finally $ae = ea^{-1}$ in the fundamental groupoid. This gives the following formula for a acting on the three subtrees beginning $\bar{f}, a\bar{f}, e$:

$$a = (\sigma; 1, b^2, a^{-1}),$$

where σ is the transposition swapping \bar{f} and $a\bar{f}$. Thus we see one part of the family of embeddings that shows the groups G_e form a self-similar family. See figure 2.4, which gives pictures to illustrate the action of a on the subtree. It shows how the action of a swaps over two of the edges below it, and acts below that by the triple $(1, b^2, a^{-1})$, as claimed.

2.4.2 Colour-preserving Nekrashevych-Röver groups

Let \mathbf{C} be a finite set of colours and let \mathbf{S} be a starting set drawn from \mathbf{C} . Let $\mathcal{T}_{\mathbf{C}, \mathbf{S}}$ be a self-similarly coloured family of trees, and let $G_{\mathbf{C}} = \{G_c : c \in \mathbf{C}\}$ be a self-similar family of groups. Whenever e is an edge of $\mathcal{T}_{\mathbf{C}, \mathbf{S}}$ directed away from the root, with $t(e)$ of colour c , then G_c acts on the tree $T_e = \beta_e(T_c)$. So we can define a permutation $\phi(g, e)$ of $\partial \mathcal{T}_{\mathbf{C}, \mathbf{S}}$ for any $g \in G_c$ as follows: $\phi(g, e)$ is the identity outside ∂T_e , and acts on ∂T_e by $\phi(g, e)(x) = \beta_e(g \cdot (\beta_e^{-1}(x)))$. This allows us to make the following definition:

Definition 2.4.2. *The colour-preserving Nekrashevych-Röver group associated*

to the data $\mathbf{C}, \mathbf{S}, G_{\mathbf{C}}$ will be written $V_{\mathbf{C}, \mathbf{S}, G}$ and will be defined as the group of permutations of $\partial\mathcal{T}_{\mathbf{C}, \mathbf{S}}$ generated by the colour-preserving Thompson group $V_{\mathbf{C}, \mathbf{S}}$, and the permutations $\phi(g, e)$ for all $e \in \Gamma^1$ and $g \in G_{c(e)}$.

Proposition 2.4.3. *Every element X of the colour-preserving Nekrashevych-Röver group $V_{\mathbf{C}, \mathbf{S}, G}$ can be written in the following form: let L_1, L_2 be leaf sets of edges of $\mathcal{T}_{\mathbf{C}, \mathbf{S}}$, of size m , with a colour-preserving bijection ψ between them (so that ψ defines an element $\bar{\psi}$ of the colour-preserving Thompson group $V_{\mathbf{C}, \mathbf{S}}$). Let e_1, e_2, \dots, e_m be the edges whose targets lie in L_1 . Then*

$$X = \bar{\psi}\phi(g_1, e_1)\phi(g_2, e_2) \dots \phi(g_m, e_m),$$

for some appropriate choices of L_1, L_2, ϕ and $g_i \in G_{c(e_i)}$.

As a picture, we would draw two finite subtrees with a bijection between their leaves, and a tree automorphism acting below each leaf, similar to Figure 1.16. The proof of this result is technical to write down, but really just comes down to showing that expansion rules exist so that we can multiply these pictures in the usual manner. The expansion rule is similar to what we have seen before.

Proof. It is clear that the generators given for $V_{\Gamma, G_{\mathbf{C}}}$ are all of the required form. We will verify that products and inverses of elements of the given form can also be put in that form, by finding an expansion rule. First of all, we remark that since ψ maps e_i to $\psi(e_i)$, it is easy to see that:

$$X = \bar{\psi}\phi(g_1, e_1)\phi(g_2, e_2) \dots \phi(g_m, e_m) = \phi(g_1, \psi(e_1))\phi(g_2, \psi(e_2)) \dots \phi(g_m, \psi(e_m))\bar{\psi}.$$

In particular,

$$X^{-1} = \bar{\psi}^{-1}\phi(g_m^{-1}, \psi(e_m))\phi(g_{m-1}^{-1}, \psi(e_{m-1})) \dots \phi(g_1, \psi(e_1)),$$

which is of the required form.

Now we show how to do a simple expansion of X . Consider a particular e_i . Let the edges of Γ whose source is $t(e_i)$ be $f_{i,1}, f_{i,2}, \dots, f_{i,d_i}$. Recall that $\phi(g_i, e_i) \in G_{e_i}$, so can be expanded:

$$\phi(g_i, e_i) = (\sigma_i; \phi_{i,1}, \phi_{i,2}, \dots, \phi_{i,k_i}),$$

say. Here σ_i is a colour-preserving permutation of the edges $f_{i,j}$, so $\sigma_i \in V_{\Gamma}$; meanwhile, each $\phi_{i,j}$ is a tree automorphism, of T_{f_j} , which by self-similarity can

be written as $\phi(h_{i,j}, f_{i,j})$ (for $h_{i,j} \in G_c$, appropriate c). Thus we can write

$$\phi(g_i, e_i) = \bar{\sigma}_i \phi(h_{i,1}, f_{i,1}) \dots \phi(h_{i,k_i}, f_{i,k_i}),$$

as a permutation of $\partial\mathcal{T}_{\mathbf{C},\mathbf{S}}$. Moreover, σ_i commutes with all $\phi(g_j, e_j)$ for $j \neq i$, as their support is disjoint. This gives:

$$X = \bar{\psi} \sigma_i \phi(g_1, e_1) \dots \phi(g_{i-1}, e_{i-1}) \phi(h_{i,1}, f_{i,1}) \dots \phi(h_{i,k_i}, f_{i,k_i}) \dots \phi(g_m, e_m).$$

This writes X as a product of an element of V_Γ and various tree automorphisms of the form $\phi(g, e)$. Also, σ_i permutes the simple expansion L'_1 of L_1 at e_i , and $\bar{\psi}$ has a representative whose source leaf set is L'_1 . So $\bar{\psi} \sigma_i$ can be defined by a bijection whose source leaf set is L'_1 . This writes X in the form

$$X = \bar{\psi}' \phi(g'_1, e'_1) \dots \phi(g'_m, e'_m),$$

where $\{e'_1, \dots, e'_m\}$ are the edges whose targets are in L'_1 and $\bar{\psi}'$ is a bijection between leaf sets, with domain leaf set equal to L'_1 . This is the correct form for a simple expansion.

Now suppose we want to form a product XY , where X and Y are given in the form above; we can assume by taking expansions that X is given mapping leaf set L_2 to L_3 , and Y mapping L_1 to L_2 . Then we can write:

$$X = \bar{\psi}_X \phi(g_1, e_1) \phi(g_2, e_2) \dots \phi(g_m, e_m)$$

$$Y = \phi(g'_1, e_1) \phi(g'_2, e_2) \dots \phi(g'_m, e_m) \bar{\psi}_Y,$$

where the edges e_i here have targets in the leaf set L_2 and $\bar{\psi}_X$ maps L_2 to L_3 , $\bar{\psi}_Y$ maps L_1 to L_2 (and we've used the result at the beginning of the proof to put $\bar{\psi}_Y$ on the right). The product XY is then:

$$XY = \bar{\psi}_X \phi(g_1 g'_1, e_1) \phi(g_2 g'_2, e_2) \dots \phi(g_m g'_m, e_m) \bar{\psi}_Y,$$

or alternatively

$$XY = \bar{\psi}_X \bar{\psi}_Y \phi(g_1 g'_1, \bar{\psi}_Y^{-1}(e_1)) \phi(g_2 g'_2, \bar{\psi}_Y^{-1}(e_2)) \dots \phi(g_m g'_m, \bar{\psi}_Y^{-1}(e_m)).$$

This is of the required form, since $\bar{\psi}_X \bar{\psi}_Y$ is a colour-preserving bijection from L_1 to L_3 , whilst the $\bar{\psi}_Y^{-1}(e_i)$ are the edges of the leaf set L_1 . \square

2.4.3 Colour-preserving Nekrashevych-Röver groups from Bass-Serre trees

Finally we show how to produce a colour-preserving Nekrashevych-Röver group from a graph of groups satisfying property (H). This completes the plan in Figure 1.10 and connects the various different objects we have been considering.

Theorem 2.4.4. *Let \mathcal{G} be a locally finite graph of groups whose underlying graph is finite. Let $G_{\mathcal{G}}$ be its path groupoid and let $L_K(\mathcal{G})$ be its Leavitt graph of groups algebra. Let $V_{\mathcal{G}}$ be the set of elements X of $L_K(\mathcal{G})$ of the form:*

$$X = \sum_{i=1}^n S_{\mu_i} U_{s(\mu_i), g_i} S_{\mu'_i}^* = \sum_{i=1}^n 1_{[(\mu'_i, g_i, \mu_i, P), \mu'_i \mathcal{G}^\omega]}$$

where each (μ'_i, g, μ_i, P) is a full element of $S_{\mathcal{G}}$, and additionally X is invertible with

$$X^{-1} = \sum_{i=1}^n S_{\mu'_i} U_{s(\mu_i), g_i^{-1}} S_{\mu_i}^*.$$

We will say X is unitary if X^{-1} is of the given form. Then $V_{\mathcal{G}}$ is a group, and is isomorphic to the topological full group of $G_{\mathcal{G}}$. It is also a colour-preserving Nekrashevych-Röver group, acting on the family of Bass-Serre trees $\{\mathcal{T}_{\mathcal{G}, v}\}_{v \in \Gamma^0}$, where the set of colours is Γ^1 , and for each $\mu \in \mathcal{G}^*$, the vertex $\mu G_{s(\mu)}$ is coloured with the colour of edge $r_1(\mu)$.

In stating this theorem, we're using the result of Proposition 2.3.2 to identify the term $S_{\mu_i} U_{s(\mu_i), g_i} S_{\mu'_i}^*$ of the Leavitt graph-of-groups algebra with the indicator $1_{[(\mu'_i, g_i, \mu_i, P), \mu'_i \mathcal{G}^\omega]}$ in the isomorphic Steinberg algebra. This is useful, because sometimes it will be easier to do algebra calculations, and sometimes it will be easier to think about open bisections acting on \mathcal{G}^ω . We prove the theorem in several steps.

Proof. Step 1: We analyse the structure of unitary elements.

Suppose that

$$X = \sum_{i=1}^n 1_{[(\mu'_i, g_i, \mu_i, P), \mu'_i \mathcal{G}^\omega]}$$

is unitary and of the given form. We write T_i for $1_{[(\mu'_i, g_i, \mu_i, P), \mu'_i \mathcal{G}^\omega]}$, and T_i^* for $1_{[(\mu_i, g_i^{-1}, \mu'_i, P), \mu_i \mathcal{G}^\omega]}$. We will understand X using its action on $K\mathcal{G}^\omega$. Let $\rho \in \mathcal{G}^\omega$. Then the term T_i acts on infinite \mathcal{G} -paths by sending paths $\mu'_i \rho'$ to $\mu_i g \cdot \rho'$ (for $\rho' \in \mathcal{G}^\omega$), and sending other paths to 0. Since $X = \sum T_i$ is unitary,

it is invertible, and so does not vanish on any $\rho \in \mathcal{G}^\omega$. Thus, every infinite \mathcal{G} -path must be an extension of at least one μ'_i on the right.

Now consider the action of some single term $T_i^*T_j$ on $\rho \in X^\omega$. This is one monomial in the product $X^{-1}X$. This monomial sends ρ either to zero or to another end. Moreover, if we consider the action of $T_i^*T_i$ on a \mathcal{G} -path $\rho = \mu'_i\rho'$ on which T_i doesn't vanish, we have:

$$T_i^*T_i\mu'_i\rho' = T_i^*\mu_i g \cdot \rho' = \mu'_i\rho'.$$

Thus $T_i^*T_i$ fixes ρ whenever T_i does not vanish on ρ . So we see $X^{-1}X\rho$ is equal to a sum of paths including $k\rho$, where k is the number of indices i such that $T_i\rho$ does not vanish. We know this has to equal ρ . This is only possible if $k = 1$ for every possible end ρ . This tells us that every infinite \mathcal{G} -path must be an extension of a unique \mathcal{G} -path μ'_i , so that the μ'_i form a leaf set (of the set \mathcal{G}^ω , which is the end space of the family of trees $\{\mathcal{T}_{\mathcal{G},v} : v \in \Gamma^0\}$). Similarly (considering XX^{-1}) the μ_i must also form a leaf set.

Conversely, suppose that

$$X = \sum_{i=1}^n S_{\mu_i} U_{s(\mu_i), g_i} S_{\mu'_i}^*$$

is of the given form, where the μ_i and μ'_j form leaf sets. Then $S_{\mu_i}^* S_{\mu_j}$ is zero unless $i = j$ (since leaf sets are incomparable). So, with T_i and T_i^* as before

$$\sum_{i=1}^n \sum_{j=1}^n T_i^* T_j = \sum_{i=1}^n T_i^* T_i = \sum_{i=1}^n S_{\mu_i} S_{\mu_i}^* = 1,$$

and a similar argument shows also that $XX^* = 1$. Thus X is a unitary.

Step 2: We show that $V_{\mathcal{G}}$ is the topological full group of $G_{\mathcal{G}}$.

Suppose $X \in V_{\mathcal{G}}$. We know X is a sum of terms $T_i = S_{\mu_i} U_{s(\mu_i), g_i} S_{\mu'_i}^*$. Fix i , and let the common rightmost edge of μ_i and μ'_i be e . Looking at T_i as an element of the Steinberg algebra (and using the proof of Theorem 2.3.2), T_i corresponds to 1_{B_i} , where B_i is the open bisection $[(\mu'_i, g, \mu_i, P), \mu_i \mathcal{G}^\omega]$ - where (μ'_i, g, μ_i, P) is a full element of $S_{\mathcal{G}}$. By step 1 of this proof, the domains $\mu_i \mathcal{G}^\omega$ of B_i are disjoint and have union \mathcal{G}^ω . Similarly, their ranges are disjoint open sets with union \mathcal{G}^ω . Thus, the union of the B_i is an open bisection B defined on all of \mathcal{G}^ω , and $\sum 1_{B_i} = 1_B$. This tells us that X is an element of the topological full group corresponding to the open bisection B .

Conversely, let $B \in [[G_{\mathcal{G}}]]$. We know that locally about $x \in \mathcal{G}^\omega$, B restricts to an open bisection $[(\mu'_x, \alpha_{\bar{e}_x}(g_x), \mu_x, P_x), \mu_x \mathcal{G}^\omega]$, where $(\mu'_x, \alpha_{\bar{e}_x}(g_x), \mu_x, P_x)$ is a full element of $S_{\mathcal{G}}$. By the usual compactness argument, finitely many elements of $S_{\mathcal{G}}$ are needed to cover B , which on restricting can be assumed to have disjoint domains (and hence disjoint ranges). Thus B corresponds to a sum of its restrictions B_i which then have the form $[(\mu'_i, \alpha_{\bar{e}}(g), \mu_i, P), \mu_i \mathcal{G}^\omega]$; the inverse of B_i is $[(\mu_i, \alpha_{\bar{e}}(g^{-1}), \mu'_i, P), \mu'_i \mathcal{G}^\omega]$. This gives a unitary X as described in the previous paragraph.

Step 3: We show that this gives a colour-preserving Nekrashevych-Röver group.

We have already seen that the family of Bass-Serre trees is self-similarly coloured. Write it as $\mathcal{T}_{\mathbf{C}, \mathbf{S}}$; first we will argue that the colour-preserving Thompson group $V_{\mathbf{C}, \mathbf{S}}$ is equal to the group of unitaries of $L_K(\mathcal{G})$ of the form

$$X = \sum_{i=1}^n S_{\mu_i} S_{\mu'_i}^*$$

where $r_1(\mu_i) = r_1(\mu'_i)$. Indeed, we've seen that such an element is unitary if and only if the μ_i and μ'_i both form leaf sets. This means that X gives the permutation of \mathcal{G}^ω sending $\mu'_i x$ to $\mu_i x$ (for $x \in \mathcal{G}^\omega$), which means it equals the colour-preserving Thompson element specified by the bijection $\mu'_i \mapsto \mu_i$. The fact that $r_1(\mu_i) = r_1(\mu'_i)$ is precisely what is required for $\mu'_i \mapsto \mu_i$ to be colour-preserving. It's clear that this construction can be reversed to form $X \in L_K(\mathcal{G})$ from an element of $V_{\mathbf{C}, \mathbf{S}}$, establishing the result.

To finish, it just remains to verify that $U_{s(\mu_i), g_i}$ acts by a colour-preserving automorphism of an appropriate subtree. The action is on the tree of (finite or infinite) \mathcal{G} -paths $x = g_1 e_1 g_2 e_2 \dots$ where $l_1(x) \neq 1\bar{e}$ (so that $\mu_i x, \mu'_i x$ are \mathcal{G} -paths). We must verify that g_i sends this tree to itself and preserves colour, under

$$g_1 e_1 g_2 e_2 \dots \mapsto g_i \cdot g_1 e_1 g_2 e_2 \dots$$

The edges never change under the action of a vertex group on \mathcal{T}_v , so the colouring is certainly preserved. Also, since (μ'_i, g, μ_i, P) was assumed full, we have $g \in \alpha_{\bar{e}}(G_e)$. We point out that the group $\alpha_{\bar{e}}(G_e)$ is the stabilizer of the \mathcal{G} -path $1\bar{e}$. Thus, g stabilizes $1\bar{e}\mathcal{G}^\omega$ and so also sends the set $\{x \in s(e)\mathcal{G}^* : l_1(x) \neq 1\bar{e}\}$ to itself. This gives the required action of g on a subtree (which is clearly by tree automorphisms) and we're done. \square

2.5 An example of the groups $V_{\mathcal{G}}$

We conclude this section by describing the group $V_{\mathcal{G}}$ for a particular graph of groups \mathcal{G} . We won't be able to say many group-theoretic properties of the group constructed (we'll prove some finiteness and simplicity results in the next chapter, but questions like the subgroup structure and isomorphism classes of Thompson-style groups are obscure). Hopefully, though, the description will make the abstract group easier to imagine. At the end of this section, we'll illustrate how theorems in the next chapter allow us to characterize all quotients of this $V_{\mathcal{G}}$.

The graph of groups \mathcal{G} will be the graph with one vertex v and one edge e , with $G_v = G_e = \mathbb{Z}$. We will write $G_v = \langle a \rangle$ and $G_e = \langle e \rangle$, and take embeddings $\alpha_e : e \mapsto a^2$ and $\alpha_{\bar{e}} : e \mapsto a^3$. This means that the fundamental group $\pi_1(\mathcal{G}, v)$ has presentation:

$$\pi_1(\mathcal{G}, v) = \langle a, e : a^2e = ea^3 \rangle.$$

This makes $\pi_1(\mathcal{G}, v)$ into a Baumslag-Solitar group $BS(2, 3)$.

The Bass-Serre tree $\mathcal{T}_{\mathcal{G}, v}$ for \mathcal{G} is drawn out in Figure 2.5. We write f for \bar{e} . Each vertex of $\mathcal{T}_{\mathcal{G}, v}$ are shown coloured red or blue, according to whether its label ends in e or f . This is a self-similar colouring: one can see that each red vertex leads to two red vertices and two blue vertices (moving outwards), whereas each blue vertex leads to one red and three blue vertices.

In the Bass-Serre tree, we have marked a small dot in the middle of the edge between \emptyset and e . Considering the tree as starting from that point rather than from \emptyset , the vertex \emptyset looks like a vertex of type f (a blue vertex), since it leads to three blue vertices and one red vertex. We can use this to define colour-preserving Thompson elements more efficiently, by giving \emptyset a colour, and will do this below. We could mark such a dot in the centre of any edge, and it would change the colours of finitely many vertices.

The fundamental group $BS(2, 3)$ acts on this tree in the usual manner, with a fixing the root and e translating along the infinite ray $\dots, ff, f, \emptyset, e, ee, \dots$. We can write the action of e or f as a colour-preserving Thompson group element, shown in Figure 2.6. This figure shows two leaf sets of size 2, which have a unique colour-preserving bijection between them. This Thompson element shows that f maps vertices ew to w and other vertices x to fx , which is the correct formula.

The action of a is not by a Thompson element, but we can study the action of its powers on subtrees to get a self-similar family of groups. That is, we

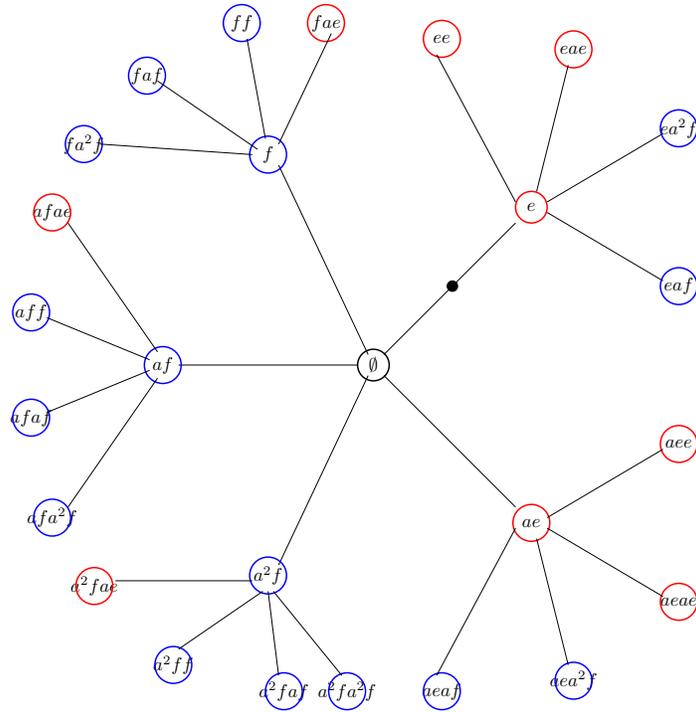


Figure 2.5: The Bass-Serre tree for the Baumslag-Solitar group

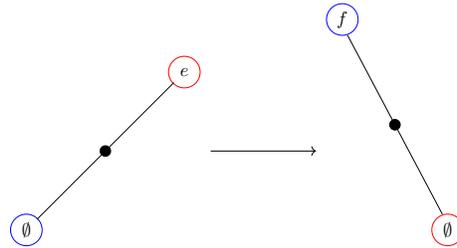


Figure 2.6: Drawing f as a Thompson element

should look at $\alpha_e(G_e) = \langle a^2 \rangle$ acting on the subtree below a blue vertex, and $\alpha_{\bar{e}}(G_e) = \langle a^3 \rangle$ acting on the subtree below a red vertex. The effect of these is drawn out in Figure 2.7.

The figure shows that a^3 acts by swapping the vertices labelled e and ae and fixing the other two; it acts by a^6 on the subtree (originally) below e , by a^3 on the subtree below ae , and so on. Indeed, we observe that in the fundamental group, $a^3 \cdot ae = (a^2)(a^2)e = ea^6$, as shown. Observe that the action below each

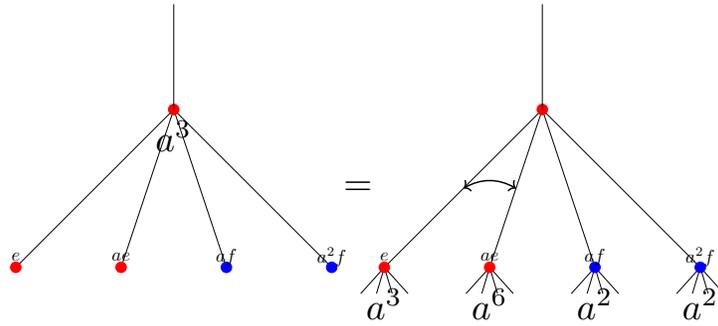


Figure 2.7: The action of a^3 on subtrees

red vertex is in $\langle a^3 \rangle$ and the action below each blue vertex is in $\langle a^2 \rangle$, as we require. Similarly, we also have an action on blue-rooted subtrees in Figure 2.8.

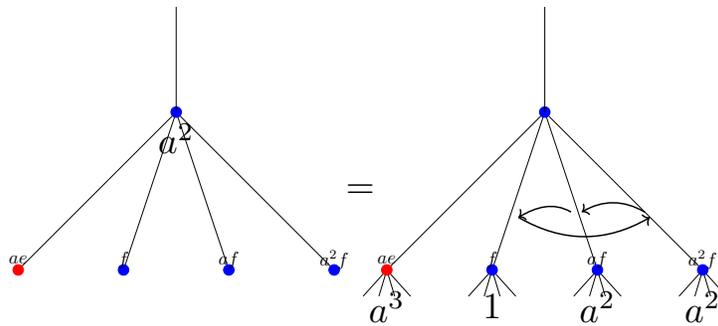
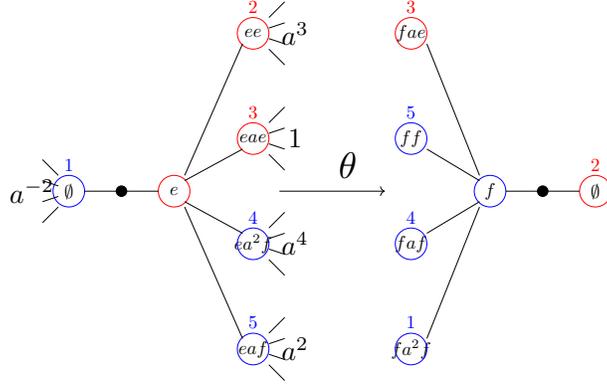


Figure 2.8: The action of a^2 on subtrees

These two figures define the action of the self-similar family of groups $\langle a^2 \rangle$, $\langle a^3 \rangle$. The entire group V_G is generated by Thompson elements with these self-similar groups acting on subtrees. An example element θ is shown below:

The small numbers in this figure show the Thompson bijection of θ , and we illustrate how it permutes the ends by finding the image of $eeaffaeaffae \dots$ under θ . The initial segment ee is mapped to \emptyset by θ (leaf 2 in each leaf set). The edge stabilizer a^3 acts on the remainder $affaeaff \dots$. We can calculate, using the self-similarity:

$$\begin{aligned} a^3 \cdot affae \dots &= af \circ (a^2 \cdot fae \dots) \\ &= afa^2fae \dots \end{aligned}$$



So overall, θ maps $eeaffae \dots$ to $afa^2faeeaff \dots$. Here we were lucky in that the self-similarity calculation terminated. This will not happen in general, but the action will still be well-defined, and we get θ acting as a permutation of \mathcal{G}^ω .

Finally we make some comments about quotients of this $V_{\mathcal{G}}$. In Theorem 3.2.3, we will see that any non-trivial quotient q of $V_{\mathcal{G}}$ is abelian. In particular, the fundamental group $\pi_1(\mathcal{G}, v) \leq V_{\mathcal{G}}$ must factor through an abelian quotient. The abelianization of $\langle a, e : a^2e = ea^3 \rangle$ is \mathbb{Z} , generated by the image of e . In particular, the image of the self-similar family of groups $\langle a^2 \rangle$ and $\langle a^3 \rangle$ is trivial in any quotient of $V_{\mathcal{G}}$, and so the image of $V_{\mathcal{G}}$ under q is the same as the image of the subgroup $V_{\mathcal{C}, \mathbf{s}}$. We will study quotients of colour-preserving Thompson groups in the next chapter; in this case, the abelianization of $V_{\mathcal{C}, \mathbf{s}}$ is $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, as identified in Theorem 3.1.6 (we discuss how to find this quotient explicitly). So any quotient of $V_{\mathcal{G}}$ must in fact be a quotient of $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$.

However, this doesn't tell us that $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ is a possible quotient of $V_{\mathcal{G}}$ - this will depend on how the fundamental group interacts with the Thompson elements. We solve this in Theorem 3.2.6 where we find a presentation for colour-preserving Nekrashevych-Röver groups, making it easy to find the abelianization. The presentation is given in terms of elements $\phi(g, e)$, which correspond to the self-similar family of groups $\langle a^2 \rangle, \langle a^3 \rangle$ acting on subtrees. We have argued that all these $\phi(g, e)$ must become trivial in any quotient q . Looking at Theorem 3.2.6, we just need to make sure that the expansion relation r3 holds in the quotient. This gives us the extra relations that the red transposition and blue 3-cycle of Figures 2.7 and 2.8 are trivial (and no extra relations). We will see that these relations actually already hold in the abelianization of $V_{\mathcal{C}, \mathbf{s}}$. So overall, we have proved that any non-trivial quotient of $V_{\mathcal{G}}$ factors through its

abelianization, which is $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$.

Chapter 3

Properties of the graph of groups constructions

In this section we establish some results about the constructions made in the previous chapter. We will study the colour-preserving Nekrashevych-Röver groups coming from a graph of groups, and see which properties of the graph of groups and of Thompson's group transfer to it. We will also comment on the Leavitt graph-of-groups algebra and discuss how similar it is to usual Leavitt path algebras.

3.1 Elementary results about colour-preserving Thompson groups

First we are going to establish some technical results about colour-preserving Thompson groups, leading to a proof that their derived subgroup is simple, and a characterization of their abelianization. These results were shown by Matui in [39], where he proved that the abelianization G' of a colour-preserving Thompson group G is equal to $(H_0(G) \otimes \mathbb{Z}/2\mathbb{Z}) \oplus H_1(G)$ (see Corollary 6.24 of [39]), giving explicit formulae for the homology groups. We will reprove these results by more combinatorial methods, because we're interested in describing the results in terms of graph-theoretic properties of coloured trees. In particular, we'll be able to describe two reasons why a colour-preserving Thompson group $V_{C,S}$ might have non trivial quotients. One is the parity of a permutation, and

the other was described in the introduction after Figure 1.14, and gives non-trivial homomorphisms to \mathbb{Z} by counting interior vertices. In [39], these are described as the first two homology groups, but it seems the description given here is new.

3.1.1 Quotients of $V_{\mathbf{C},\mathbf{S}}$

Let \mathbf{C} be a finite set of colours and let \mathbf{S} be a starting set drawn from \mathbf{C} . Let p be a production rule on \mathbf{C} . Let $V_{\mathbf{C},\mathbf{S}}$ be the corresponding colour-preserving Thompson group. We will study quotients of $V_{\mathbf{C},\mathbf{S}}$, in particular answering the question of when $V_{\mathbf{C},\mathbf{S}}$ is simple.

Throughout this section, we will make two assumptions on \mathbf{C} . First, we will assume that for each colour $c \in \mathbf{C}$, then every other colour c' appears in one of the tuples $p(c), p(p(c)), p(p(p(c))), \dots$. We will say that \mathbf{C} is *transitive* if this happens. Moreover, we will say that \mathbf{C} is *growing* if $|p(c)| > 1$ for some $c \in \mathbf{C}$.

In the case of a (locally finite non-singular) graph of groups \mathcal{G} where $\mathbf{C} = \mathcal{G}^1$, then \mathbf{C} will be transitive if for any $e, f \in \Gamma^1$ there is a \mathcal{G} -path p with $l_1(p) = e$ and $r_1(p) = f$. It is possible for this to fail. Indeed, suppose that \mathcal{G} is the graph with one vertex and one edge, whose vertex group is $\mathbb{Z} = \langle a \rangle$ and whose edge group is $\mathbb{Z} = \langle e \rangle$, with embeddings $\alpha_e(e) = a^2, \alpha_e(\bar{e}) = a$. Then any \mathcal{G} -path is formed by concatenating the length 1 paths e, ae and \bar{e} , without $e\bar{e}$ or $\bar{e}e$ appearing. In particular, there's no \mathcal{G} -path p with $l_1(p) = e$ and $r_1(p) = \bar{e}$. But this case is unusual, and we will normally find that the set of colours is transitive. For example, if \mathcal{G} is connected and $\alpha_e(G_e)$ is a proper subgroup of $G_{t(e)}$ for all e , then the set of colours is transitive.

The condition that \mathbf{C} is growing is true almost always, and is just there to rule out some degenerate cases where the tree $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ becomes a disjoint collection of rays instead of a branching tree.

The colour count homomorphism: Here we define a homomorphism $\text{CC} : V_{\mathbf{C},\mathbf{S}} \rightarrow \mathbb{Z}^{\mathbf{C}}$. Let $X \in V_{\mathbf{C},\mathbf{S}}$, and suppose that X is represented by a bijection of leaf sets $\mathcal{L}_1 \rightarrow \mathcal{L}_2$. Let $\mathcal{T}_1, \mathcal{T}_2$ be the full subforest of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ with leaves $\mathcal{L}_1, \mathcal{L}_2$ respectively. For each $c \in \mathbf{C}$ and $i = 1, 2$, define $I_i(c, X)$ to be the number of interior vertices (that is, vertices that are not leaves) of \mathcal{T}_i that have the colour c . Finally, we define

$$\text{CC}(X)_c = I_2(c, X) - I_1(c, X).$$

As an example, suppose that $\mathbf{C} = \{A, B\}$ with production rules $p(A) = (A, B, A)$, $p(B) = (B, A, B)$ and that $\mathbf{S} = \{A\}$. We consider again the example from the introduction, Figure 1.14:

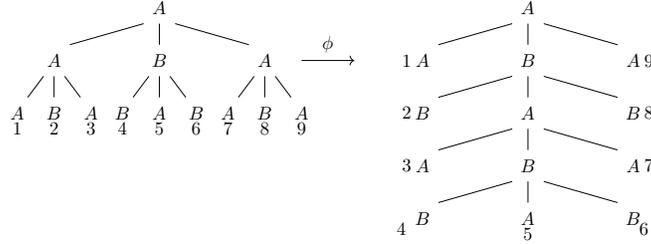


Figure 3.1: An example of non-trivial $\text{CC}(X)$

Here $I_1(A, \phi) = 3$, $I_1(B, \phi) = 1$, $I_2(A, \phi) = 2$ and $I_2(B, \phi) = 2$. So $\text{CC}(\phi) = (-1, 1)$.

In the general case, it is clear that if $\phi : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ is a bijection defining an element $\bar{\phi}$ of $V_{\mathbf{C}, \mathbf{S}}$, then performing a simple expansion of ϕ does not change $\text{CC}(\bar{\phi})$ (a simple expansion at colour c just adds 1 to $I_1(c, \phi)$ and $I_2(c, \phi)$). So CC is well-defined. Moreover if $\psi : \mathcal{L}_2 \rightarrow \mathcal{L}_3$ is another bijection, clearly $\text{CC}(\psi \circ \phi) = \text{CC}(\psi) + \text{CC}(\phi)$. This is enough to show that CC is a homomorphism, which we will call the *colour count homomorphism*. We write $V_{\mathbf{C}, \mathbf{S}}^0$ for the kernel of CC .

Proposition 3.1.1. *Let \mathbf{C} be a transitive growing set of colours and let \mathbf{S} be a starting set. Let $V_{\mathbf{C}, \mathbf{S}}$ be the associated colour-preserving Thompson group, and let $V_{\mathbf{C}, \mathbf{S}}^0$ be as described above. Then $V_{\mathbf{C}, \mathbf{S}}^0$ is the subgroup of $V_{\mathbf{C}, \mathbf{S}}$ generated by its transpositions (those elements defined by a transposition $\phi : \mathcal{L} \rightarrow \mathcal{L}$, for some leaf set \mathcal{L}).*

Proof. If $\phi : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ is a transposition, then $\mathcal{L}_1 = \mathcal{L}_2$ and so $I_1(c, \phi) = I_2(c, \phi)$ for all colours c , and so $\bar{\phi} \in V_{\mathbf{C}, \mathbf{S}}^0$.

Conversely, suppose $\text{CC}(X) = 0$. We seek to write X as a product of transpositions. Suppose that X is defined by a bijection $\mathcal{L}_1 \rightarrow \mathcal{L}_2$. By expanding, we can assume that all leaves of \mathcal{L}_1 have the same depth, so that $\mathcal{L}_1 = \mathbf{S}\mathbf{C}^k$ for some k (the set of all leaves of depth k). If $\mathcal{L}_1 = \mathcal{L}_2$, then X is given by a permutation of \mathcal{L}_1 , which is a product of transpositions since transpositions generate the symmetric group (in each colour). Otherwise, let x be a leaf of minimal depth in \mathcal{T}_2 , which has depth $k' < k$. Then x is an interior vertex of \mathcal{T}_1 but not of \mathcal{T}_2 . Since $I_1(\chi(x), X) = I_2(\chi(x), X)$ (for χ the colour function),

there exists y with $\chi(y) = \chi(x)$ such that y is an interior vertex of \mathcal{T}_2 but not of \mathcal{T}_1 . Notice that x and y must be incomparable (since the vertices above an interior vertex of tree T are also interior vertices), so there exists a transposition $\tau \in V_{\mathbf{C},\mathbf{S}}^0$ swapping x and y . Then replacing X with τX preserves \mathcal{L}_1 and decreases the number of vertices of \mathcal{L}_2 of minimal depth. Repeating this algorithm, we reach the case $\mathcal{L}_1 = \mathcal{L}_2$. \square

An example of this proof is shown below (where $\mathbf{C} = \{A, B\}$, $p(A) = (A, B, B)$, $p(B) = (B, A)$ and $\mathbf{S} = \{A\}$). We don't specify the bijection between leaf sets since it doesn't actually matter for the proof. The vertices x and y are circled.

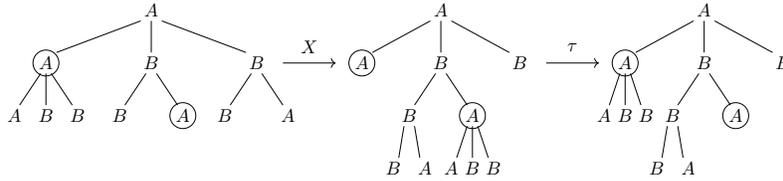


Figure 3.2: An example of simplification with transpositions

We prove a useful corollary.

Corollary 3.1.2. *Let $V = V_{\mathbf{C},\mathbf{S}}$ be a colour-preserving Thompson group and let $X \in V^0$. Suppose that $X = \bar{\phi}$ for some $\phi : L_1 \rightarrow L_2$, and suppose that ϕ has a fixed point x . Then X is a product of transpositions of V which are defined on, and fix, x .*

Proof. Let T be the set of all points in \mathbf{SC}^* that equal or lie below some point of L_1 or L_2 other than x . Then T is equal to \mathbf{SC}^* with $x\mathbf{C}^*$ and finitely many other points removed, and if $v \in T$ then all words vw lie in T also. This means that T is a self-similarly coloured family of trees, with some starting set S_T and the same set of colours \mathbf{C} . Moreover if L is a leaf set for T , then $L \cup \{x\}$ is a leaf set for $\mathcal{T}_{\mathbf{C},\mathbf{S}}$.

Now look at the restriction of X to ∂T , which is an element of $V_{S_T,\mathbf{C}}$, the colour-preserving Thompson group acting on the ends on T . This restriction is clearly still in the kernel of CC , so can be written as a product of transpositions of $V_{S_T,\mathbf{C}}$. If we extend these transpositions to \mathbf{SC}^ω by defining them to be the identity outside ∂T (equivalently, by adding x to the domain and range leaf sets), then they remain transpositions, and they fix x . So we're done. \square

Now we identify the quotient $V_{\mathbf{C},\mathbf{S}}/V_{\mathbf{C},\mathbf{S}}^0$. Define a matrix $M = M_{\mathbf{C},\mathbf{S}}$ as follows: the rows and columns of M are labelled by \mathbf{C} , and for $c, d \in \mathbf{C}$, then $M_{c,d}$ is equal to the number of times d appears in $p(c)$. This M can also be defined as an adjacency matrix of the graph whose vertex set is \mathbf{C} and with edges from c to each colour of $p(c)$, with multiplicity: this is done in [39]. For example, in the pictures of Figure 3.2, we would have:

$$M = \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}.$$

The first column (and first row) represents colour A , and says that it gives an A and two B s on expanding, whilst the second column tells us that B expands to a B and two A s.

Proposition 3.1.3. *Let $V = V_{\mathbf{C},\mathbf{S}}$ be a colour-preserving Thompson group, where \mathbf{C} is a transitive growing set of colours. Then the quotient of V by $V^0 = V_{\mathbf{C},\mathbf{S}}^0$ is isomorphic to $\text{Ker}(M_{\mathbf{C},\mathbf{S}} - \text{Id}_{|\mathbf{C}|})$, when $M_{\mathbf{C},\mathbf{S}}$ is seen as a map from $\mathbb{Z}^{\mathbf{C}}$ to $\mathbb{Z}^{\mathbf{C}}$.*

Proof. We study the image of CC as a subgroup of $\mathbb{Z}^{\mathbf{C}}$, and show that it is $\text{ker}(M_{\mathbf{C},\mathbf{S}} - \text{Id}_{|\mathbf{C}|})$.

Suppose that $X \in V$ is represented by a bijection between leaf sets $\mathcal{L}_1, \mathcal{L}_2$. Notice that \mathcal{L}_1 is formed from \mathbf{S} by repeatedly expanding: expanding at colour c converts a leaf of colour c into an interior vertex, and adds leaves coloured by $p(c)$. Let $s \in \mathbb{Z}^{\mathbf{C}}$ be the vector whose c -coordinate counts how many times the colour c appears in \mathbf{S} . Then the vector of colours in \mathcal{L}_1 can be given by:

$$v = s + \sum_{c \in \mathbf{C}} (M_c - \text{Id}_c) I_1(c, X),$$

where M_c is the column vector of $M_{\mathbf{C},\mathbf{S}}$ in column c , and Id_c is the c th column of the identity matrix. This works because the d -component of $M_c - \text{Id}_c$ counts how many additional leaves of colour d are formed by expanding a leaf of colour c .

Now recall that \mathcal{L}_1 and \mathcal{L}_2 have the same number of leaves of each colour, so we must have:

$$s + \sum_{c \in \mathbf{C}} (M_c - \text{Id}_c) I_1(c, X) = s + \sum_{c \in \mathbf{C}} (M_c - \text{Id}_c) I_2(c, X),$$

so that the $\mathbb{Z}^{\mathbf{C}}$ -vector $(I_1(c, X) - I_2(c, X))_{c \in \mathbf{C}}$ lies in the kernel of $M - \text{Id}$. Moreover, $X \in V^0$ if and only if $I_1(c, X) - I_2(c, X) = 0$ for all c . So we see that the colour count homomorphism defines a map to the abelian group $\text{Ker}(M - \text{Id})$, whose kernel is precisely V^0 . It remains to show that this homomorphism is surjective.

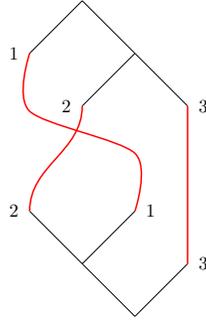
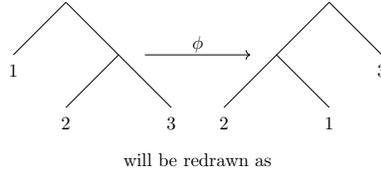
Let $v \in \text{Ker}(M - \text{Id})$. We seek a bijection $\phi : L_1 \rightarrow L_2$ between leaf sets, such that $I_2(c, \bar{\phi}) - I_1(c, \bar{\phi}) = v_c$ for all colours c . Define vectors v^+, v^- by $v_c^+ = v_c$ if $v_c > 0$ and $v_c^+ = 0$ otherwise, and $v_c^- = -v_c$ if $v_c < 0$ and $v_c^- = 0$ otherwise. Then v^+, v^- are nonnegative vectors of disjoint support with $v = v^+ - v^-$. Let N be the largest number appearing in v^+ or v^- , and let \mathcal{L} be a leaf set chosen large enough that each colour appears at least N times (which is definitely possible for \mathbf{C} transitive and expanding). Define $\mathcal{L}_1, \mathcal{L}_2$ by expanding \mathcal{L} according to the vectors v^-, v^+ respectively (that is, if $v_c^+ = k$, we expand k of the colour c leaves in \mathcal{L} when forming \mathcal{L}_2 , and likewise for \mathcal{L}_1). Then let ϕ be any bijection $\mathcal{L}_1 \rightarrow \mathcal{L}_2$ (these exist because $v \in \text{Ker}(M - \text{Id})$.) Then $I_2(c, X) - I_1(c, X) = v_c^+ - v_c^- = v_c$, so $CC(X) = v$ as required. \square

Parity: The second obstruction to $V_{\mathbf{C}, \mathbf{S}}$ being simple is parity, just like in the symmetric groups. Parity is also an invariant for Higman-Thompson groups $V_{n,d}$, when n is odd. If n is even (such as in Thompson's original group V) then a transposition can be expanded to a product of an even number of transpositions, so parity is not defined for $V_{n,d}$.

In our case, we could potentially have more parities to worry about. For a bijection $\phi : L_1 \rightarrow L_2$ between leaf sets, we will define a parity for each pair of colours (counting how many times they are interchanged). Thus we will define a parity function $f_p : V \rightarrow \mathbb{F}_2^N / R_{\mathbf{C}}$, where $N = \frac{1}{2}|\mathbf{C}|(|\mathbf{C}| + 1)$, and $R_{\mathbf{C}}$ is a subspace of relations among the colours \mathbf{C} . In this section we will work with trees drawn in the plane, so we fix a left-to-right ordering on the starting set \mathbf{S} and on the expansions $p(c)$ for each $c \in \mathbf{C}$.

Let $X \in V_{\mathbf{C}, \mathbf{S}}$, and suppose that $\mathcal{T}_1, \mathcal{T}_2$ are full subtrees of $\mathcal{T}_{\mathbf{C}, \mathbf{S}}$ with leaf sets $\mathcal{L}_1, \mathcal{L}_2$, so that X is defined by a bijection $\phi : \mathcal{L}_1 \rightarrow \mathcal{L}_2$. We draw \mathcal{T}_1 as normal, but we draw \mathcal{T}_2 below it, reflected in a horizontal axis, and draw ϕ by connecting appropriate vertices. We colour the connecting lines with the colour of the vertices they join. An example is below, for the usual Thompson's group V .

Given such a diagram for the bijection ϕ and $c, d \in \mathbf{C}$, we define $F_p(\phi)$ to be the \mathbb{F}_2 -vector with $F_p(\phi)_{c,d}$ to be the number of crossings between a line



of colour c and a line of colour d , taken mod 2. We allow $c = d$, so there are $N = \frac{1}{2}|\mathbf{C}|(|\mathbf{C}|+1)$ pairs of colours. We assume, as is usual, ϕ is drawn smoothly in general position, so that no three lines pass through a point, no two lines touch without crossing, no line crosses itself, and with only finitely many crossings. This means that F_p is defined, and we call F_p the *pre-parity function*. As with \mathfrak{S}_n , $F_p(\phi)$ does not depend on how we draw ϕ - in fact, $F_p(\phi)_{c,d}$ is equal to the number of pairs of leaves $\{\ell, \ell'\} \subset \mathcal{L}_1$ whose colours are c and d , and where $\phi(\ell), \phi(\ell')$ are oriented the opposite way to ℓ, ℓ' . Moreover, if $\phi : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ and $\psi : \mathcal{L}_2 \rightarrow \mathcal{L}_3$ are bijections, then it's clear that $F_p(\psi \circ \phi) = F_p(\psi) + F_p(\phi)$. However, F_p may change upon replacing ϕ by an expansion.

We make expansions work by taking a quotient of \mathbb{F}_2^N . First we define \mathbb{F}_2^N another way. Let $A(\mathbf{C})$ be the polynomial \mathbb{F}_2 -algebra $\mathbb{F}_2[X_c : c \in \mathbf{C}]$, with one variable for each colour. Then \mathbb{F}_2^N can be seen as the space $A^{(2)}(\mathbf{C})$ of homogeneous degree two polynomials in $A(\mathbf{C})$. So we will write $F_p(\phi)$ as an element of $A^{(2)}(\mathbf{C})$, where the $X_c X_d$ coefficient is the number of crossings between colour c and colour d , mod 2.

Now take $c \in \mathbf{C}$ and suppose ϕ defines $X \in V$. Consider expanding $\phi : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ at a vertex $\ell_1 \in \mathcal{L}_1$ of colour c . Suppose that the line joining ℓ_1 to $\phi(\ell_1)$ meets a line of colour d . This intersection contributes $X_c X_d$ to $\mathbf{v}_p(\phi)$ (which we're now seeing as a homogeneous quadratic polynomial). Expanding ℓ_1 , we instead get $\sum_{c' \in \mathbf{C}} M_{c,c'} X_{c'} X_d$, where $M_{c,c'}$ as before is equal to the number of

times c' appears in c . So we define $R_{\mathbf{C}}$ to be the set of relations:

$$X_c X_d = \sum_{c' \in \mathbf{C}} M_{c,c'} X_{c'} X_d.$$

The *parity function* $f_p(X)$ can be defined as the image of $F_p(\phi)$ in $A^{(2)}(\mathbf{C})/R_{\mathbf{C}}$, and it is then a homomorphism. It takes values in the degree 2 component of the graded algebra $B(\mathbf{C}) = A(\mathbf{C})/I$, where I is the ideal generated by homogeneous polynomials $X_c - \sum_{c' \in \mathbf{C}} M_{c,c'} X_{c'}$. We say $X \in V$ is *even* if its parity vector is zero.

For example, consider a Higman-Thompson group $V_{n,1}$. This is a colour-preserving Thompson group with one colour, so its preparity function take values in $\langle X_c^2 \rangle$. The set of relations is $X_c^2 = nX_c^2$. If n is even, this is $X_c^2 = 0$, so the parity function is always trivial. If n is odd, though, there are no non-trivial relations, so the parity function takes two values, either 0 or X_c^2 . This is what we've seen happens for Higman-Thompson groups.

We remark that we can simplify the parity function on the subgroup V^0 .

Lemma 3.1.4. *Let $X \in V^0$ where $V = V_{\mathbf{C},\mathbf{S}}$ is a colour-preserving Thompson group such that \mathbf{C} is growing and transitive. Then $f_p(X)$ is contained in the \mathbb{F}_2 -span of the terms $\{X_c^2 : c \in \mathbf{C}\}$. Moreover, there is an isomorphism:*

$$\frac{\langle X_c^2 : c \in \mathbf{C} \rangle}{I \cap \langle X_c^2 : c \in \mathbf{C} \rangle} \cong \frac{\langle X_c : c \in \mathbf{C} \rangle}{I \cap \langle X_c : c \in \mathbf{C} \rangle}.$$

This means that we can define the parity function as an element of $B^{(1)}(\mathbf{C}) = \langle X_c \rangle / I \cap \langle X_c \rangle$.

Proof. First we recall that V^0 is generated by transpositions, and it's clear that if τ is a transposition swapping two leaves of colour c , then $f_p(\tau) = X_c^2$. This gives the first statement.

For the second statement, recall I is defined as the ideal of $A(\mathbf{C})$ generated by polynomials $F_c = X_c - \sum_{c' \in \mathbf{C}} M_{c,c'} X_{c'}$, and write $B(\mathbf{C}) = A(\mathbf{C})/I$ as before. The algebras $A = A(\mathbf{C}), B = B(\mathbf{C})$ and the ideal I are graded, so we write $A^{(n)}, B^{(n)}, I^{(n)}$ for the graded components in degree n . f_p is defined to take values in $A^{(2)}/I^{(2)} = B^{(2)}$, and we're trying to show that a subspace of $B^{(2)}$ is isomorphic to $B^{(1)}$.

Define a group homomorphism ϕ from $B^{(1)}$ to $B^{(2)}$ by $\phi : f \mapsto f^2$ (this is the Frobenius endomorphism). Then ϕ sends $B^{(1)}$ onto $\langle X_c^2 : c \in \mathbf{C} \rangle \subset B^{(2)}$.

It remains to show that ϕ is injective, as then ϕ will give the vector space isomorphism we seek. We show that for $f \in A^{(1)}$, $f^2 \in I$ implies $f \in I$, which implies the result on descending to the quotient. To do this, we change basis. Let $f_{c_1}, f_{c_2}, \dots, f_{c_r}$ be a maximal linearly independent subset of the polynomials $\{f_c : c \in \mathbf{C}\} \subset A^{(1)}$, so that I is generated by $f_{c_1}, f_{c_2}, \dots, f_{c_k}$. Let g_1, g_2, \dots, g_n be a basis of $A^{(1)}$ (where $n = |\mathbf{C}|$) such that $g_i = f_{c_i}$ for $i = 1, 2, \dots, k$. Then observe that

$$A = \mathbb{F}_2[g_1, g_2, \dots, g_n],$$

and

$$I = \mathbb{F}_2[g_1, g_2, \dots, g_k].$$

We need to show that $f^2 \in I$ implies $f \in I$, and in this new basis this is clear. So \mathcal{F} gives an isomorphism between $B^{(1)}$ and the span of X_c^2 in $B^{(2)}$, as required. \square

This lemma tells us that on V^0 , the parity function takes values in

$$B^{(1)} = \frac{\langle X_c : c \in \mathbf{C} \rangle}{X_c = \sum_{c' \in \mathbf{C}} M_{c,c'} X_{c'}} = \text{Coker}(M - \text{Id}),$$

which agrees with a result of [39] (although that paper uses the transpose of M because of differences in definitions).

It is not true in general that Lemma 3.1.4 holds on all of $V_{\mathbf{C},\mathbf{S}}$: there can be elements of $V_{\mathbf{C},\mathbf{S}}$ whose parity function is not in the span of the terms X_c^2 . For example, consider the case where $\mathbf{C} = \{A, B\}$, $\mathbf{S} = (A, B)$, and where $p(A) = (A, B, A, B, A)$, $p(B) = (B, A, A, B, B)$. Consider the element X shown in Figure 3.3.

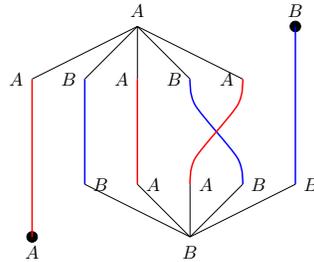


Figure 3.3: An example of why Lemma 3.1.4 does not hold on all of $V_{\mathbf{C},\mathbf{S}}$

This example shows that the parity function of the element pictured is

$X_A X_B$. Moreover, the space $R_{\mathbf{C}}$ of relations is trivial for this set of colours. So this gives an example of a parity function not contained in the span of X_A^2, S_B^2 .

Simplicity: Now we prove a simplicity result. We will prepare by arguing that we can pass to the case where \mathbf{C} is minimal.

Lemma 3.1.5. *Let $V_{\mathbf{C},\mathbf{S}}$ be a colour-preserving Thompson group, and assume that for some $c \in \mathbf{C}$, the tuple $p(c)$ does not contain c . Let $\mathbf{C}_1, \mathbf{S}_1$ be the sets formed by replacing c with $p(c)$ wherever it occurs in \mathbf{S} or in $p(d)$ (for $d \in \mathbf{C}$). Then $V_{\mathbf{C}_1, \mathbf{S}_1}$ is a colour-preserving Thompson group isomorphic to $V_{\mathbf{C}, \mathbf{S}}$. Moreover, $CC(V_{\mathbf{C}_1, \mathbf{S}_1})$ is isomorphic to $CC(V_{\mathbf{C}, \mathbf{S}})$ and $f_p(V_{\mathbf{C}_1, \mathbf{S}_1})$ is isomorphic to $f_p(V_{\mathbf{C}, \mathbf{S}})$, by isomorphisms commuting with the isomorphism between $V_{\mathbf{C}, \mathbf{S}}$ and $V_{\mathbf{C}_1, \mathbf{S}_1}$.*

Proof. We have already seen the isomorphism drawn out in Figure 1.15. Recall that it is defined by first replacing \mathbf{S} with an expansion of \mathbf{S} at each vertex of colour c - which keeps the same space of ends whilst changing \mathbf{S} to \mathbf{S}_1 - and then whenever $\phi : L_1 \rightarrow L_2$ defines an element $\bar{\phi}$ of $V_{\mathbf{C}, \mathbf{S}_1}$, expanding L_1 at each vertex of colour c to get another bijection ϕ_1 defining the same permutation $\bar{\phi}$, but which doesn't involve any vertices of colour c . This lets us drop c from the set of colours \mathbf{C} and get an isomorphic group.

For parity functions, we just observe that if $\phi : L_1 \rightarrow L_2$ defines an element of $V_{\mathbf{C}, \mathbf{S}}$, and colour c does not appear in L_1 , then the preparity function of ϕ never includes the variable X_c . So the parity is the same whether it's calculated with respect to \mathbf{C} or $\mathbf{C} \setminus \{c\}$ which is enough for the result.

Finally, we consider the colour count homomorphism. Suppose that $\phi : L_1 \rightarrow L_2$ defines an element of $V_{\mathbf{C}, \mathbf{S}}$ and that the colour c never appears in L_1 , so that ϕ also defines an element of $V_{\mathbf{C}_1, \mathbf{S}_1}$. Then $CC(\phi)$ is the same when evaluated with respect to the set \mathbf{C} or \mathbf{C}_1 , except that the c -coordinate is dropped for \mathbf{C}_1 . So we get a homomorphism from $CC(V_{\mathbf{C}, \mathbf{S}})$ onto $CC(V_{\mathbf{C}_1, \mathbf{S}_1})$ by ignoring the c -coordinate. It's easy to see that this homomorphism is surjective and commutes with the isomorphism between $V_{\mathbf{C}, \mathbf{S}}$ and $V_{\mathbf{C}_1, \mathbf{S}_1}$. We just need to see that it is injective; if not, then 1_c (the vector that is 1 in the c -coordinate and zero elsewhere) must be in the image of CC on $V_{\mathbf{C}, \mathbf{S}}$, which is $\text{Ker}(M_{\mathbf{C}, \mathbf{S}} - Id_{|\mathbf{C}|})$. This could only happen if $p(c)$ is the singleton $\{c\}$, which does not occur. \square

The point of this lemma is that we can pass to minimal \mathbf{C} without loss of generality. We use this to prove a simplicity result. Let V' be the intersection of the kernels of the homomorphisms CC and f_p .

Theorem 3.1.6. *Suppose that $V = V_{\mathbf{C},\mathbf{S}}$ is a colour-preserving Thompson group where the set \mathbf{C} of colours is growing and transitive. Then V' is a simple group, equal to the derived subgroup of V .*

Proof. During this proof, we will assume that the set \mathbf{C} of colours is minimal: by the previous lemma, we can do this by passing to an isomorphic colour-preserving Thompson group $V_{\mathbf{C}_1,\mathbf{S}_1}$, with

$$V'_{\mathbf{C},\mathbf{S}} = V'_{\mathbf{C}_1,\mathbf{S}_1}.$$

Note also that the construction of a minimal set of colours from \mathbf{C} preserves the fact that \mathbf{C} is growing and transitive. So without loss of generality we can assume that \mathbf{C} is minimal, growing, and transitive. In particular, this implies that there exists $n \in \mathbb{N}$ such that for each $c \in \mathbf{C}$, every colour appears in $p^k(c)$, for all $k > n$. If \mathcal{L} is a leaf set, we will write $\mathfrak{S}(\mathcal{L})$ for the group of colour-preserving permutations of \mathcal{L} . We will write \mathbf{SC}^n for the leaf set whose elements are all vertices of depth n .

Since V' is the kernel of a homomorphism (CC, f_p) from V to an abelian group, it must contain the derived subgroup. Conversely, if we can show that V' is simple, it will equal its own derived subgroup (which is a non-trivial normal subgroup). Since the derived subgroup of V contains the derived subgroup V' of V' , we get that the derived subgroup of V equals V' . We prove simplicity in several stages:

Step 1: generating a layer-preserving permutation:

Let H be a non-trivial normal subgroup of V' , and let $\bar{\sigma} \in H, \bar{\sigma} \neq 1$, where σ is a bijection $\mathcal{L}_1 \rightarrow \mathcal{L}_2$. Our first step will be to use σ to generate a non-trivial element of H which preserves depth of vertices. Assume σ does not have this property, so that $\mathcal{L}_2 \neq \mathcal{L}_1$. We find a particularly nice pair of leaf sets $\mathcal{L}_1, \mathcal{L}_2$ to work in.

We can assume (by expanding) that $\mathcal{L}_1 = \mathbf{SC}^k$ for some k . Let v be some element of \mathcal{L}_1 such that $\sigma(v)$ is of maximal depth, and suppose that $\sigma(v)$ lies below $w \in \mathbf{SC}^k$. Now expand $\mathcal{L}_1 = \mathbf{SC}^k$ to $\mathcal{L}'_1 = \mathbf{SC}^{k'}$ (for some $k' > k$) and expand σ to $\sigma' : \mathcal{L}'_1 \rightarrow \mathcal{L}'_2$ also. For large enough k' , there exist $vm_1, vm_2 \in \mathbf{SC}^{k'}$, both of colour c and lying below v . Let $\ell_1 = \sigma'(vm_1), \ell_2 = \sigma'(vm_2)$. Then ℓ_1, ℓ_2 are of maximal depth in the image of σ' . Moreover, if we write $\ell_1 = v_1\ell'_1, \ell_2 = v_2\ell'_2$ for $v_1, v_2 \in \mathcal{L}'_1$, then v_1, v_2 both lie below $w \in \mathcal{L}_1$. So in particular $\sigma'(v_1)$ and $\sigma'(v_2)$ have the same length, as they are both found by expanding below $\sigma'(w)$.

In summary (and now dropping the primes ' from k' , \mathcal{L}'_i and σ') we have found $\sigma \in H$, defined by a bijection $\sigma : \mathcal{L}_1 \rightarrow \mathcal{L}_2$, such that $\mathcal{L}_1 = \mathbf{SC}^k$. In addition, there are two leaves ℓ_1, ℓ_2 of maximal depth in \mathcal{L}_2 , both of colour c , such that if v_1, v_2 are the leaves of \mathcal{L}_1 above ℓ_1, ℓ_2 , then $\sigma(v_1)$ and $\sigma(v_2)$ have the same length. We can also assume that $|\mathcal{L}_2| \geq 4$. We are now ready to begin.

First suppose that σ fixes some $\ell \in \mathcal{L}_1$. Let τ_1 be a transposition swapping ℓ_1 and ℓ_2 as described. Take τ_2 to be a transposition of two vertices of $\mathcal{T}_{\mathbf{C}, \mathbf{S}}$ that are labelled c , have the same length, and lie below ℓ (this is possible - in a large enough \mathbf{SC}^m , there are two leaves v_1, v_2 of the same colour below ℓ . Since \mathbf{C} is transitive, there exists $x \in \mathbf{C}^*$ such that v_1x has colour c , and then v_2x has colour c also). Observe that τ_2 then commutes with σ .

Now consider τ_1 . Write $\ell_1 = v_1\ell'_1, \ell_2 = v_2\ell'_2$, where v_1, v_2 have depth k (as before), and then ℓ'_1, ℓ'_2 also have the same length. Since \mathcal{L}_1 consists of all depth k vertices, $\sigma(v_1), \sigma(v_2)$ are defined, and $\sigma\tau_1\sigma^{-1} \in V$ is given by the transposition swapping $\sigma(v_1)\ell'_1$ and $\sigma(v_2)\ell'_2$. We check that $\sigma\tau_1\sigma^{-1} \neq \tau_1$. If this were to happen, we would need to have $\sigma(v_1)$ equal to v_1 or v_2 . But this is not possible because ℓ_1, ℓ_2 are in the image of σ , and lie below v_1 and v_2 , so v_1, v_2 cannot be in the image of σ . Thus $\sigma(v_1)\ell'_1$ is not equal to $v_1\ell'_1$ or $v_2\ell'_2$, so $\sigma\tau_1\sigma^{-1} \neq \tau_1$. Moreover, by choice of ℓ_1, ℓ_2 we have that $\sigma(v_1)$ and $\sigma(v_2)$ have the same depth.

Putting this together, define $c = \tau\sigma\tau^{-1}\sigma^{-1} \in H$. Then $c = \tau_1\sigma\tau_1^{-1}\sigma^{-1}$, and we see:

$$c = (\ell_1 \ \ell_2)(\sigma(v_1)\ell'_1 \ \sigma(v_2)\ell'_2).$$

Our study of τ_1 tells us that this is a non-identity, level-preserving element. This is what we wanted to find.

Now for general H , it suffices to find $\sigma \in H$ which fixes some point of $\mathcal{T}_{\mathbf{C}, \mathbf{S}}^\omega$. Let $\rho : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ define a non-identity element of H , and assume that $\mathcal{L}_1 = \mathbf{SC}^k$ as before. After expanding, we can also take k large enough that \mathcal{L}_1 contains five leaves of the same colour, $\ell_1, \ell_2, \ell_3, \ell_4, \ell_5$, so that all ℓ_i and all $\rho(\ell_j)$ are pairwise incomparable. Let τ be the double transposition $(\rho(\ell_1) \ \rho(\ell_2))(\rho(\ell_3) \ \rho(\ell_4))$ (as a permutation of \mathcal{L}_2), so that $\rho^{-1}\tau\rho = (\ell_1 \ \ell_2)(\ell_3 \ \ell_4) \neq \tau$. Then $\tau^{-1}\rho^{-1}\tau\rho \in H$, and this element is not the identity but fixes ℓ_5 , so this generates an element of H with a fixed point.

Step 2: generating all even permutations of some leaf set:

We have produced some $\sigma \in H$ which permutes the set \mathbf{SC}^k of vertices of some particular depth k . The group of all colour-preserving permutations

preserving this layer is $\mathfrak{S}_{k_1} \times \dots \times \mathfrak{S}_{k_n}$ for k_1, \dots, k_n the number of vertices of each colour, and the subgroup V' of even permutations meets it at least at the product $\mathfrak{A}_{k_1} \times \dots \times \mathfrak{A}_{k_n}$ of alternating groups. We first argue that this product of alternating groups is contained in H (for k large enough).

Suppose that $\sigma \in H$ permutes $\mathbf{S}\mathbf{C}^k$, and moves some $\ell \in \mathbf{S}\mathbf{C}^k$ of colour c . Since we assumed that \mathbf{C} was minimal and transitive, every colour of \mathbf{C} appears below ℓ at all sufficiently large depths. So by taking k large enough, after expanding, we can assume that σ moves at least one vertex of each colour. Then the intersection of H with the group $\mathfrak{A}_{k_1} \times \dots \times \mathfrak{A}_{k_n}$ is a normal subgroup that is non-trivial in each coordinate (by choice of σ), so H contains at least $\mathfrak{A}_{k_1} \times \dots \times \mathfrak{A}_{k_n}$, since the only normal subgroups of this product of alternating groups are direct products of either 1 or all of \mathfrak{A}_{k_i} in each coordinate.

Our next task is to show that there exists a leaf set \mathcal{L} such that H contains all permutations of \mathcal{L} whose parity function vanishes. Instead of using the parity function $f_p \in B(\mathbf{C}, \mathbf{S})$, we will use the preparity function F_p . We will use the result of Lemma 3.1.4 and view F_p as an element of $A^{(1)}(\mathbf{C}) = \langle X_c : c \in \mathbf{C} \rangle$. We find a leaf set where it takes all values in $I^{(1)}$. Recall that for \mathcal{L} a leaf set, F_p is a group homomorphism when restricted to $\mathfrak{S}(\mathcal{L})$. Choose $\mathcal{L} = \mathbf{S}\mathbf{C}^k$ for large enough k (it will suffice that every colour appears at least 5 times in \mathcal{L} , and this will happen eventually for \mathbf{C} minimal and growing). We know that $H \cap \mathfrak{S}(\mathcal{L})$ contains all permutations with preparity function $0 \in A^{(1)}$, which are permutations that are even on each colour. We will expand \mathcal{L} to get permutations of different preparity function.

Let $\mathbf{C} = \{c_1, c_2, \dots, c_n\}$. For $c \in \mathbf{C}$, write

$$F_c = X_c - \sum_{c' \in \mathbf{C}} M_{c,c'} X_{c'},$$

so that $I^{(1)}$ is spanned by the various F_c . We inductively find leaf sets $\mathcal{L}^{(0)} = \mathcal{L}, \mathcal{L}^{(1)}, \mathcal{L}^{(2)}, \dots, \mathcal{L}^{(n)}$, such that $H \cap \mathfrak{S}(\mathcal{L}^{(k)})$ contains all permutations σ where $F_p(\sigma) \in \langle F_{c_1}, \dots, F_{c_k} \rangle$. Moreover, we will construct $\mathcal{L}^{(k)}$ from \mathcal{L} by expanding at two leaves of each colour c_1, c_2, \dots, c_k . The base case is $\mathcal{L}^{(0)} = \mathcal{L}$, which we have shown contains all permutations with preparity function 0.

Given $\mathcal{L}^{(k)}$, find four leaves $\ell_{k+1,1}, \ell_{k+1,2}, \ell_{k+1,3}, \ell_{k+1,4}$ of \mathcal{L} with colour c_{k+1} . Because $\mathcal{L}^{(k)}$ is formed from \mathcal{L} by expanding at different colours to c_{k+1} , these will also be leaves of $\mathcal{L}^{(k)}$. Define $\tau_{k+1} \in H$ to be the double transposition $(\ell_{k+1,1} \ell_{k+1,2})(\ell_{k+1,3} \ell_{k+1,4})$ on $\mathcal{L}^{(k)}$, and define $\mathcal{L}^{(k+1)}$ to be the expansion

of $\mathcal{L}^{(k)}$ at leaves $\ell_{k+1,1}, \ell_{k+1,2}$. Observe that $\mathcal{L}^{(k+1)}$ shares at least 3 vertices of each colour with \mathcal{L} , and so $H \cap \mathfrak{S}(\mathcal{L}^{(k+1)})$ contains at least a 3-cycle in each colour, and so contains all permutations of preparity function 0 (by normality). Moreover, if $\tau_l^{(k+1)}$ is the expansion of τ_l to $\mathcal{L}^{(k+1)}$ (for each $l \leq k+1$), then $\tau_l^{(k+1)}$ has preparity function $F_{c_{k+1}}$ - one of its two transpositions is expanded in $\mathcal{L}^{(k+1)}$. Since F_p is a homomorphism, $H \cap \mathfrak{S}(\mathcal{L}^{(k+1)})$ contains all permutations whose preparity function is in the span of $F_{c_1}, \dots, F_{c_{k+1}}$. We are then done inductively.

Step 3: generating all even permutations:

We now have produced a leaf set $\mathcal{L}^{(n)}$ such that H contains all permutations of $\mathcal{L}^{(n)}$ whose parity function vanishes. We now claim the same is true for any sufficiently deep leaf set \mathcal{M} . We will prove that whenever \mathcal{M} is an expansion of $\mathcal{L}^{(n)}$, then $H \cap \mathfrak{S}(\mathcal{M})$ contains all the permutations whose parity function vanishes, and we will do this inductively, by using simple expansions. Notice that $\mathcal{L}^{(n)}$ was formed by expanding a leaf set \mathcal{L} consisting of all vertices of some depth, and that any sufficiently large \mathcal{L} would work. So we can choose \mathcal{L} large enough that in $\mathcal{L}^{(n)}$, every colour appears N times in $\mathcal{L}^{(n)}$, for N to be chosen.

Suppose that \mathcal{M} is a leaf set such that $H \cap \mathfrak{S}(\mathcal{M})$ contains all the permutations of \mathcal{M} whose parity function vanishes, and also that every colour appears at least N times in \mathcal{M} . For example, $\mathcal{M} = \mathcal{L}^{(n)}$ serves as a base case. Let \mathcal{M}' be a simple expansion of \mathcal{M} . Since \mathbf{C} was assumed minimal, every colour appears at least N times in \mathcal{M}' . For every preparity function $f \in I^{(1)}$ that is the preparity function of some permutation, let σ_f be a permutation realizing f (with σ_f in the abstract group $\mathfrak{S}_{k_1} \times \dots \times \mathfrak{S}_{k_n}$, some integers $k_i = k_i(\sigma)$). Choose $N \geq \max(k_i(\sigma_f) + 1)$, the maximum taken over $f \in I^{(1)}$ and $1 \leq i \leq n = |\mathbf{C}|$. Then every colour c_i appears at least k_i times in $\mathcal{M} \cap \mathcal{M}'$ (which differs from \mathcal{M} at a single point). By choice of k_i , every even preparity function is realized by some permutation supported on $\mathcal{M} \cap \mathcal{M}'$, which means every preparity function in I is realized on \mathcal{M}' . Moreover, if we also insist $N \geq 4$, then $\mathcal{M} \cap \mathcal{M}'$ contains at least 3 vertices of each colour, so $H \cap \mathfrak{S}(\mathcal{M}')$ contains at least a 3-cycle on each colour. As usual, this implies it contains all permutations whose preparity function vanishes, and combining this with the previous fact, every permutation of \mathcal{M}' whose preparity function lies in I is contained in the normal subgroup H . This completes the induction.

This shows that all sufficiently deep even permutations are contained in H . Since we can expand any even permutation and it remains even, we see that all even permutations lie in H .

Step 4: generating V^0 :

We have shown that all even permutations are contained in the normal subgroup H . It remains to show that V^0 is generated by these permutations. Let $X = \bar{\phi} \in V^0$ be non-trivial, where $\bar{\phi}$ is represented by a bijection $\phi : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ between leaf sets. We want to write X as a product of permutations whose parity function vanishes. We can assume there exists $v \in \mathcal{L}_1$ such that $v, \phi(v)$ are incomparable (by taking \mathcal{L}_1 sufficiently large, since there exists an open subset U of the space $\partial\mathcal{T}_{\mathbf{C},\mathbf{S}}$ such that $U \cap X(U) = \emptyset$). Let τ be a double transposition of V^0 which interchanges v and $\phi(v)$. Then τX fixes v . Thus (replacing X with τX) we can assume that X has a fixed point v in \mathcal{L}_1 .

Next, we know that X can be written as a product $X = \bar{\tau}_1 \bar{\tau}_2 \dots \bar{\tau}_k$, where each τ_i is a transposition on some leaf set. Moreover we can assume that none of the transpositions $\bar{\tau}_i$ move the set $v\mathbf{C}^\omega$ of ends below v , by Corollary 3.1.2. Consider the element

$$Y = \bar{\tau}_1 \bar{\tau}'_1 \dots \bar{\tau}_k \bar{\tau}'_k,$$

where τ'_i is a transposition of two leaves $\ell_{i,1}, \ell_{i,2}$, with the same colour as the leaves swapped by τ_i , such that all $\ell_{i,j}$ lie below v and are pairwise incomparable. Then each $\tau_i \tau'_i$ is an even permutation, so Y is a product of even permutations, and $Y \in H$. Moreover, all τ'_i commute with all τ_i , so $Y = X \bar{\tau}'_1 \dots \bar{\tau}'_k$. Since Y and X are both even, then $X^{-1}Y = \bar{\tau}'_1 \dots \bar{\tau}'_k$ is even, and since all the $\ell_{i,j}$ are incomparable, it is a permutation, so $X^{-1}Y \in H$. Since $Y, X^{-1}Y \in H$ it follows that $X \in H$, and we're done - the normal subgroup H contains all of V^0 . So V^0 is simple and equal to its own derived subgroup. \square

Finally we prove a result about bijections between leaf sets that we'll find useful later. Define a *partial leaf set* to be a finite set of incomparable vertices of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ that is not a leaf set. It is clear that every partial leaf set can have vertices added to form a leaf set.

Proposition 3.1.7. *Let \mathbf{C} be a transitive growing set of colours. Let ϕ be a colour-preserving bijection between two partial leaf sets L_1, L_2 . Then ϕ can be extended to a colour-preserving bijection between leaf sets.*

We show an example of the statement in Figure 3.4. Two partial leaf sets for the same $\mathbf{C} = \{A, B, C\}, \mathbf{S} = \{A\}$ are shown in red. For Thompson's group V , this result is almost obvious, but the need to get a colour-preserving bijection makes it more difficult in this case.

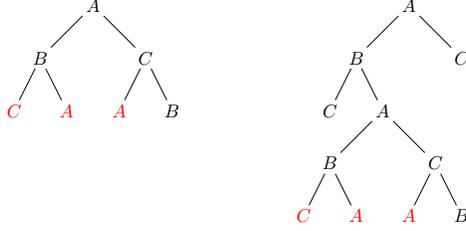


Figure 3.4: An example of two partial leaf sets in bijection

Proof. Let $\mathbf{S} = (s_1, s_2, \dots, s_m)$. There exist m vertices u_1, \dots, u_m , incomparable with L_1 and each other, such that $c(u_i) = c(s_i)$ (since \mathbf{C} is transitive and growing). It is possible to extend $L_1 \cup \{u_1, \dots, u_m\}$ to a leaf set

$$M_1 = \{m_1, \dots, m_k, u_1, \dots, u_m\},$$

where $L_1 \subset \{m_1, \dots, m_k\}$. Similarly extend L_2 to a leaf set $M_2 = \{m'_1, \dots, m'_e, u'_1, \dots, u'_m\}$ where u'_i has the same colour as s_i . Now suppose that

$$l_i = s_j x_{c_1, i_1} \dots x_{c_n, i_n} \in \mathbf{SC}^*.$$

We write $u'_i l_i$ for the string $u'_j x_{c_1, i_1} \dots x_{c_n, i_n}$. Define the set $u' M_1$ to be $\{u' m_i : 1 \leq i \leq m\}$. Informally, $u' M_1$ is formed by hanging the leaves m_i of M_1 off the vertices u'_1, \dots, u'_m . Define $u M_2$ similarly, as $\{u m'_i : 1 \leq i \leq m\}$. Then it is easy to verify that $u' M_1 \cup \{m_i : 1 \leq i \leq k\}$ and $u M_2 \cup \{m'_i : 1 \leq i \leq m\}$ are leaf sets, with a colour-preserving bijection, so we're done. □

An example of this proof is drawn below for the partial leaf sets of Figure 3.4. In Figure 3.5, the leaf sets M_1 and M_2 are shown, with the points m_i in red and the points u_j in blue (in this example, $m = 1$ because \mathbf{S} is a single point). Figure 3.6 shows the construction of the final leaf sets, which you can see are produced by hanging one of the two trees of Figure 3.6 off the vertex u_1 or u'_1 of the other. To make this clearer, one tree has been drawn in red and one in black, with the connecting vertex in blue.

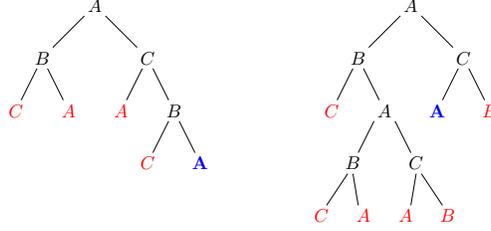


Figure 3.5: The leaf sets M_1 and M_2

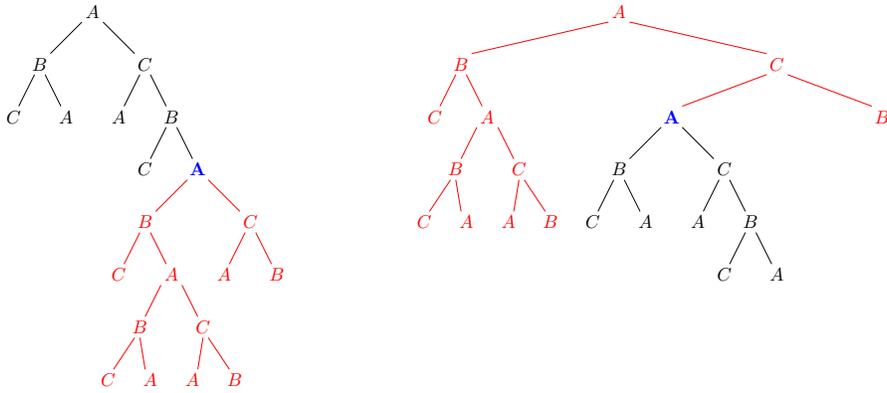


Figure 3.6: Leaf sets with the same colours, extending the partial leaf sets of Figure 3.4

3.2 Results about colour-preserving Nekrashevych-Röver groups

We now use these results to tell us about colour-preserving Nekrashevych-Röver groups. Let \mathbf{C} be a finite set of colours, let \mathbf{S} be a finite starting set drawn from \mathbf{C} , and let $\{G_c : c \in \mathbf{C}\}$ be a self-similar family of groups.

Proposition 3.2.1. *Let \mathcal{G} be a finite graph of groups, and let $V_{\mathcal{G}}$ be its colour-preserving Nekrashevych-Röver group. Then there is a homomorphism ϕ from the fundamental group $\pi_1(\mathcal{G}, v)$ to $V_{\mathcal{G}}$, whose kernel is given by elements of $\pi_1(\mathcal{G}, v)$ that act trivially on $v\mathcal{G}^\omega$.*

We first find a map from the fundamental groupoid $F(\mathcal{G})$ embeds in $L_K(\mathcal{G})$, and then consider $\pi_1(\mathcal{G}, v)$ as a subset of the fundamental groupoid.

First recall the definition of the fundamental groupoid (eg from [13] section 2.4). $F(\mathcal{G})$ has set of objects Γ^0 , and set of morphisms generated by $g \in G_v$

(for each $v \in \Gamma^0$, embedding G_v into the isotropy group of $F(\mathcal{G})$ at v) and $e \in \Gamma^1$. The domains and ranges are as follows: $d(g) = r(g) = v$ for $g \in G_v$, and $d(e) = s(e), r(e) = t(e)$. We take relations $e\bar{e} = 1_{t(e)}, \bar{e}e = 1_{s(e)}$, and also $e\alpha_{\bar{e}}(g) = \alpha_e(g)e$ for each $g \in G_e$. Then $F(\mathcal{G})$ can be identified with the set of reduced \mathcal{G} -words, where the given relations are used to simplify a product of two reduced words.

We also make a useful definition of some elements of $L_K(\mathcal{G})$.

Definition 3.2.2. *Let $\mathcal{G} = (\Gamma^0, \Gamma^1, G, \alpha)$ be a graph of groups. Let $e \in \Gamma^1$. Then we define $T_e \in L_K(\mathcal{G})$ by $T_e = S_e + S_{\bar{e}}^*$. We extend this definition to all \mathcal{G} -paths by defining, for $p = g_1e_1 \dots g_n e_n$ in usual notation,*

$$T_p = U_{t(e_1), g_1} T_{e_1} \dots U_{t(e_n), g_n} T_{e_n}.$$

For intuition, T_e acts on \mathcal{G}^ω either by adding e or removing \bar{e} : in other words, you can think of it as ‘adding e on the left, then cancelling with \bar{e} if possible’. T_p similarly can be thought of as adding p on the left of an infinite path and making whatever simplifications you can. If $s(p) = t(p) = v$ say, then p represents an element of the fundamental group $\pi_1(\mathcal{G}, v)$ and we will use T_p as the image of the element p of the fundamental group in $L_K(\mathcal{G})$.

Now we can return to the proof of Proposition 3.2.1.

Proof. We build a subset F of $L_K(\mathcal{G})$ with a bijection $\phi : F(\mathcal{G}) \rightarrow F$, with the property that $\phi(f_1 f_2) = \phi(f_1)\phi(f_2)$ whenever f_1, f_2 are composable in $F(\mathcal{G})$, and $\phi(f_1)\phi(f_2) = 0$ otherwise.

For $e \in \Gamma^1$, recall $T_e = S_e + S_{\bar{e}}^*$. We define ϕ by $\phi(g) = U_{v,g}$ for $g \in G_v$, and $\phi(e) = T_e$, and we extend to $F(\mathcal{G})$ using the multiplicative property. Notice that $U_{v, \alpha_e(g)} T_e = T_e U_{v, \alpha_{\bar{e}}(g)}$ and $T_e T_{\bar{e}} = (S_e + S_{\bar{e}}^*)(S_{\bar{e}} + S_e^*) = S_e S_{\bar{e}}^* + S_{\bar{e}}^* S_e = U_{t(e), 1}$, so the defining relations of $F(\mathcal{G})$ hold under applying ϕ , and we get a homomorphism. We also verify that $\phi(f_1)\phi(f_2) = 0$ whenever $d(f_1) \neq r(f_2)$ in $F(\mathcal{G})$. To do this, let w be a \mathcal{G} -word, and observe that $\phi(w) = U_{t(w), 1} T_w U_{s(w), 1}$ (this is easy to check on the generators, and so works on all of $F(\mathcal{G})$ multiplicatively).

Finally it’s easy to see that the kernel of ϕ is the set of $p \in \pi_1(\mathcal{G}, v)$ where p acts trivially on \mathcal{G}^ω . This completes the proof. □

3.2.1 Quotients of colour-preserving Nekrashevych-Röver groups

We now strengthen one of Nekrashevych's results about quotients of Nekrashevych-Röver groups, to the colour-preserving case. The basic result is Theorem 9.11 of [40]. We make some preparatory definitions: let \mathbf{C} be a finite set of colours, let \mathbf{S} be a starting set, and let $\{G_c : c \in \mathbf{C}\}$ be a self-similar family of groups. Let $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ be the tree associated to the sets \mathbf{C} and \mathbf{S} , coloured by a function χ . Let $V_{\mathbf{C},\mathbf{S},G}$ be the colour-preserving Nekrashevych-Röver group associated to this data, and let $V_{\mathbf{C},\mathbf{S}}$ be the colour-preserving Thompson group associated to \mathbf{C} and \mathbf{S} . We will borrow Nekrashevych's notation, so if $\{v_i\}_{i=1}^n$ and $\{w_i\}_{i=1}^n$ are leaf sets of the family of trees $\mathcal{T}_{\mathbf{C},\mathbf{S}}$, with $\chi(v_i) = \chi(w_i) = c_i$, say, and $g_i \in G(c_i)$, we will write:

$$X = \begin{pmatrix} v_1 & v_2 & \dots & v_n \\ g_1 & g_2 & \dots & g_n \\ w_1 & w_2 & \dots & w_n \end{pmatrix}$$

for the element of $V_{\mathbf{C},\mathbf{S},G}$ that acts on \mathcal{G}^ω by sending $v_i x$ to $w_i g_i \cdot x$, for $x \in \mathcal{G}^\omega$.

Notice that every element of $V_{\mathbf{C},\mathbf{S},G}$ can be factorized:

$$\begin{pmatrix} v_1 & v_2 & \dots & v_n \\ g_1 & g_2 & \dots & g_n \\ w_1 & w_2 & \dots & w_n \end{pmatrix} = \begin{pmatrix} w_1 & w_2 & \dots & w_n \\ g_1 & g_2 & \dots & g_n \\ w_1 & w_2 & \dots & w_n \end{pmatrix} \cdot \begin{pmatrix} v_1 & v_2 & \dots & v_n \\ 1 & 1 & \dots & 1 \\ w_1 & w_2 & \dots & w_n \end{pmatrix}.$$

This factors a general element of $V_{\mathbf{C},\mathbf{S},G}$ into an element of $V_{\mathbf{C},\mathbf{S}}$ (where all $g_i = 1$) and a tree automorphism (where the two leaf sets are identical). We write $\text{Aut}_{\mathbf{C},\mathbf{S},G}$ for the group of all such tree automorphisms (defined by a table with two identical leaf sets) - this is closed under composition, because the family G is self-similar so expansions exist in $\text{Aut}_{\mathbf{C},\mathbf{S},G}$. Then V_G is a product of its subgroups $V_{\mathbf{C},\mathbf{S}}$ and Aut_G (although these two subgroups need not have trivial intersection). As before, we write $V'_{\mathbf{C},\mathbf{S}}$ for the subgroup of $V_{\mathbf{C},\mathbf{S}}$ generated by its even permutations, which we have seen is a simple normal subgroup of $V_{\mathbf{C},\mathbf{S}}$. Finally, we write $V'_{\mathbf{C},\mathbf{S},G}$ for the subgroup of $V_{\mathbf{C},\mathbf{S},G}$ generated by $V'_{\mathbf{C},\mathbf{S}}$ and Aut_G . It's worth being aware that $V_{\mathbf{C},\mathbf{S},G}$ could include elements of $V_{\mathbf{C},\mathbf{S}} \setminus V'_{\mathbf{C},\mathbf{S}}$, because Aut_G might include some non-even permutations.

Now we state the result, which will be about quotients of $V_{\mathbf{C},\mathbf{S},G}$.

Theorem 3.2.3. *Let \mathbf{C} be a finite set of colours, let \mathbf{S} be a finite starting set drawn from \mathbf{C} , and let $\{G_c : c \in \mathbf{C}\}$ be a self-similar family of groups. Suppose*

that \mathbf{C} is growing and transitive. Then any proper quotient of $V_{\mathbf{C},\mathbf{S},G}$ or $V'_{\mathbf{C},\mathbf{S},G}$ is abelian.

This could be proved from Theorem 4.16 of [39], which says that if G is an essentially principal ample groupoid which is purely infinite and minimal (definitions in [39]), and $[[G]]$ is its topological full group, then any non-trivial subgroup of $[[G]]$ normalized by the derived subgroup $[[G]]'$ contains $[[G]]'$. In particular, any normal subgroup contains $[[G]]'$ so any quotient is abelian. One can construct a groupoid for which $V_{\mathbf{C},\mathbf{S},G}$ is the topological full group and check it has the properties required. We will instead give a more group-theoretic proof, following [40], but with extra care taken in several places because of the colouring. In particular, we've already seen that Proposition 3.1.7 is harder.

Proof. The proof will adopt the argument of [40] and prove the result for both $V_{\mathbf{C},\mathbf{S},G}$ and $V'_{\mathbf{C},\mathbf{S},G}$ by the same argument. We will call whichever group is under consideration V .

Let N be a normal subgroup of V , and let $\phi \in N, \phi \neq 1$ be represented by the table:

$$\phi = \begin{pmatrix} v_1 & v_2 & \dots & v_n \\ g_1 & g_2 & \dots & g_n \\ u_1 & u_2 & \dots & u_n \end{pmatrix}.$$

Let $\rho \in \partial\mathcal{T}_{\mathbf{C},\mathbf{S}}$ such that $\phi(\rho) \neq \rho$, and take U a neighbourhood of ρ such that $\phi(U) \cap U = \emptyset$ (which is possible because ϕ is not the identity, and acts continuously). After refining ϕ and restricting U if necessary, we can assume that U is a cylinder set $v_i\mathbf{C}^\omega$. Then ϕ maps $v_i\mathbf{C}^\omega$ to $u_i\mathbf{C}^\omega$. We will also assume, by restricting further, that $v_i\mathbf{C}^\omega \cup u_i\mathbf{C}^\omega$ is not all of $\mathbf{S}\mathbf{C}^\omega$. Moreover, since \mathbf{C} is transitive, we can restrict still further and assume that v_i and u_i are both of colour c , for any particular $c \in \mathbf{C}$. We write v for v_i , u for u_i and h for g_i .

We now prove a lemma that finds useful elements of N .

Let $r \in \mathbf{S}\mathbf{C}^*$, where r has colour c ; let $f \in V_{\mathbf{C},\{c\},G}$. Then we define a map $\Lambda_r(f) \in V_{\mathbf{C},\mathbf{S},G}$ by the following permutation of $\mathbf{S}\mathbf{C}^\omega$: $\Lambda_r(f)$ fixes all ends $\rho \in \mathbf{S}\mathbf{C}^\omega$ where ρ does not begin with r , and $\Lambda_r(f)$ maps $r\rho$ to $rf(\rho)$. Since the tree $r\mathbf{C}^*$ is isomorphic as a coloured tree to $\mathcal{T}_{\mathbf{C},\{c\}}$, it is easy to verify that $\Lambda_r(f)$ is an element of $V_{\mathbf{C},\mathbf{S},G}$. Intuitively, $\Lambda_r(f)$ is just the element f acting on the subtree below r . We remark that if $f \in V'_{\mathbf{C},\{c\},G}$, then $\Lambda_r(f) \in V'_{\mathbf{C},\{c\},G}$.

Lemma 3.2.4 ([40], Lemma 9.13). *Suppose that $\{r, s\}$ is a partial leaf set where r, s are both of colour c . Suppose that $f \in V_{\mathbf{C},\{c\},G}$ (or $f \in V'_{\mathbf{C},\{c\},G}$, depending*

on what we're using for V). Define ψ to be $\Lambda_s(f)\Lambda_r(f^{-1})$ (so that in particular ψ fixes all ends not beginning in r or s). Then $\psi \in N$.

Proof. It is possible to extend both $\{v, u\}$ and $\{r, s\}$ to leaf sets with a colour preserving bijection between them (by Proposition 3.1.7); moreover, we can assume (by first choosing larger partial leaf sets) that each leaf set contains at least two vertices of each colour not equal to u, v, r, s , and so choose this bijection to be even. So V contains an element p defined by a table:

$$p = \begin{pmatrix} u & v & \dots \\ h^{-1} & 1 & \dots \\ r & s & \dots \end{pmatrix},$$

where h is defined as above, such that ϕ contains a column (v, h, u) .

We study $p\phi p^{-1} \in N$ (where ϕ is as described earlier, as a non-trivial element of N). Notice that for $x \in \mathbf{SC}^\omega$, then

$$p\phi p^{-1}(sx) = p\phi(vx) = p(uh \cdot x) = r \circ (h^{-1}h) \cdot x = rx.$$

Define $q = p\phi p^{-1}$ and consider $\Lambda_r(f)^{-1}q^{-1}\Lambda_r(f)q \in N$. We have that:

$$q^{-1}\Lambda_r(f)q(sx) = q^{-1}\Lambda_r(f)(rx) = q^{-1}r(f \cdot x) = s(f \cdot x).$$

Moreover, if $x \in \mathbf{SC}^\omega$ does not begin with s , then qx does not begin with r , so $\Lambda_r(x)$ fixes qx . So we see that $q^{-1}\Lambda_r(f)q$ sends sx to $s(f \cdot x)$, and fixes all other ends: that is,

$$q^{-1}\Lambda_r(f)q = \Lambda_s(f).$$

So we see that

$$\Lambda_r(f^{-1})\Lambda_s(f) = \Lambda_r(f^{-1})q^{-1}\Lambda_r(f)q \in N,$$

proving the lemma. \square

Returning to Theorem 3.2.3, we see in particular that N contains a non-trivial element $\Lambda_s(f)\Lambda_r(f^{-1})$ of $V'_{\mathbf{C}, \mathbf{S}}$, so contains all of the simple group $V'_{\mathbf{C}, \mathbf{S}}$.

Now let $\pi : V \rightarrow H = V/N$ be the quotient homomorphism. Lemma 3.2.4 tells us that $\pi(\Lambda_r(f)) = \pi(\Lambda_s(f))$ whenever $\{r, s\}$ is an incomplete antichain. It's easy to see that if $r_1, r_2 \in \mathbf{SC}^*$ are any words with $\chi(r_1) = \chi(r_2)$, then there exist s_1, s_2 (chosen sufficiently deep) such that each of $\{r_1, s_1\}, \{s_1, s_2\}, \{s_2, r_2\}$

is an incomplete antichain (except possibly if r_i is the unique element of \mathbf{S} - we will only need the result for sufficiently deep r_i). So

$$\pi(\Lambda_{r_1}(f)) = \pi(\Lambda_{s_1}(f)) = \pi(\Lambda_{s_2}(f)) = \pi(\Lambda_{r_2}(f)),$$

and so $\pi(\Lambda_r(f))$ is constant as r ranges over all non-empty words of a particular colour.

We can now prove the theorem in the special case of elements defined by even tables. Consider an element ψ defined by a table

$$\psi : \begin{pmatrix} v_1 & v_2 & \dots & v_n \\ g_1 & g_2 & \dots & g_n \\ w_1 & w_2 & \dots & w_n \end{pmatrix}.$$

Define

$$\phi : \begin{pmatrix} v_1 & v_2 & \dots & v_n \\ 1 & 1 & \dots & 1 \\ w_1 & w_2 & \dots & w_n \end{pmatrix}.$$

Then $\phi^{-1}\psi$ is an element of $\text{Aut}_{\mathbf{C},\mathbf{S},G}$. Suppose that $\phi \in V'_{\mathbf{C},\mathbf{S}}$. Then $\pi(\phi^{-1}\psi) = \pi(\psi)$, since we know $V'_{\mathbf{C},\mathbf{S}}$ is contained in the kernel of π . But

$$\phi^{-1}\psi = \Lambda_{v_1}(g_1) \dots \Lambda_{v_n}(g_n),$$

and since $\Lambda_r(g)$ and $\Lambda_s(h)$ commute for incomparable r, s , and all $\Lambda_r(g)$ are the same under π as r varies, then all $\pi(\Lambda_r(g))$ commute. This proves that the image of any two elements of V defined by even tables commute. These elements generate the group $V'(G)$ (it is generated by V' , where the tables are even, and $\text{Aut}_{\mathbf{C},\mathbf{S},G}$ where the two leaf sets are equal) so in particular, the image of $V'(G)$ under π is abelian.

Now for the general case. We know that if $V(G) \neq V'(G)$, then $V(G)$ is generated by V and $V'(G)$, and that V/V' is isomorphic to $(\mathbb{Z}/2\mathbb{Z})^a \times \mathbb{Z}^b$, some a, b . Moreover, we know that V' is contained in the kernel of π . We study lifts of generators of V/V' . The torsion part of V/V' is generated by a transposition in each colour. We can choose a transposition τ_c of each colour c of disjoint support, so they all commute, and moreover τ_c commutes with $\Lambda_r(g)$ whenever $r\mathbf{C}^*$ is disjoint from the support of τ_c . Since we know that $\pi(\Lambda_r(g))$ is independent of r , we obtain that $\pi(\tau_c)$ commutes with all of $\pi(V'(G))$. This is enough to show that $\pi(V^0(G))$ is abelian (where V^0 is the subgroup of V

generated by transpositions).

Finally suppose that $\sigma_1, \dots, \sigma_k \in V$ are chosen such that $\text{CC}(\sigma_i)$ generate the quotient $\text{CC}(V)$, where CC is the colour count homomorphism. Let $\mathbf{S} = \{s_1, s_2, \dots, s_m\}$, and for each $1 \leq i \leq k$, let $S_i = \{s_{i,1}, \dots, s_{i,m}\}$ be a partial leaf set where $c(s_{i,j}) = c(s_i)$. We can choose the S_i such that the set of all $s_{i,j}$ is itself a partial leaf set (so all of them are incomparable, and they don't form a leaf set taken together). Define σ'_i to be the permutation of \mathbf{C}^ω defined by $\sigma'_i(s_{i,j}x) = s_{i,j}\sigma_i(x)$ and $\sigma'_i(x) = x$ if x does not lie below $s_{i,j}$ for some j (informally, σ'_i is ' σ_i , but just acting on the forest below the set S_i '). It's easy to see that $\text{CC}(\sigma'_i) = \text{CC}(\sigma_i)$ and that all the σ'_i commute. So $\pi(V(G))$ is generated by $\pi(V^0(G))$ and the $\pi(\sigma'_i)$. Moreover, given the σ'_i , one can choose transpositions in each colour and elements $\Lambda_r(g)$ for all $g \in G$ that commute with all σ'_i , so that σ'_i commutes with elements that under π , generate $\pi(V^0(G))$. This is enough to show that $\pi(V(G))$ is abelian. □

3.2.2 Presentations for colour-preserving Nekrashevych-Röver groups

We now turn our attention to finiteness properties. We will establish that colour-preserving Nekrashevych-Röver groups inherit the property of being finitely generated or finitely presented from the groups G_c . This is sensible to ask, because we know $V_{\mathbf{C},\mathbf{S}}$ is finitely presented (in fact, FP_∞) by results of Matui in [39]. In this section, it will sometimes be easier to think about $X \in V_{\mathbf{C},\mathbf{S}}$ acting on sufficiently deep edges of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ as well as on vertices - the action works the same way, acting on all edges below some domain leaf set defining X . Recall in particular that we write G_e for the image of G_c under $\phi(g, e)$, so that G_e is a group of automorphisms of the tree below edge e .

Proposition 3.2.5. *Let $V_{\mathbf{C},\mathbf{S},G}$ be a colour-preserving Nekrashevych-Röver group. Suppose that \mathbf{C} is transitive and growing, and that each group G_c (for $c \in \mathbf{C}$) is finitely generated. Then $V_{\mathbf{C},\mathbf{S},G}$ is also finitely generated.*

Proof. We take as generators the following:

- A finite set of generators for $V_{\mathbf{C},\mathbf{S}}$.
- For each edge e of depth less than d , the set $\phi(g, e)$, where g ranges over a generating set for $G_{\chi(e)}$. We will choose d later.

Notice that the second set of generators are enough to generate G_e for each sufficiently high e . In this proof, we'll write the edges directed away from the roots of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$. Notice also that if $X \in V_\Gamma$ has an action defined on edge e , mapping e to $X(e)$, then X conjugates G_e to $G_{X(e)}$. So if d is large enough that every edge e of Γ is V_Γ -conjugate to some edge e' of depth less than d , then the generators given also generate all G_e and so generate $V_{\mathbf{C},\mathbf{S},G}$ by definition. We just need to find such a d .

Choose d large enough that for every vertex v of depth less than d (except maybe elements of \mathbf{S}) there exists a second vertex w of depth less than d that is incomparable with v and has the same colour as v . Let e be any edge of X . Then there is some leaf set containing both $t(e)$ and $t(e')$, and so there's an element of V_Γ which interchanges T_e and $T_{e'}$ and fixes the rest of the leaf set. So we're done. \square

This result can be strengthened (by weakening the condition that \mathbf{C} is transitive and growing) but we don't do that now. Instead we go on to presentations.

Proposition 3.2.6. *In the usual notation, $V_{\mathbf{C},\mathbf{S},G}$ has an infinite presentation as follows. The generators are:*

- *A finite set X_1, X_2, \dots of generators for $V_{\mathbf{C},\mathbf{S}}$.*
- *The set $\phi(g, e)$, where g ranges over a generating set for $G_{\chi(e)}$, for each edge e .*

and we take relations

r1 A finite set of relations R_1, R_2, \dots such that $\langle X_1, X_2, \dots \mid R_1, R_2, \dots \rangle$ is a presentation for $V_{\mathbf{C},\mathbf{S}}$.

r2 For each e , a set of relations such that the various $\phi(g, e)$ present G_e .

r3 Whenever $g \in G_c$ satisfies:

$$g = (\sigma_g; g_1, g_2, \dots, g_k),$$

according to the self-similarity, and e is an edge with $\chi(e) = c$, take relation:

$$\phi(g, e) = \sigma_{g,e} \phi(g_1, f_1) \phi(g_2, f_2) \dots \phi(g_k, f_k),$$

where the f_i are the edges whose source is $t(e)$, and where $\sigma_{g,e}$ is the permutation σ_g acting on the edges f_i , as an element of $V_{\mathbf{C},\mathbf{S}}$.

- r4 Commutative relations $\phi(g, e)\phi(h, f) = \phi(h, f)\phi(g, e)$ whenever the edges e, f are incomparable, and $g \in G_{\chi(e)}, h \in G_{\chi(f)}$.
- r5 Conjugation $\sigma\phi(g, e)\sigma^{-1} = \phi(g, \sigma(e))$, whenever $\sigma \in V_{\mathbf{C}, \mathbf{S}}$ can be defined on the edge e . In particular σ commutes with $\phi(g, e)$ whenever σ fixes the edge e .

Notice that all these relations clearly hold in $V_{\mathbf{C}, \mathbf{S}, G}$.

Proof. We study the proof of Proposition 2.4.3 which proved that every element X of $V_{\mathbf{C}, \mathbf{S}, G}$ could be put in a canonical form:

$$X = \bar{\psi}\phi(g_1, e_1) \dots \phi(g_n, e_n).$$

All the generators of $V_{\mathbf{C}, \mathbf{S}, G}$ are in this canonical form, and we showed that you could multiply two elements of this form by expanding so they had a leaf set in common. Moreover, one can check that every relation used in the proof of Proposition 2.4.3 follows easily from (r1)-(r5). So every element of the group generated by this presentation can be written in the canonical form of Proposition 2.4.3, and two such elements are the same when they have a common expansion. Thus elements of the group presented by these relations are in bijection with equivalence classes of canonical form under common expansion, as are elements of $V_{\mathbf{C}, \mathbf{S}, G}$, and the multiplication of canonical forms is the same in both cases. So the two groups are isomorphic. \square

We now improve this presentation to a finite presentation. It uses only the edges of depth at most d , where the depth of an edge is the depth of its source:

Theorem 3.2.7. *Let $V_{\mathbf{C}, \mathbf{S}, G}$ be a colour-preserving Nekrashevych-Röver group, where \mathbf{C} is transitive and growing. Suppose each group G_c for $c \in \mathbf{C}$ is finitely presented. Let $d \in \mathbb{N}$ be large enough (determined in the course of the proof). Then $V_{\mathbf{C}, \mathbf{S}, G}$ has a finite presentation with generators:*

- A finite set X_1, X_2, \dots of generators for $V_{\mathbf{C}, \mathbf{S}}$.
- The set $\phi(g, e)$, where g ranges over a generating set $\gamma(e)$ for $G_{\chi(e)}$, for each edge e of depth less than or equal to d .

and relations:

- R1 A finite set of relations R_1, R_2, \dots such that $\langle X_1, X_2, \dots \mid R_1, R_2, \dots \rangle$ is a presentation for $V_{\mathbf{C}, \mathbf{S}}$.

R2 For each e of depth less than d , a finite set of relations such that the various $\phi(g, e)$ present G_e .

R3 Relations $\phi(g, e)\phi(h, f) = \phi(h, f)\phi(g, e)$, whenever e and f are edges of depth less than d and incomparable, and $g \in \gamma(e), h \in \gamma(f)$.

R4 For each e an edge of $\mathcal{T}_{\mathbf{C}, \mathbf{S}}$ of depth less than d , let $V_{\mathbf{C}, \mathbf{S}}[e]$ be the subgroup of $V_{\mathbf{C}, \mathbf{S}}$ that is supported outside ∂T_e . This is itself a colour-preserving Thompson group (take any leaf set L containing $t(e)$, and remove $t(e)$ to form L' ; $V_{\mathbf{C}, \mathbf{S}}[e]$ is then the colour-preserving Thompson group on the self-similar tree below L'). Under the assumptions that \mathbf{C} is transitive and growing, $V_{\mathbf{C}, \mathbf{S}}[e]$ is then finitely generated, so add in relations to say that each generator of $V_{\mathbf{C}, \mathbf{S}}[e]$ commutes with each generator of G_e .

R5 For all edges e of depth less than d , and for $g \in G_{\chi(e)}$ satisfying $g = (\sigma_g; g_1, g_2, \dots, g_k)$, we will take as relations the expansion formulae:

$$\phi(g, e) = \sigma_{g, e} \phi(g_1, f_1) \phi(g_2, f_2) \dots \phi(g_k, f_k),$$

where the f_i are the edges whose source is $t(e)$, and where $\sigma_{g, e} \in V_{\mathbf{C}, \mathbf{S}}$ is the permutation that σ_g gives on the edges f_i .

R6 Whenever edges e, f are of the same colour and of depth less than d , and there exists $X_{e, f} \in V_{\mathbf{C}, \mathbf{S}}$ such that $X_{e, f}$ is defined on $t(e)$ and $X_{e, f} \cdot t(e) = t(f)$, then choose one particular $X_{e, f}$ with this property, and take relations $X_{e, f} \phi(g, e) X_{e, f}^{-1} = \phi(g, f)$, for each g in a generating set for G_e .

In this presentation, we have carefully taken generators $\phi(g, e)$ where e has depth less than or equal to d , but we have only taken relations for depth less than d . This is done because we need R5 to give an expansion in terms of generators $\phi(g_i, f_i)$, so it must be the case that R5 only is given for levels less than d .

It will turn out that d is large enough if the following hold. First, any two vertices v, w of depth less than d , but greater than zero, and the same colour can be linked by a chain $v_0 = v, v_1, v_2, \dots, v_n = w$ where all v_i are of depth less than d and the same colour, and v_i is incomparable to v_{i+1} . Second, for any vertices v_1, v_2 of the same colour, incomparable, and with depth greater than d , there exist v_3, v_4 of depth less than d , incomparable with v_1, v_2 and each other, and of the same colour as v_1 and v_2 . Both of these properties hold for d large enough.

Proof. Let the group presented by the relations R1-R6 be W . It is easy to check that these relations hold in $V = V_{\mathbf{C},\mathbf{S},G}$. Conversely, we will show that the relations R1-R6 imply r1-r5, which will be enough to give the result. Our strategy will be to conjugate the given relations down the tree to generate the infinite set of relations we've already seen giving a presentation for $V_{\mathbf{C},\mathbf{S},G}$.

Notice in particular that W contains a quotient of $V_{\mathbf{C},\mathbf{S}}$ generated by the X_i . Moreover, since $V_{\mathbf{C},\mathbf{S},G}$ is a quotient of W containing $V_{\mathbf{C},\mathbf{S}}$ (generated by the image of the X_i), it follows that the X_i generate a group isomorphic to $V_{\mathbf{C},\mathbf{S}}$ inside W , which we identify with $V_{\mathbf{C},\mathbf{S}}$. This argument is valid since we know that no non-trivial quotient of $V_{\mathbf{C},\mathbf{S}}$ is isomorphic to $V_{\mathbf{C},\mathbf{S}}$.

First, we claim the following relation family (R6') holds in W : suppose that e, f are the same colour and of depth less than d . Suppose that $Y \in V_{\mathbf{C},\mathbf{S}}$ is defined on e , with $Y \cdot e = f$. Then $Y\phi(g, e)Y^{-1} = \phi(g, f)$. Note that (R6) gives (R6') for one particular choice of Y , namely $Y = X_{e,f}$.

To prove (R6'), we need to show that $X_{e,f}^{-1}Y$ commutes with G_e in W . We study $X_{e,f}^{-1}Y$ as an element of $V_{\mathbf{C},\mathbf{S}}$. By choice of $X_{e,f}$, it can be written with a range leaf set containing $t(f)$, and so $X_{e,f}^{-1}$ can be written with domain leaf set containing $t(f)$ (which maps to $t(e)$ under the bijection defining $X_{e,f}^{-1}$). Similarly, Y has a range leaf set containing $t(f)$ (the image of $t(e)$). So there's a common expansion of the range of Y and the domain of $X_{e,f}^{-1}$ that contains $t(f)$. Composing, we get that $X_{e,f}^{-1}Y$ is represented by a map between two leaf sets that sends $t(e)$ to $t(e)$, so $X_{e,f}^{-1}Y$ lies in $V_{\mathbf{C},\mathbf{S}}[e]$ (since the set of ends below e all lie below $t(e)$, and so anything fixing $t(e)$ lies in $V_{\mathbf{C},\mathbf{S}}[e]$). We are given in (R4) that this commutes with G_e so we're done.

Next, we will define a subgroup G_f of W for every edge f of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$. We are given these groups G_e by R2 for sufficiently high e . Given an edge f , find an edge e of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$ of depth less than d , with $\chi(e) = \chi(f)$ and $t(e)$ and $t(f)$ incomparable. Then there is a leaf set L containing $t(e)$ and $t(f)$. Let σ be the transposition of L interchanging $t(e)$ and $t(f)$. Then we define $\phi(g, f) = \sigma\phi(g, e)\sigma^{-1}$ as elements of W . We claim this doesn't depend on the choice of e . Indeed, suppose e' is another edge of depth less than d , with $\chi(e') = \chi(e)$. Let σ' be the transposition interchanging e' and f (in any leaf set containing both). We require that:

$$\sigma\phi(g, e)\sigma^{-1} = \sigma'\phi(g, e')(\sigma')^{-1}.$$

Now, if e and e' are incomparable, then there's a leaf set containing $t(e), t(f)$ and $t(e')$, on which $\sigma^{-1}\sigma'$ acts as the 3-cycle sending e' to e to f and back

to e . By (R6'), $\sigma^{-1}\sigma'$ conjugates $\phi(g, e')$ to $\phi(g, e)$, as required. In general (for d large enough) we find a sequence $t(e) = v_0, v_1, \dots, v_n = t(f)$, all of the same colour and depth less than d , with v_i and v_{i+1} incomparable. Let e_i be the unique edge directed away from the roots with target v_i , and let σ_i be the transposition exchanging v_i and $t(e)$. By the previous case, we have

$$\sigma_i \phi(g, e_i) \sigma_i^{-1} = \sigma_{i+1} \phi(g, e_{i+1}) \sigma_{i+1}^{-1},$$

and this is enough to establish the result.

Thus there is a well-defined subgroup G_f of W , isomorphic to $G_{\chi(f)}$, for each edge f . We will show that these subgroups satisfy the relations of Proposition 3.2.6.

1. Relation family r1 is the same as R1.
2. The relations in r2 are given by R2 for all e of depth less than d , and given by conjugation for other edges.
3. We do r5 next. Suppose that f is an edge and $\sigma \in V_{\mathbf{C}, \mathbf{S}}$ is defined on f . We split into cases.

Case 1: Suppose that f and $\sigma(f)$ are both of depth at least d . Choose e of the same colour as $f, \sigma(f)$, of depth less than d , and incomparable with both (possible by choice of d). Let τ, τ' be transpositions swapping $t(e)$ with $t(f), t(\sigma(f))$ respectively. Then by definition:

$$\phi(g, f) = \tau \phi(g, e) \tau^{-1},$$

and

$$\phi(g, \sigma(f)) = \tau' \phi(g, e) (\tau')^{-1}.$$

So it's enough to show that $(\tau')^{-1}\sigma\tau$ commutes with $\phi(g, e)$. Analysing $(\tau')^{-1}\sigma\tau$ as an element of V , it sends $t(e)$ to itself (since τ, τ' both map $t(e)$ to $t(f)$, which σ fixes). Thus $(\tau')^{-1}\sigma\tau \in V_{\mathbf{C}, \mathbf{S}}[e]$. Relation R4 tell us that this element commutes with G_e .

Case 2: Suppose that both f and $\sigma(f)$ are of depth less than d . Then this case is R6'.

Case 3: Suppose that f is of depth less than d , and $\sigma(f)$ is of depth at least d (this case will also cover when f has depth at least d , and $\sigma(f)$ has depth less than d , by looking at σ^{-1}). Put $f' = \sigma(f)$. There exists

e of depth less than d , incomparable with f' ; let τ be the transposition swapping $t(e)$ and $t(f')$. We're required to show:

$$\sigma\phi(g, f)\sigma^{-1} = \phi(g, f') = \tau\phi(g, e)\tau^{-1}.$$

In other words, it's enough to show $\tau^{-1}\sigma$ conjugates $\phi(g, f)$ to $\phi(g, e)$. But by the usual argument, $\tau^{-1}\sigma$ is an element of $V_{\mathbf{C}, \mathbf{S}}$ which can be written to send $t(f)$ to $t(e)$, and so we're done by R6'. This completes this part of the proof.

Now we return to r3.

4. The expansions in r3 are given for all sufficiently high edges e by R5. For a general edge e' , conjugate this formula by some transposition σ which swaps e and e' , where e is chosen to be of depth less than d with $\chi(e) = \chi(e')$. The conjugation sends $\sigma_{g,e}$ to $\sigma_{g,e'}$ (because this is true in $V_{\mathbf{C}, \mathbf{S}}$, which we're including a presentation for in R1) and sends $\phi(g, e)$ to $\phi(g, e')$ by definition of the group $G_{e'}$. We need that the conjugation sends $\phi(g_i, f_i)$ to $\phi(g_i, f'_i)$ (where f_i, f'_i are the i th edges whose sources are respectively $t(e)$ and $t(e')$, all edges directed away from the roots), but this holds by property r5 since σ sends f_i to f'_i .
5. For r4, suppose that e_1 and e_2 are incomparable. We wish to show that $\phi(g, e_1)$ and $\phi(h, e_2)$ commute. If e_1 and e_2 are both of depth less than d , this is R3. If e_1 and e_2 are both of depth d or more, then there exist edges e'_1 and e'_2 of depth at most d such that e_1, e_2, e'_1, e'_2 are all incomparable and have the same colour (by choice of d). Conjugating $\phi(g, e_1)$ and $\phi(h, e_2)$ by the double transposition swapping e_1 and e'_1 , and swapping e_2 and e'_2 , we get $\phi(g, e'_1)$ and $\phi(h, e'_2)$, which commute.

The only case remaining is when e_1 is of depth less than d , and e_2 is of depth greater than d (or the reverse). But now expand:

$$\phi(g, e_1) = \sigma_{g, e_1} \phi(g_1, f_1) \dots \phi(g_k, f_k),$$

where f_1, \dots, f_k are the edges below e_1 . Since σ_{g, e_1} only permutes edges below e , it commutes with $\phi(h, e_2)$ by a special case of r5. We will be done if we can show that each $\phi(g_i, f_i)$ commutes with $\phi(h, e_2)$. But f_i and e_2 are incomparable, and repeating the expansion, we eventually expand enough that both edges being compared are of depth d at least.

We've already done this case, so this completes the proof.

□

The conditions asked for on \mathbf{C} are stronger than they need to be, but are applicable in most cases, so we don't try to improve them here. It seems reasonable to expect that similar results hold for other finiteness properties, since Matui showed that colour-preserving Thompson groups are FP_∞ . It seems probable that if all groups G_c for $c \in \mathbf{C}$ are FP_n , then $V_{\mathbf{C},\mathbf{S}}$ inherits this property. Proving this will require some different techniques from those shown here, since higher FP_n properties are defined in terms of cell complexes on which the groups act. We don't go into those constructions here.

3.3 Results on the structure of Leavitt graph of groups algebras

Here we adapt the theorems of [54] to prove analogous uniqueness theorems for Leavitt graph of groups algebras. Let \mathcal{G} be a graph of groups. Let $L_K(\mathcal{G})$ be its Leavitt path algebra, as defined by the presentation in Proposition 2.3.2. Our aim is to produce two theorems showing that an algebra homomorphism $\pi : L_K(\mathcal{G}) \rightarrow A$ is injective provided that π does not vanish on a certain subset S of $L_K(\mathcal{G})$. In one theorem, we will assume that π is a graded homomorphism; in the other, we will not, but will require more conditions on the groups of \mathcal{G} .

We saw in the introduction (Section 1.3.5) that there are uniqueness theorems that work in the general context of Steinberg algebras. The uniqueness theorem we prove here is based on the Steinberg algebra result, but we take care in the case when $G_{\mathcal{G}}$ is non-Hausdorff. The graded uniqueness theorem requires fewer conditions on the graph of groups, but asks for the homomorphism to be graded and not vanish on a larger set. We do this so that it will be true for a wider variety of graphs of groups.

3.3.1 The uniqueness theorems

We begin our uniqueness theorems with the Steinberg uniqueness theorems (of Theorem 1.3.1). We won't be able to apply the result directly, since the groupoid of germs $G_{\mathcal{G}}$ is not Hausdorff. Instead, we work through the proof given in [19] and show that $G_{\mathcal{G}}$ is still close enough to being Hausdorff that the theorems

stand. We first prove that Theorem 1.3.1 still applies in this case and then later consider how to improve or refine it.

Suppose \mathcal{G} is a graph of groups with underlying graph Γ . Let $x \in \mathcal{G}^*$. We will write P_x for $S_x^*S_x$ and Q_x for $S_xS_x^*$ in this section. Then P_x and Q_x are idempotents of $L_K(\mathcal{G})$, which we describe as projections. In terms of the action on infinite \mathcal{G} -paths, P_x fixes paths μ where $x\mu$ is a \mathcal{G} -path, and sends other paths to zero, whilst Q_x fixes paths of the form $x\mu$ (and sends others to zero). That is, Q_x gives projection onto the set $x\mathcal{G}^\omega$.

Now consider the subalgebra of $L_K(\mathcal{G})$ generated by the various projections Q_x . Since the sets $x\mathcal{G}^\omega$ form a basis of clopen sets for \mathcal{G}^ω , it is easy to see that this algebra is isomorphic to $LC_K(\mathcal{G}^\omega)$, the algebra of locally constant K -valued functions on \mathcal{G}^ω . Under this isomorphism, Q_x corresponds to the indicator of $x\mathcal{G}^\omega$. We use this isomorphism to treat $LC_K(\mathcal{G}^\omega)$ as a subalgebra of $L_K(\mathcal{G})$.

We're now ready to discuss the uniqueness theorems.

Theorem 3.3.1 (Uniqueness for Leavitt graph of groups algebras). *Let \mathcal{G} be a locally finite non-singular graph of groups whose underlying graph Γ is finite. Suppose either that the groups of \mathcal{G} are countable or that the groupoid $G_{\mathcal{G}}$ is Hausdorff. Let $L_K(\mathcal{G})$ be its Leavitt graph of groups algebra over a field K . Suppose that $\pi : L_K(\mathcal{G}) \rightarrow A$ is a homomorphism of K -algebras which is injective on $LC_K(\mathcal{G}^\omega)$. Then π is injective.*

First we do the case when $G_{\mathcal{G}}$ is Hausdorff. Notice that, since $G_{\mathcal{G}}$ is a groupoid of germs, then the interior of the isotropy bundle of $G_{\mathcal{G}}$ is just the unit space $G_{\mathcal{G}}^{(0)}$. Indeed, suppose $[s, x] \in G_{\mathcal{G}}$ is a non-identity element of the isotropy group of $G_{\mathcal{G}}$ at x . Then $s \in S_{\mathcal{G}}$ satisfies $s(x) = x$, but any open neighbourhood of x contains a point y with $s(y) \neq x$ (or else the germ of s at x would be trivial). Thus, no neighbourhood $[s, U]$ of $[s, x]$ is contained in the isotropy bundle, as required. The result of this is that Theorem 1.3.1 applies and gives the result in the Hausdorff case.

Now consider the countable case. The proof works through a number of lemmas, as in [19].

Lemma 3.3.2 (cf [19] Lemma 3.2). *Let \mathcal{G} be a locally finite non-singular graph of countable groups with finite underlying graph Γ . Let X be the set of elements of \mathcal{G}^ω whose isotropy group in $G_{\mathcal{G}}$ is trivial. Then X is dense in \mathcal{G}^ω .*

Proof. We use the Baire category theorem. Consider any $s \in S_{\mathcal{G}}$. Write I_s for the set of points $x \in \mathcal{G}^\omega$ such that $s(x) = x$, but the germ of s at x is not the

identity (and so the isotropy group at x contains a non-trivial germ of s). We will show that I_s is a nowhere dense closed set, and so taking the (countable) union of all I_s gives a nowhere dense set, which is the set of all $x \in \mathcal{G}^\omega$ with non-trivial isotropy group.

Since s acts continuously on a subset of \mathcal{G}^ω , I_s is certainly closed (it's an intersection of $\{x : s(x) = x\}$ with the complement of the open set $\{x : s(y) = y \text{ on a neighbourhood of } x\}$). It remains to show that the complement of I_s is dense. Consider any open set $U \subset \mathcal{G}^\omega$; we show that we cannot have $U \subset I_s$. If s is not defined on some point x of U , then x is not contained in I_s . Also, if $s(x) \neq x$ for some $x \in U$, then $x \notin I_s$. The final case is when s is the identity on its restriction to U , but then the germ of s at x is the identity for any $x \in U$, so U does not meet I_s . This completes the proof. \square

Lemma 3.3.3 (cf [19] Lemma 3.3). *Let \mathcal{G} be a locally finite graph of countable groups whose underlying graph is finite. Let $x \in \mathcal{G}^\omega$ be such that the isotropy group of $G_{\mathcal{G}}$ at x is trivial. Let $f \in L_K(\mathcal{G})$ such that f does not vanish on $[1, x] \in G_{\mathcal{G}}$. Then there exists a compact open neighbourhood U of x such that $1_U f 1_U = c 1_{[1, U]}$ for nonzero constant $c \in K$.*

Proof. We reiterate the proof from [19], pointing out that we don't need $G_{\mathcal{G}}$ to be Hausdorff in the proof because it's enough that its unit space is Hausdorff. So we write $f = \sum_{i=1}^N a_i 1_{D_i}$, where each D_i is a compact open bisection and $a_i \in K$ (here we're thinking of $L_K(\mathcal{G})$ as a Steinberg algebra, so it is spanned by indicators of compact open bisections). Our assumption on x means that the only open bisections containing $[1, x]$ are $[1, U]$ for open U containing x . So by splitting the D_i further if necessary, we can assume that $[1, x]$ appears in a unique D_i , which without loss of generality is D_1 . For each i , we choose a compact open neighbourhood V_i of x as follows:

- We have chosen D_1 such that $[1, x] \in D_1$. So it is possible to choose V_1 to be a compact open neighbourhood of x such that $[1, V_1] \subset D_1$, so that $1_{V_1} D_1 1_{V_1} = 1_{V_1}$.
- Suppose that $[s, y] \in D_i$ where either $y = x$ or $s(y) = x$ (but not both, because that is covered by the previous case). Then there is a compact open subset D' of D , containing $[s, y]$, whose domain $d(D')$ and range $r(D')$ are disjoint. Choose V_i to be either $d(D')$ or $r(D')$, whichever one contains x . Then $1_{V_i} D_i 1_{V_i} = \emptyset$.

- If neither the range nor domain of D_i include x , we will choose V_i to be a compact open neighbourhood of x not meeting the range or domain of D_i . Then $1_{V_i}D_i1_{V_i} = \emptyset$.

Now define U to be the intersection of all the V_i . Then we have that

$$1_U f 1_U = 1_U \left(\sum_{i=1}^N a_i 1_{D_i} \right) 1_U = a_1 1_U 1_{D_1} 1_U = a_1 1_{[1,U]},$$

as required. □

Lemma 3.3.4 (cf [18] Lemma 3.1). *Let \mathcal{G} be a locally finite non-singular graph of groups. Let $f \in L_K(\mathcal{G})$ with $f \neq 0$, and suppose that f is in the homogeneous component of degree k . Then there exists a compact open bisection B such that $g = 1_{B^{-1}}f$ is homogeneous of degree 0, and $\text{supp}(g)$ meets $G_{\mathcal{G}}^{(0)}$ non-trivially.*

Proof. We write f as

$$f = \sum_{i=1}^N a_i 1_{D_i},$$

where each $a_i \in K \setminus \{0\}$ and each D_i is a homogeneous compact open bisection. Since f is non-zero and continuous on $G_{\mathcal{G}}$, there must be a basic open set $[g, U] \in G_{\mathcal{G}}$ on which f does not vanish. Choose $x \in U$ with trivial isotropy group. Refining as necessary, we can assume that $[g, x]$ appears in D_1 only (it cannot also appear as $[h, x]$ in some other bisection for $h \neq g$). Putting $B = D_1 = [g, U]$, we have that $1_B^{-1}f$ is certainly homogeneous of degree zero. It remains to show that $1_B^{-1}f$ meets $G_{\mathcal{G}}^{(0)}$ non-trivially. Indeed, we have insisted that $[g, x] \in B$; then $[1, x]$ appears in $B^{-1}B$, and we claim it does not appear in any $B^{-1}D_i$ for $i > 1$. Suppose then that $\beta^{-1}\gamma = \alpha^{-1}\alpha = [1, x]$ for $\beta \in B, \gamma \in D_i$. Since $\beta \in B$, we must have $\beta = [g, y]$. Then $\gamma = \beta[1, x]$, and if γ is to be non-zero we must have $x = y$, and $\gamma = [g, x]$, a contradiction as we assumed $[g, x]$ appeared only in D_1 . Thus overall, $[1, x]$ appears in $1_B^{-1}f$, so in particular $1_B^{-1}f$ does not vanish on $G_{\mathcal{G}}^{(0)}$, and we're done. □

Proof of Theorem 3.3.1: Suppose that π is not injective, and let $f \in \text{Ker } \pi$. Write $f = \sum_{i=1}^N a_i 1_{B_i}$, where the B_i are disjoint homogeneous compact open bisections, and $a_i \in K \setminus \{0\}$. By multiplying by appropriate $1_{B^{-1}}$ (as in Lemma 3.3.4) we can assume that $B_1 = [1, V]$ for some compact open V . By Lemma

3.3.3 we have that $1_U f 1_U \in \text{Ker } \pi$, where $1_U f 1_U = c 1_{[1,U]}$. This is a contradiction, because we've assumed that such elements do not lie in $\text{Ker } \pi$.

Corollary 3.3.5. *Let \mathcal{G} be a graph of groups satisfying the conditions of the Uniqueness Theorem (3.3.1). Let K be a field, and let $\pi : L_K(\mathcal{G})$ be an algebra homomorphism. Suppose that π does not vanish on any P_e , for $e \in \Gamma^1$. Then π is injective. Alternatively, if π does not vanish on any Q_e for $e \in \Gamma^1$, π is again injective.*

Proof. Suppose that π is not injective. Let $I = \text{Ker } \pi$. By the uniqueness theorem, π must vanish on some $1_{[1,U]}$ for an open bisection U . Take $x \in \mathcal{G}^*$ such that $x\mathcal{G}^\omega \subset U$ (and x has length at least 1). Then $1_{[1,x\mathcal{G}^\omega]} \in L_K(\mathcal{G})$ (it can be written as $S_x S_x^*$) and

$$1_{[1,x\mathcal{G}^\omega]} 1_{[1,U]} = 1_{[1,x\mathcal{G}^\omega]},$$

so $S_x S_x^* \in I$. Thus we also have $S_x^* S_x S_x^* S_x = S_x^* S_x \in I$, since I is an ideal. But we have seen before that $S_x^* S_x = S_e^* S_e$, where $e = r_1(x)$ (this is true because the set of paths $\rho \in \mathcal{G}^\omega$ for which $x\rho$ is a \mathcal{G} -path is the same as the set for which $e\rho$ is a \mathcal{G} -path). So we get that $P_e = S_e^* S_e \in I$. Finally, note that $Q_e = S_e S_e^* S_e S_e^*$ must then also lie in I , and we're done. \square

3.3.2 The graded uniqueness theorem

Now we study graded homomorphisms. All the proofs in this section parallel the proofs of [54] for Leavitt path algebras, but with some more care taken around vertex groups. Our theorem will work for any locally finite non-singular graph of groups \mathcal{G} , but we will ask for more conditions on the homomorphism than in the previous theorem.

Proposition 3.3.6. *Let \mathcal{G} be a locally finite nonsingular graph of groups. Suppose that I is a \mathbb{Z} -graded ideal of the Leavitt graph of groups algebra $L_K(\mathcal{G})$ (so that $I = \bigoplus_{n \in \mathbb{Z}} I_n$, where $I_n = I \cap L_K(\mathcal{G})_n$.) Then I is generated by its degree zero subspace I_0 .*

Proof. We have that I is a direct sum of the subspaces I_n , so it's enough to show that each I_n lies in the ideal generated by I_0 . Let $X \in I_n$ for $n > 0$. Then X is a linear combination of monomials $S_\mu U_{s(\mu),g} S_\nu^*$ for \mathcal{G} -paths μ, ν , where

$\ell(\mu) = \ell(\nu) + n$ in each case. Then we can write this as:

$$X = \sum_{p:\ell(p)=n} S_p X_p,$$

a sum over all paths of length p , so that each $X_p \in I_0$ (since $X_p = S_p^* X$). Then X is clearly in the ideal generated by I_0 . The case for $n < 0$ is similar, but we write X as a sum of terms $(X S_q) S_q^*$, where each $X S_q$ is of degree zero. \square

Proposition 3.3.7 (Graded Uniqueness Theorem). *Let \mathcal{G} be a locally finite non-singular graph of groups whose underlying graph is finite. Let $\pi : L_K(\mathcal{G}) \rightarrow A$ be a graded homomorphism of graded K -algebras. Let \mathcal{S} be the set:*

$$\mathcal{S} = \bigcup_{v \in \Gamma^0} K G_v \cup \bigcup_{v \in \Gamma^0} \{P_\mu X P_\nu : X \in K G_v; \mu, \nu \in \mathcal{G}^1\}.$$

Suppose that π does not vanish on any nonzero point of \mathcal{S} . Then π is injective.

In other words, any nontrivial graded ideal of π must meet the set \mathcal{S} non-trivially. We remark that \mathcal{S} is a strict subset of I_0 , so this is a strengthening of the previous proposition. We also remark that if π vanishes on $X \in K G_v$, then it also vanishes on $P_\mu X P_\nu$ whenever P_μ, P_ν are length 1 \mathcal{G} -paths such that $P_\mu X P_\nu$ is non-zero. So the theorem could have been stated without including the group algebras $K G_v$ in the set \mathcal{S} . We've included them in the statement first because they are simpler than terms $P_\mu X P_\nu$, and so easier to check, and second because the two parts of \mathcal{S} appear naturally in the proof.

Proof. Let I be the kernel of π . We need to show that $I = \{0\}$. Since I is a graded ideal it's enough to show that I_0 is zero. Define:

$$\mathcal{S}_n = \langle S_\mu U_{s(\mu),g} S_\nu^* : \ell(\mu) = \ell(\nu) = n \rangle_K,$$

so that $L_K(\mathcal{G})_0$ is the union of the \mathcal{S}_n (the notation \mathcal{S}_n is chosen to suggest that these are sets extending \mathcal{S} , on which π still does not vanish). We will show, inductively on n , that $I \cap \mathcal{S}_n = \{0\}$. First observe that since we have

$$U_{s(\mu),1} = \sum_{r(f)=s(\mu),g \in \Sigma_f} U_{g,s(\mu)} S_f S_f^* U_{g^{-1},s(\mu)},$$

we can write $S_\mu U_{s(\mu),g} S_\nu^* = S_\mu U_{s(\mu),g} U_{s(\mu),1} S_\nu^*$ as a sum of terms $S_{\mu'} U_{s(\mu'),g'} S_{\nu'}^*$, where μ' and ν' are paths that extend μ, ν . This embeds \mathcal{S}_n into \mathcal{S}_{n+1} .

Now we begin the induction. For $n = 0$, we must show $I \cap \mathcal{S}_0 = \{0\}$. If $X \in I \cap \mathcal{S}_0$, then X is a sum of terms $\lambda_g U_{v,g}$ (for K -coefficients λ_g). So for each $v \in \Gamma^0$

$$P_v X = \sum_{g \in G_v, \lambda_g \neq 0} \lambda_g U_{v,g} \in KG_v \subset \mathcal{S}.$$

Thus each $P_v X$ is zero, and so X is zero (since $\sum_{v \in \Gamma^0} P_v = 1$).

Now for the inductive step. Let $X \in \mathcal{S}_{n+1} \cap \ker \pi$, and write:

$$x = \sum_{i=1}^N \lambda_i S_{\mu_i} U_{\nu_i, g_i} S_{\nu_i}^*,$$

where each $\lambda_i \in K$, and μ_i, ν_i are \mathcal{G} -paths of length $n + 1$. Recall that (since Γ is finite) we also have

$$1 = \sum_{e \in \Gamma^1, h \in \Sigma_e} S_{he} S_{he}^*,$$

so that $x = \sum S_{he} S_{he}^* x = \sum x S_{he} S_{he}^*$. This means that there exist length 1 \mathcal{G} -paths he and kf such that $S_{he}^* x S_{kf}$ is non-zero. But then if $n + 1 = \ell(\mu) > 1$, observe that $S_{he}^* S_\mu = S_{\mu'}$ if $\mu = he\mu'$, and $S_{he}^* S_\mu$ is zero otherwise. Similarly, $S_{\nu'}^* S_{kf} = S_{\nu'}$ if $\nu = kf\nu'$ and is zero otherwise. So overall, $S_{he}^* x S_{kf}$ is a sum of terms $S_{\mu'} U_{v,g} S_{\nu'}^* \in \mathcal{S}_n \cap I$, which inductively is zero - a contradiction. If $n = 1$,

$$S_{he}^* x S_{kf} = S_{he}^* S_{he} y S_{kf}^* S_{kf} = P_{he} y P_{kf}$$

for $y \in KG_v$, where $v = s(e) = s(f)$. Then $P_{he} y P_{kf} \in \mathcal{S}$ also. \square

One can use this theorem to show that the natural embedding $U_{v,g} \mapsto U_{v,g}, S_e \mapsto S_e$ defines an injection from $L_K(\mathcal{G})$ into the C^* -algebra $C^*(\mathcal{G})$; we don't give the details. Alternatively, one could try to extend the general (non-graded) uniqueness theorem to more graphs of groups. I proved such a result by generalizing the proof of a Leavitt uniqueness theorem, but it required adding conditions to imply that $G_{\mathcal{G}}$ being Hausdorff. Since we got a stronger result than this by adapting the Steinberg uniqueness theorem, we don't repeat this theorem here.

Chapter 4

Representations and Hecke algebras for Thompson groups

In this chapter we move to think about representations of Thompson's group V and its variants. First of all we discuss some of the results known about representation theory of V , and discuss what happens for colour-preserving Thompson variants. One might hope that some of the well-understood theory of representations of \mathfrak{S}_n transfers to V . It seems that this is too much to hope for, and there is no general classification of all representations of V . Instead, most people working on representing V either just find families of representations where they can, or show that representations of other kinds do not exist. All the representations considered will be unitary representations on Hilbert spaces (that is, we look for homomorphisms $\rho : V \rightarrow \mathcal{U}(\mathcal{H})$, the group of unitary operators on Hilbert space \mathcal{H}), and we put the discrete topology on V .

In the second part of this chapter, we try to generalize the theory of (Iwahori)-Hecke algebras, which are formed by deforming a presentation of the group algebra of \mathfrak{S}_n , to give an algebra that can be thought of as a deformation of the group algebra of V . We will try to produce an algebra which deforms V in its action on the infinite binary tree, by taking a quotient of the group algebra of a braided Thompson group. We will see why this exercise cannot work perfectly, but we will push the analogy between V and an infinite version of \mathfrak{S}_n as far as we can. This will result in an interesting representation of a braided Thompson

group, which locally behaves like a finite Iwahori-Hecke algebra. But we will see that it cannot have the global properties we'd ideally want from a Hecke algebra.

4.1 Representations of Thompson's group V and its variants

There are three families of representations we will consider for V : quasi-regular representations, Koopman representations, and representations of finite factor type. We will state some known results about representations of V and generalize them to colour-preserving Thompson or Nekrashevych-Röver groups.

4.1.1 Quasi-regular representations

The first and easiest examples of representations of V will be the quasi-regular representations. These are the natural generalization of permutation representations of finite groups. In general, let G be a group acting transitively on a set X . Fix $x \in X$, and let P be the subgroup of G equal to the stabilizer of x (we call P a parabolic subgroup). Then X is in bijection with the set G/P of left cosets of P . The quasi-regular representation $\rho_{G/P}$ is defined on the \mathbb{C} -vector space $\ell^2(X)$, via the formula:

$$\rho_{G/P}(g) \cdot f(y) = f(g^{-1} \cdot y),$$

for $g \in G, y \in X$ and $f \in \ell^2(X)$. One can easily verify that $\rho_{G/P}$ is indeed a representation, and equals the induced representation, $\rho_{G/P} = \text{Ind}_P^G 1_P$.

If X is a finite set (of size greater than 1), then $\rho_{G/P}$ cannot be irreducible, because the constant functions on X form an invariant subspace. In fact, for G and X finite, we get that $\rho_{G/P} = 1_G \oplus \rho'$, where ρ' is irreducible if and only if the action of G on X is 2-transitive. For infinite X , the constant functions are not ℓ^2 -integrable, so it is possible for $\rho_{G/P}$ to be irreducible. The relevant condition is given by the following theorem ([38], given in this form in [11]).

Theorem 4.1.1 (Mackey). *Let G be a discrete group and let H be a subgroup of G . Define the commensurator of H , $\text{Comm}_G(H)$, to be:*

$$\text{Comm}_G(H) = \{g \in G : H \cap gHg^{-1} \text{ has finite index in } H \text{ and } gHg^{-1}\}.$$

1. Suppose that $H = \text{Comm}_G(H)$. Let σ be a finite-dimensional irreducible unitary representation of H . Then $\text{Ind}_H^G \sigma$ is an irreducible representation of G .
2. Suppose that H, K are two different subgroups of G , each equal to their own commensurator. Let σ_1, σ_2 be irreducible unitary representations of H, K respectively. Then $\text{Ind}_H^G \sigma_1$ and $\text{Ind}_K^G \sigma_2$ are equivalent representations of G if and only if there exists $g \in G$ such that $K = gHg^{-1}$ and σ_2 is equivalent to the conjugate of σ_1 by g .

In particular, quasi-regular representations are irreducible whenever the parabolic subgroup P is its own commensurator. In the case that $P = \text{Stab}(x)$, then $P \cap gPg^{-1} = \text{Stab}(x) \cap \text{Stab}(g \cdot x)$, which has infinite index in P so long as the P -orbit of $g \cdot x$ is infinite. So $\rho_{G/P}$ is irreducible so long as all P -orbits on X are infinite, apart from the one orbit $\{x\}$. This is applied to branch groups in [9]. We apply it to colour-preserving Thompson groups.

Theorem 4.1.2. *Let $V_{\mathbf{C}, \mathbf{S}}$ be a colour-preserving Thompson group with transitive growing set of colours. Let G be a group of permutations of $\mathbf{S}\mathbf{C}^\omega$ containing $V_{\mathbf{C}, \mathbf{S}}$ (for example, a colour-preserving Nekrashevych-Röver group). Let P be the stabilizer of some point of $\mathbf{S}\mathbf{C}^\omega$. Then the quasi-regular representation $\rho_{G/P}$ is irreducible. If G is countable, then there are uncountably many equivalence classes of such representation.*

Proof. Let $P = \text{Stab}(x)$, and let y be a point of $\mathbf{S}\mathbf{C}^\omega$ distinct from x . We must show that the orbit of y under P is infinite. Let $x = vx'$ and $y = wy'$, where $v, w \in \mathbf{S}\mathbf{C}^*$ form a partial leaf set. Since \mathbf{C} is transitive and growing, we can inductively choose a family $w_0 = w, w_1, w_2, \dots$ of elements of $\mathbf{S}\mathbf{C}^*$ such that w_i and w have the same colour, and v, w_0, w_1, \dots, w_k form a partial leaf set. Extending each set v, w_0, \dots, w_k to a leaf set \mathcal{L} , we see that there exists a permutation of \mathcal{L} fixing v and interchanging w_0 and w_k , and hence there exists an element of G fixing x but interchanging y and $w_k y'$. By choice of w_k , all the $w_k y'$ are distinct, and so the P -orbit of y is infinite, as required.

Finally, suppose that G is countable. Then since there are uncountably many ends in $\mathbf{S}\mathbf{C}^\omega$ but all G -orbits are countable, there are uncountably many conjugacy classes of parabolic subgroups P . By the second part of Mackey's theorem, each class of parabolic subgroups gives rise to an inequivalent irreducible representation. \square

4.1.2 Representations of finite factor type

Here we briefly survey the results of [27] and [28] on other ways to find representations of Thompson's group V and its variants. First we describe factor representations of finite type. The definition of this family of representations is given in [28]. In brief, representations of finite factor type are the representations of infinite groups for which a character theory exists, similar to classical character theory of finite groups.

Here are the essentials of the theory. A character on group G is defined as a class function $\chi : G \rightarrow \mathbb{C}$ that satisfies a non-negative definiteness property, and where $\chi(1) = 1$. (These characters are not equal to the usual characters for finite groups, but are the same after rescaling so $\chi(1)$ has the correct value). A character is said to be indecomposable if it cannot be written as a positive linear combination of distinct characters. Extended linearly, a character χ defines a state ϕ on the C^* -algebra $C^*(G)$ (that is, ϕ is a positive linear functional with $\phi(1) = 1$). States in turn give rise to cyclic representations ρ_χ of $C^*(G)$ via the Gelfand-Naimark-Siegel construction standard in C^* -theory (see for example section 1.12 of [47]). The image $\rho_\chi(C^*(G))$ then forms a von Neumann algebra of operators on a space \mathcal{H}_χ . We say ρ_χ is a factor representation if this von Neumann algebra is a factor. This happens if and only if χ is indecomposable (see eg [28] Proposition 2.21), and the factor is then of finite type, which explains the name finite-type factor representations. The definition of von Neumann algebras and their classification into factors is standard in C^* -theory, but not important in this thesis.

Overall, this process produces a bijection between indecomposable characters and finite-type factor representations. Two such representations ρ_1 and ρ_2 are quasi-equivalent if and only if they have equal characters. Quasi-equivalence is the correct notion of equivalence here: it says that the von Neumann algebras $\rho_1(C^*(G))$ and $\rho_2(C^*(G))$ are isomorphic, even if the spaces they act on might not be. So the study of representations of finite factor type reduces to finding indecomposable characters.

Every group G has a character given by the constant function 1, which is called the identity character. It is indecomposable and corresponds to the trivial representation. We can also define the regular character to be 1 on the identity and zero elsewhere. This is also a character, and is indecomposable if all non-identity conjugacy classes of G are infinite. In that case, the left regular representation of G is a factor representation. We say that G has no proper

characters if these are the only two indecomposable characters.

We can now state the main results of [27].

Theorem 4.1.3 ([27] Corollary 3.6, Corollary 3.9). *Thompson's group V has no proper characters. More generally, let G be a colour-preserving Thompson group where the set of colours is transitive and growing. Then every finite-type factor representation of G is either the regular representation, or is a 1-dimensional representation factoring through the abelianization $q : G \rightarrow G/G'$.*

In [27], this theorem is stated in terms of topological full groups of one-sided shifts, but we have been understanding these groups as colour-preserving Thompson groups. It tells us that knowledge of G/G' determines everything about finite-type factor representations of G .

We indicate how to generalize this to colour-preserving Nekrashevych-Röver groups.

Theorem 4.1.4. *Let G be a colour-preserving Nekrashevych-Röver group whose set of colours is transitive and growing. Then any indecomposable character of G is either regular or comes from a one-dimensional representation of G .*

Proof. We imitate the proof given in [27] for colour-preserving Thompson groups, in the paragraph preceding Corollary 3.9. Let $V = V_{\mathbf{C}, \mathbf{S}}$ be the underlying colour-preserving Thompson group of G , with sets of colours \mathbf{C}, \mathbf{S} . Let R be the subgroup of V defined analogously to the subgroup [27], as the set of elements of the derived subgroup V' which are the identity on some neighbourhood of some fixed $x_0 \in \mathbf{S}\mathbf{C}^\omega$. Then R is simple (as an ascending union of restrictions of V' to clopen subsets, which are simple) and has no proper characters.

Now we check that Theorem 2.10 of [27] holds for G and its subgroup R . This will tell us that every finite factor representation is either regular or factors through a quotient of G , which we have shown to be abelian. We need to find, for $g \in G$ not equal to the identity, an infinite sequence $\{g_i\}_{i \geq 1}$ of elements of the R -conjugacy class of g , such that $g_i^{-1}g_j \in R$. We will seek $g_i = r_i g r_i^{-1}$ (for $r_i \in R$). We have:

$$g_i^{-1}g_j = r_i^{-1}g^{-1}r_i r_j g r_j^{-1},$$

which we need to be the identity on a neighbourhood of the point x_0 . Since $r_i, r_j \in R$, it suffices that $g^{-1}r_i r_j g$ is the identity on a neighbourhood of x_0 , or that $r_i r_j g$ and g are locally equal around x_0 . Thus, we need that $r_i r_j$ is locally the identity around $g \cdot x_0$, which we clearly can achieve by taking $r_i \in R$ equal

to the identity locally around $g \cdot x_0$. There are infinitely many such r_i , so this works. Moreover, we need that the $r_i g r_i^{-1}$ are distinct, or in other words no $r_j^{-1} r_i$ commutes with g . To do this, take U_i disjoint open sets in the support of g (for $i \in \mathbb{N}$) such that x_0 is not in the closure of any U_i , and take r_i supported on U_i . such that g does not commute with r_i (which is certainly possible, since $V|_{U_i}$ has trivial centre). Then $r_j^{-1} r_i$ does not commute with g either, and this completes the proof. \square

4.1.3 Koopman representations

Here we build on the work in [29] to show that another family of representations of colour-preserving Thompson groups are irreducible. The representations here are Koopman representations, which are a natural class of representations defined from a group acting on a measure space. Indeed, suppose that G is a group acting on a measure space (X, μ) by transformations that preserve the measure class (that is, sets of measure zero remain of measure zero under the G -action). Then the Radon-Nikadym derivative $\frac{d\mu(g \cdot x)}{d\mu(x)}$ is defined, and we can define the Koopman representation κ , where for $g \in G$, $f \in L^2(X, \mu)$ and $x \in X$:

$$(\kappa(g)f)(x) = \sqrt{\frac{d\mu(g^{-1} \cdot x)}{d\mu(x)}} f(g^{-1} \cdot x).$$

The paper [29] proves that a certain Koopman representation of a Higman-Thompson group $V_{n,d}$ is irreducible (as well as some other related representations twisted by a cocycle). We will point out that the results generalize to colour-preserving Thompson groups. We will also use the work by one of the same authors in [26] to show that some of the Koopman representations we construct are not unitarily equivalent.

Let $V = V_{\mathbf{C}, \mathbf{S}}$ be a colour-preserving Thompson group, which acts on the space \mathbf{SC}^ω isomorphic to the ends of a tree $\mathcal{T} = \mathcal{T}_{\mathbf{C}, \mathbf{S}}$. Let $\mathbf{S} = \{s_1, s_2, \dots, s_m\}$, and let the production rule be p . We will define a probability measure on \mathbf{SC}^ω by first putting probability measures on each of the finite sets \mathbf{SC}^n . Intuitively, we choose a random end of \mathcal{T} by choosing a random root vertex and then choosing one of the edges at random to go down from each vertex. We don't insist that the edges are chosen with equal probability, but we do insist that the probabilities are consistent across vertices of the same colour.

Formally, let p_1, p_2, \dots, p_m be probabilities assigned to s_1, \dots, s_m , with each $p_i > 0$ and $\sum_{i=1}^m p_i = 1$. If $c \in \mathbf{C}$ with $p(c) = (c_1, c_2, \dots, c_r)$, then assign c_i the

probability $p_{c,i}$, where each $p_{c,i} > 0$ and $\sum_{i=1}^r p_{c,i} = 1$ for each colour c (here r is a function of c). This lets us define a Borel probability measure $\mu_{\mathbf{p}}$ on the end space of $T^{\mathcal{L}}$ as follows. Let $v = s_i x_{c_1, i_1} \dots x_{c_n, i_n} \in \mathbf{SC}^n$ (so that $s_i \in \mathbf{S}$, of colour c_1 , and thereafter $s_i x_{c_1, i_1} \dots x_{c_k, i_k}$ is the i_k th vertex below the vertex $s_i x_{c_1, i_1} \dots x_{c_{k-1}, i_{k-1}}$, with its colour the i_k th term of $p(c_k)$). Then we define:

$$\mu_{\mathbf{p}}(v\mathbf{C}^{\omega}) = \mathbf{p}(s_i)\mathbf{p}(x_{c_1, i_1}) \dots \mathbf{p}(x_{c_n, i_n}).$$

Recall that we write $v\mathbf{C}^{\omega}$ for the set $\{x \in \mathbf{SC}^{\omega} : x = vx'\}$ and call it a cylinder set (at v). Here the function \mathbf{p} is defined by $\mathbf{p}(s_i) = p_i$ and $\mathbf{p}(x_{c,i}) = p_{c,i}$. This defines μ on cylinder sets, and we check this extends to a probability measure:

Proposition 4.1.5. $\mu_{\mathbf{p}}$ extends uniquely to a well-defined probability measure on \mathbf{SC}^{ω} .

Proof. This is almost obvious from the probabilistic description, but we give a proof using Carathéodory's extension theorem. We extend the definition of $\mu_{\mathbf{p}}$ to the ring \mathcal{C} of clopen sets, which are finite unions of cylinder sets, by using the finite additivity property of measures (so the measure of a disjoint union of cylinder sets is the sum of the measures of the individual sets). Carathéodory's theorem tells us that $\mu_{\mathbf{p}}$ extends to a measure on the Borel σ -algebra, provided $\mu_{\mathbf{p}}$ is countably additive on \mathcal{C} .

Suppose that a countable disjoint union of clopen sets is clopen. By compactness, the union must be finite. So we just need to check $\mu_{\mathbf{p}}$ is finitely additive (so that the measure on \mathcal{C} is well-defined, independent of how a clopen set is broken up into cylinder sets). It's clear that if we refine a cylinder set $v\mathbf{C}^{\omega}$ into a disjoint union of sets $vx_{c,i}\mathbf{C}^{\omega}$, then $\mu_{\mathbf{p}}(v\mathbf{C}^{\omega})$ is equal to the sum of $\mu_{\mathbf{p}}(vx_{c,i}\mathbf{C}^{\omega})$ as i varies from 1 to $|p(c)|$. So given a clopen set X , and two partitions of X into disjoint cylinder sets, $X = X_1 \cup \dots \cup X_m = Y_1 \cup \dots \cup Y_n$, we can choose a common refinement of the X_i and Y_j into smaller cylinder sets. Since the refinement process preserves measure, then $\sum \mu_{\mathbf{p}}(X_i) = \sum \mu_{\mathbf{p}}(Y_j)$. This completes the proof. \square

We will write $\mu_{\mathbf{p},n}$ for the probability measure on \mathbf{SC}^n defined by:

$$\mu_{\mathbf{p},n}(\{v\}) = \mathbf{p}(s_i)\mathbf{p}(x_{c_1, i_1}) \dots \mathbf{p}(x_{c_n, i_n}),$$

for v as above. These measures on finite sets can be thought of as approximating the measure $\mu_{\mathbf{p}}$ on the limit \mathbf{SC}^{ω} . Sometimes we will just write $\mu(X)$

for $\mu_{\mathbf{p}}(X)$ or $\mu_{\mathbf{p},n}(X)$, since it should be clear from the set X which measure is meant. In the case of measures on finite sets, we'll also write $\mu(v)$ when we should formally write $\mu(\{v\})$.

To define a Koopman representation, we need to show that the action of $V_{\mathbf{C},\mathbf{S}}$ is measure class preserving. In fact we do more and work with a Nekrashevych-Röver group $V_{\mathbf{C},\mathbf{S},G}$ whose underlying colour-preserving Thompson group is $V_{\mathbf{C},\mathbf{S}}$ and with tree automorphism groups G_c for each colour $c \in \mathbf{C}$. This was done for the Higman-Thompson group in [29], and for certain groups of tree automorphisms in [26]: the work here is an extension of this to involve colours.

Theorem 4.1.6. *Let $V_{\mathbf{C},\mathbf{S},G}$ be a colour-preserving Nekrashevych-Röver group, as above. Suppose that either each group G_c is subexponentially bounded (in the sense of Definition 2 of [26]), or that $p_{c,i} = p_{c,j}$ whenever the i th and j th colours of $p(c)$ are equal. Then the action of G on $\mathbf{S}\mathbf{C}^\omega$ with the measure $\mu_{\mathbf{p}}$ is measure class preserving. In particular, the action of $V_{\mathbf{C},\mathbf{S}}$ is always measure class preserving.*

Proof. First we observe that each group G_c is measure class preserving (in its action on any subtree of $\mathcal{T}_{\mathbf{C},\mathbf{S}}$). Indeed, if the groups G_c have subexponential growth, they are measure class preserving by the argument of [26], Proposition 2. If instead $p_{c,i} = p_{c,j}$ whenever the i th and j th colours of $p(c)$ are equal, then G_c preserves the measure itself, so is definitely measure class preserving. Since every element of G is a product of an element of $V_{\mathbf{C},\mathbf{S}}$ with elements of G_c acting on subtrees, we now just need to show that the group $V_{\mathbf{C},\mathbf{S}}$ preserves the measure class.

Now suppose $g \in V_{\mathbf{C},\mathbf{S}}$, and suppose M is a set of positive measure in $\mathbf{S}\mathbf{C}^\omega$. Suppose g has the table:

$$g = \begin{pmatrix} a_1 & a_2 & \dots & a_n \\ b_1 & b_2 & \dots & b_n \end{pmatrix},$$

for some words a_i and b_i - so g acts on words in $a_i\mathbf{C}^\omega$ by replacing the initial segment a_i with b_i . This means that if X is a measurable set contained in the cylinder set $a_i\mathbf{C}^\omega$ of paths beginning a_i , then

$$\mu_{\mathbf{p}}(gX) = \mu_{\mathbf{p}}(X) \frac{\mu_{\mathbf{p},\ell(b_i)}(b_i)}{\mu_{\mathbf{p},\ell(a_i)}(a_i)}.$$

In particular, if Y is any set of non-zero measure, then $Y \cap a_i\mathbf{C}^\omega$ has non-zero

measure for some i , and so gY has non-zero measure. Hence the group G is measure class preserving. \square

We now check Dudko's measure contracting property.

Definition 4.1.7. (from [29]) *Let G act on a probability space (X, μ) by measure-class-preserving transformations. The action of G is measure contracting if for every measurable subset $M \subset X$ and any $\epsilon > 0$ there exists $g \in G$ such that:*

•

$$\mu(\text{supp}(g) \setminus M) < \epsilon.$$

•

$$\mu(\{x \in M : \sqrt{\frac{d\mu(g(x))}{d\mu(x)}} < \epsilon\}) > \mu(M) - \epsilon.$$

Here $\text{supp}(g) = \{x \in X : gx \neq x\}$.

Intuitively, the support of g is mostly contained in M , and the derivative of g is small on almost all of M . We care about this because of the following theorem:

Theorem 4.1.8. ([29]) *Let G be a group acting on a probability space (X, μ) via an ergodic measure-contracting action. Then the associated Koopman representation κ of G is irreducible.*

Recall that an action of a group G on a probability space is said to be ergodic when any G -invariant set has probability zero or one. So we need to check that the action of $V_{\mathbf{C}, \mathbf{S}, G}$ is ergodic and measure contracting. We'll be able to follow Dudko's work fairly closely here. It's enough to just consider the Thompson group $V_{\mathbf{C}, \mathbf{S}}$, since passing to a larger group will keep the action ergodic and measure contracting. So we can assume each automorphism group G_c is trivial.

Proposition 4.1.9. *Let $V = V_{\mathbf{C}, \mathbf{S}}$ be a colour-preserving Thompson group acting on \mathbf{SC}^ω with the measure $\mu_{\mathbf{P}}$. Suppose that \mathbf{C} is transitive and growing. Then the action of V is ergodic.*

Proof. Let X be a V -invariant measurable subset of \mathbf{SC}^ω , and suppose $0 < \mu_{\mathbf{P}}(X) < 1$. Let C_1, C_2 be any two disjoint cylinder sets with the same colour. Then there's an element of G which swaps the cylinder sets C_1 and C_2 , and

fixes everything else; on C_i , $\frac{d\mu(g \cdot x)}{d\mu(x)}$ is constant. This means that

$$\frac{\mu_{\mathbf{P}}(C_1 \cap X)}{\mu_{\mathbf{P}}(C_1)} = \frac{\mu_{\mathbf{P}}(C_2 \cap X)}{\mu_{\mathbf{P}}(C_2)}.$$

We can connect any two cylinder sets C, C' by a chain $C = C_0, C_1, \dots, C_m = C'$ where C_i, C_{i+1} are disjoint. So we get that $\frac{\mu_{\mathbf{P}}(C \cap X)}{\mu_{\mathbf{P}}(C)}$ does not depend on the choice of C of a given colour. Call this fraction m_c , for the colour c .

Now consider splitting up a cylinder set C , with colour c say. We partition C into sets C_1, C_2, \dots, C_r by a simple expansion (so if $C = v\mathbf{C}^\omega$, then $C_i = vx_{c,i}\mathbf{C}^\omega$). Let $p(c) = (c_1, c_2, \dots, c_r)$, where r is a function of c . Then

$$\mu_{\mathbf{P}}(X \cap C) = \sum_{i=1}^{r(c)} \mu_{\mathbf{P}}(X \cap C_i)$$

This gives:

$$m_c \mu_{\mathbf{P}}(C) = \sum_{i=1}^{r(c)} m_{c_i} \mu_{\mathbf{P}}(C_i) = \sum_{i=1}^{r(c)} m_{c_i} p_{c,i} \mu_{\mathbf{P}}(C).$$

Now let the colour c be such that $m_c \geq m_d$, whenever $d \in \mathbf{C}$. We then have

$$m_c = \sum_{i=1}^{r(c)} m_{c_i} p_{c,i} \leq \sum_{i=1}^{r(c)} m_c p_{c,i} = m_c,$$

since the sum of the $p_{c,i}$ is 1. To get equality, we must have that all m_{c_i} are equal, and equal m_c (for all colours c_i appearing in $p(c)$). We can now repeat the argument for any of these m_{c_i} , and by transitivity conclude that all m_d are equal.

Finally, we have shown that $\frac{\mu_{\mathbf{P}}(C \cap X)}{\mu_{\mathbf{P}}(C)}$ does not depend on the choice of cylinder set C , of any label. Since any measurable subset of $\mathbf{S}\mathbf{C}^\omega$ can be approximated arbitrarily well by finite unions of cylinder sets, then $\frac{\mu_{\mathbf{P}}(M \cap X)}{\mu_{\mathbf{P}}(M)}$ is the same for any measurable set M with $\mu_{\mathbf{P}}(M)$ non-zero. In particular, if $\mu_{\mathbf{P}}(X) \neq 0$, consider both $C = X$ and $C = \mathbf{S}\mathbf{C}^\omega$. Then

$$1 = \frac{\mu_{\mathbf{P}}(X \cap X)}{\mu_{\mathbf{P}}(X)} = \frac{\mu_{\mathbf{P}}(\mathbf{S}\mathbf{C}^\omega \cap X)}{\mu_{\mathbf{P}}(\mathbf{S}\mathbf{C}^\omega)} = \mu_{\mathbf{P}}(X).$$

This completes the proof. \square

The second thing to show is that the action is measure-contracting.

Proposition 4.1.10. *Let $V = V_{\mathbf{C},\mathbf{S}}$ be a colour-preserving Thompson group acting on $\mathbf{S}\mathbf{C}^\omega$ with the measure $\mu_{\mathbf{P}}$. Then the action of G is measure-contracting.*

Proof. First we work with cylinder sets rather than general measurable sets. Let $C = v\mathbf{C}^\omega$ be a cylinder set. We need to find, for every $\epsilon > 0$, an element g of G such that:

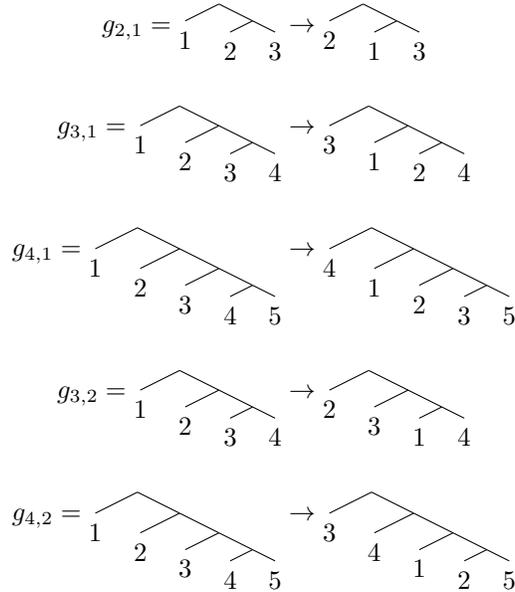
•

$$\mu_{\mathbf{P}}(\text{supp}(g) \setminus C) < \epsilon.$$

•

$$\mu_{\mathbf{P}}(\{x \in C : \sqrt{\frac{d\mu_{\mathbf{P}}(g(x))}{d\mu_{\mathbf{P}}(x)}} < \epsilon\}) > \mu(C) - \epsilon.$$

We will in fact choose g supported on C , so that the first condition is automatic. We restrict to the subgroup $V|_C$ which is isomorphic to $V_{\mathbf{C},\{c\},G}$, where $c = \chi(v)$. First we demonstrate the proof for Thompson's group V , to explain the idea. For each $k, n \in \mathbb{N}$ with $n > k$, take elements $g_{n,k}$ as demonstrated:



In general, $g_{n,k}$ is a permutation on the rightmost leaf set of $n + 1$ leaves, which moves all the leaves to the left down cyclically by k positions.

Now observe that since $g_{n,k}$ increases the depth of the top $n - k$ leaves by k . This means that their associated cylinder sets move from having measure

$q^i p$ to measure $q^{i+k} p$ (where q is the probability of moving right and p is the probability of moving left), giving a Radon-Nikodym derivative of q^{-k} across those sets. By taking n, k large, we can make this derivative as small as we please, on a set of measure $1 - q^{-(n-k)}$. We can also make this measure as close to 1 as we please, so we're done in this case.

In the general case of $V_{\mathbf{C}, \{c\}, G}$, we will use similarly defined elements $g_{n,k}$. Since \mathbf{C} is transitive, some vertex labelled c appears somewhere below the root, and without loss of generality it appears on the far right. We form finite leaf sets \mathcal{L}_n by expanding from the singleton $\mathbf{S} = \{c\}$ n times, at the rightmost vertex r_n each time. The colours of the vertices in \mathcal{L}_n apart from r_n then occur periodically, repeating every d levels (for some d). This means that whenever $n = sd$ is a multiple of d , there is an element $g_{n,rd}$ that fixes $r_n \in \mathcal{L}_n$ and moves the other elements of \mathcal{L}_n down rd layers, taken mod n . Below any leaf of depth at most $n - rd$, the Radon-Nikodym derivative of $g_{n,rd}$ is now Q^r where $Q = q_1 q_2 \dots q_d$, for q_i the probability of taking the rightmost edge at stage i . Since $0 < Q < 1$, we can make Q^r arbitrarily small and have the derivative equal to Q^r on a set of measure arbitrarily close to 1, as r increases.

This completes the proof for cylinder sets. The result extends to all measurable sets M since any measurable set can be arbitrarily well approximated by a finite union of cylinder sets, which we can assume is disjoint. Say we approximate M by $C_1 \cup \dots \cup C_n$, so that the symmetric difference of these two sets has measure less than ϵ . Choose g_i supported on C_i such that

$$\mu_{\mathbf{P}}(\{x \in C_i : \sqrt{\frac{d\mu_{\mathbf{P}}(g_i(x))}{d\mu_{\mathbf{P}}(x)}} < \frac{\epsilon}{n}\}) > \mu(C) - \frac{\epsilon}{n},$$

and define g to be the product of the g_i (which commute). Then

$$\mu_{\mathbf{P}}(\text{supp}(g) \setminus M) < \epsilon,$$

and also

$$\mu_{\mathbf{P}}(\{x \in M : \sqrt{\frac{d\mu_{\mathbf{P}}(g(x))}{d\mu_{\mathbf{P}}(x)}} < \frac{\epsilon}{n}\}) > \mu(M) - 2\epsilon.$$

This is enough to give the result. □

We summarize the result.

Proposition 4.1.11. *Let G be a colour-preserving Nekrashevych-Röver group, acting on $\mathbf{S}\mathbf{C}^\omega$ with measure $\mu_{\mathbf{P}}$ as described. Suppose that either each auto-*

morphism group G_c (for $c \in \mathbf{C}$) is subexponentially bounded, or that $p_{c,i} = p_{c,j}$ whenever the i th and j th components of $p(c)$ are equal. Then the corresponding Koopman representation κ is irreducible.

Disjointness of the representations constructed

Here we follow [26] again to argue that different choices of measure give distinct Koopman representations of colour-preserving Thompson groups. The result that we shall prove is:

Theorem 4.1.12. *Let $G = V_{\mathbf{C},\mathbf{S},G}$ for set \mathbf{C} of colours with starting set \mathbf{S} . Assume that \mathbf{C} is transitive and growing. Then:*

1. *If $\mu_{\mathbf{p}}$ is a probability measure as above, and $x \in \mathbf{S}\mathbf{C}^\omega$ is an end, then the Koopman representation $\kappa_{\mu_{\mathbf{p}}}$ and the quasi-regular representation ρ_x are not unitarily equivalent.*
2. *If \mathbf{p} and \mathbf{q} are different probability measures, then the two Koopman representations $\kappa_{\mu_{\mathbf{p}}}$ and $\kappa_{\mu_{\mathbf{q}}}$ are not unitarily equivalent.*

It's enough to prove the result for the subgroup $V_{\mathbf{C},\mathbf{S}}$ of G , whose Koopman representations are just restrictions of Koopman representations of G . Now let $\kappa = \kappa_{\mu_{\mathbf{p}}}$ be a Koopman representation of V on the Hilbert space $\mathcal{H} = L^2(\mathbf{S}\mathbf{C}^\omega, \mu)$ (so that $\kappa(g)$ is a unitary operator on \mathcal{H} , for each $g \in G$). Following [26], we make the following definitions for an open subset A of $\mathbf{S}\mathbf{C}^\omega$. Define a subset $V|_A$ of V by

$$V|_A = \{g \in G : \text{supp}(g) \subset A\}.$$

Also define a subspace \mathcal{H}_A of \mathcal{H} as the subspace fixed by $V|_A$, that is,

$$\mathcal{H}_A = \{\eta \in \mathcal{H} : \pi(g)\eta = \eta \text{ for all } g \in V|_A\}.$$

Let $\mathcal{M}_\kappa = \mathcal{M}_{\kappa_{\mathbf{p}}}$ be the von Neumann algebra generated by operators of the representation $\kappa_{\mathbf{p}}$ acting on \mathcal{H} , with commutant \mathcal{M}'_κ . We first characterize \mathcal{H}_A differently.

Lemma 4.1.13. *The set \mathcal{H}_A is equal to $\{\eta \in \mathcal{H} : \text{supp}(\eta) \subset X \setminus A\}$.*

Proof. Let $\eta \in \mathcal{H}$. It's clear that if the support of η lies outside A , then $V|_A$ fixes η . This gives a containment in one direction. Conversely, if \mathcal{H}_A were

to be strictly larger than the right hand side, there would be some $\eta \in \mathcal{H} = L^2(\mathbf{SC}^\omega, \mu)$, fixed under all of $V|_A$, and supported on A . We can also assume that η has L^2 -norm 1.

We study η by approximating with a sum of indicators of cylinder sets. For $v \in \mathbf{SC}^*$, write ξ_v for the indicator function of $v\mathbf{C}^\omega$. Given $\epsilon > 0$, it is possible to choose a finite set S of vertices v and constants α_v , such that:

$$\|\eta - \sum_{v \in S} \alpha_v \xi_v\| < \epsilon,$$

where the norm is the L^2 -norm. We can assume the cylinder sets $v\mathbf{C}^\omega$ are disjoint. On each $v\mathbf{C}^\omega$, and for each $k \in \mathbb{N}$, we show we can choose $g_{v,k}$ satisfying:

$$\text{supp}(g_{v,k}) \subset X_v \text{ and } (\kappa(g_{v,k})\xi_v, \xi_v) < 1/k,$$

where the second condition is stated using the inner product on \mathcal{H} .

To find the $g_{v,k}$, we use the same elements $g_{n,k}$ that we used to show the action is measure-contracting. We restrict the group V to $V|_{v\mathbf{C}^\omega}$, which is isomorphic to $V_{\mathbf{C}, \{\chi(v)\}, \mathcal{G}}$. Let L_n be the leaf set formed as before, by expanding the singleton $\{v\}$ n times, at the rightmost vertex r_k each time. The depth $k+1$ vertices of L_n then all lie below r_k , and have periodic colours, repeating every d levels. Let us take $n = sd, k = rd$ so that $g_{sd,rd}$ is a valid element of V (which permutes leaves of L_n by moving all vertices except r_n down rd levels, mod n). As before, there exists Q , with $0 < Q < 1$, such that the Radon-Nikodym derivative of $g_{sd,rd}$ is Q^r on a set of measure $1 - Q^{s-r}$. This set is the union of $w\mathbf{C}^\omega$, where w is a leaf of T^n moved down by $g_{sd,rd}$. On the rest of $v\mathbf{C}^\omega$, the derivative is bounded above by Q^{-s} . Using the defining formula for Koopman representations, we get:

$$(\kappa(g_{v,k})\xi_{X_v}, \xi_{X_v}) \leq (1 - Q^{s-r})\sqrt{Q^r} + Q^{s-r}\sqrt{Q^{-s}} = Q^{\frac{1}{2}r} - Q^{s-\frac{1}{2}r} + Q^{\frac{1}{2}s-r}.$$

For r fixed, taking $s \rightarrow \infty$ causes this upper bound to tend to $Q^{\frac{1}{2}r}$. Then by increasing r , we can make this as small as we please. This is enough to define the $g_{v,k}$.

Now for each $k \in \mathbb{N}$, set $h_k = \prod_{v \in S} g_{v,k}$, noticing that the $g_{v,k}$ commute since they have disjoint support. Then h_k is supported on A , and

$$\lim_{k \rightarrow \infty} (\kappa(h_k) \sum_{v \in S} \alpha_v \xi_v, \sum_{v \in S} \alpha_v \xi_v) = 0.$$

Since η is a unit vector with $\|\eta - \sum_{v \in S} \alpha_v \xi_{X_v}\| < \epsilon$, we get that:

$$\limsup_{k \rightarrow \infty} |(\kappa(h_k)\eta, \eta)| \leq 2\epsilon + \epsilon^2.$$

But we've assumed that $\kappa(h_k)\eta = \eta$ (since $h_k \in V|_A$) and that η is a unit vector. This gives a contradiction for sufficiently small ϵ . □

We use this to prove the disjointness result, following [26] again. First we quote the following well-known fact from Lemma 5 of [26], which will let us make use of the set \mathcal{H}_A :

Proposition 4.1.14. *Let π be a unitary representation of a discrete group Γ on a Hilbert space \mathcal{H} . Set $\mathcal{H}_1 = \{\eta \in \mathcal{H} : \pi(g)\eta = \eta \text{ for all } g \in \Gamma\}$. Then the orthogonal projection P onto \mathcal{H}_1 belongs to the von Neumann algebra \mathcal{M}_π .*

We apply this to the representation $\kappa = \kappa_{\mathbf{p}}$ on the subgroup $V|_A$. Then \mathcal{H}_1 in the statement is \mathcal{H}_A as defined above, and so P is the projection onto $\text{supp}(\eta) \subset X \setminus A$. This projection lies in \mathcal{M}_κ . In other words \mathcal{M}_κ contains the operator of multiplication by the characteristic function $\xi_{X \setminus A}$.

This fact is crucial for the proof of disjointness, which we prove now. This continues to follow the methods of [26].

Proof of Theorem 4.1.12. We do the two parts separately.

1. First we want to show that Koopman and quasi-regular representations are disjoint. Let $x \in \mathbf{S}\mathbf{C}^\omega$, with quasi-regular representation $\rho_x = \rho_{G/P}$, for P the stabilizer of x . If A is an open subset of $\mathbf{S}\mathbf{C}^\omega$, then let \mathcal{H}_A be as before, for the measure $\mu_{\mathbf{p}}$, and define the equivalent subspace for ρ_x :

$$\mathcal{H}_A^x = \{\eta \in l^2(Gx) : \rho_x(g)\eta = \eta \text{ for all } g \in V|_A\},$$

and let P_A^x be the orthogonal projection onto this subspace. Suppose that the representations are equivalent, and that the isometry $U : L^2(X, \mu_{\mathbf{p}}) \rightarrow l^2(Gx)$ intertwines $\kappa_{\mu_{\mathbf{p}}}$ and ρ_x (ie $U\kappa_{\mu_{\mathbf{p}}}(g) = \rho_x(g)U$). Choose a sequence of open covers A_n of Gx whose $\mu_{\mathbf{p}}$ -measure tends to 0 (which is possible, since the orbit is countable). Then for all n :

$$UP_{A_n}U^* = P_{A_n}^x.$$

The set A_n is open and meets Gx (which is dense in \mathbf{SC}^ω), and the orbit of any $y \in Gx$ under G_{A_n} is infinite. This implies that $P_{A_n}^x = 0$ - any invariant vector would need to contain an orbit sum, which cannot be in ℓ^2 . But since the measure of A_n tends to 0, $P_{A_n} \rightarrow \text{Id}$ weakly as $n \rightarrow \infty$; this gives a contradiction.

2. Let $\mu_{\mathbf{p}}, \mu_{\tilde{\mathbf{p}}}$ be distinct probability measures of the type considered; define $\mathcal{H}_A, \tilde{\mathcal{H}}_A$ as before (for each open $A \subset \partial T^\mathcal{L}$), with orthogonal projections P_A, \tilde{P}_A onto them. We claim there exists a subset $A \subset X$ with $\mu_{\mathbf{p}}(A) = 0$ and $\mu_{\tilde{\mathbf{p}}}(A) = 1$. Indeed, consider choosing a random end $vx_{c_1i_1}x_{c_2i_2}\dots \in \mathbf{SC}^\omega$ using the measure $\mu_{\mathbf{p}}$. We think of this as choosing a random path down the tree $\mathcal{T}_{\mathbf{C},\mathbf{S}}$, where at a vertex coloured c , we choose the i th edge with probability $p_{c,i}$. By the law of large numbers, with probability 1, the fraction of vertices coloured c which are followed by the i th edge tends to p_i (whereas for a different measure $\tilde{\mathbf{p}}$ it tends to \tilde{p}_i). This provides a set that has measure 1 for $\tilde{\mathbf{p}}$ and measure 0 for any $\mathbf{p} \neq \tilde{\mathbf{p}}$.

Now assume that $\kappa_{\mathbf{p}}$ and $\kappa_{\tilde{\mathbf{p}}}$ are equivalent via isometry U : that is,

$$U(\kappa_{\mathbf{p}}(g)\eta) = \kappa_{\tilde{\mathbf{p}}}(g)U(\eta)$$

holds for all $g \in G$ and $\eta \in L^2(\partial T^\mathcal{L}, \mu_{\mathbf{p}})$. Take a sequence (A_n) of open covers of A with $\mu_{\mathbf{p}}(A_n) \rightarrow 0$ as $n \rightarrow \infty$; since all the A_n contain A , $\mu_{\tilde{\mathbf{p}}}(A) = 1$ for every n . Then U conjugates \mathcal{H}_{A_n} to $\tilde{\mathcal{H}}_{A_n}$ and P_{A_n} to \tilde{P}_{A_n} . But P_{A_n} is the orthogonal projection onto the complement of A_n , so \tilde{P}_{A_n} tends weakly to the identity whilst \tilde{P}_A is zero; this gives a contradiction.

□

4.2 A Hecke algebra for Thompson's group V

In the rest of this thesis we produce an algebra $\mathcal{H}_{V,q}$ that - as far as possible - deforms the group algebra of a Higman-Thompson group $V_{n,d}$. Our inspiration is the theory of Hecke algebras of type A_n , which are deformations of the group algebra of \mathfrak{S}_n . The principal aim of this section is to see whether the parallels between \mathfrak{S}_n and V carry over into parallel deformations, which we could use to learn about representations of V . We shall see that there are many obstructions to the existence of such an $\mathcal{H}_{V,q}$ but will construct the best algebra possible.

An algebra $\mathcal{H}_{V,q}$ of this type could tell us two things about representations. First, we could generalize the question of finding representations of V into finding representations of $\mathcal{H}_{V,q}$. The question of which irreducible representations of V deform is interesting (and includes all irreducible representations in the \mathfrak{S}_n case). Second, our discussion of the Hecke algebra will produce a ‘general linear’ group $GL(\Gamma, \mathbb{F})$, as a subset of the Leavitt path algebra of suitable graph Γ , with V viewable as the ‘permutation matrices’ within it. It seems worth studying the relations between V and $GL(\Gamma, \mathbb{F})$, to see if there are parallels in their representation theory.

We give the classical theory first.

4.2.1 The classical theory of Hecke algebras

The symmetric group as a Coxeter group First we recall some classical theory of symmetric groups. A good reference for the theory of finite reflection groups and Coxeter groups is the first two chapters of [7]. Let \mathfrak{S}_n be the symmetric group on n objects. We will have \mathfrak{S}_n act on $\{1, 2, \dots, n\}$ on the left. Recall that \mathfrak{S}_n is generated by $n - 1$ transpositions $\sigma_1, \sigma_2, \dots, \sigma_{n-1}$, where σ_i is the transposition $(i \ i + 1)$. It has a presentation with these generators, and relations:

- $\sigma_i^2 = 1$ for all $1 \leq i \leq n - 1$.
- $\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$ for $1 \leq i \leq n - 2$.
- $\sigma_i \sigma_j = \sigma_j \sigma_i$ whenever $|i - j| > 1$.

Recall that a Coxeter group is a group with a presentation with generators $\{r_i\}_{i \in I}$ (for a finite set I) and relations $r_i^2 = 1$, $(r_i r_j)^{m_{i,j}} = 1$. This presentation of \mathfrak{S}_n tells us that it is a Coxeter group. Moreover, finite Coxeter groups are the same as finite groups generated by Euclidean reflections (see eg Section 2.5.4 of [7]). We explain how to realize \mathfrak{S}_n as a group of reflections, from [7] Section 1.4.7.

Let \mathcal{H} be the set of $\binom{n}{2}$ hyperplanes $x_i = x_j$ in \mathbb{R}^n (for each $i \neq j$, where x_i are the n coordinates of a point in \mathbb{R}^n). We call this configuration the braid arrangement. Let G be the group generated by reflections in the hyperplanes \mathcal{H} . The connected components of $\mathbb{R}^n \setminus \mathcal{H}$ are called chambers, and can be specified by giving an ordering of the coordinates,

$$x_{\pi(1)} > x_{\pi(2)} > \dots > x_{\pi(n)},$$

for some permutation π . In particular, there is a chamber $x_1 > x_2 > \dots > x_n$ which we call the fundamental chamber. G then acts faithfully and transitively on the chambers, giving an isomorphism from G to \mathfrak{S}_n , where the reflection in hyperplane $x_i = x_j$ corresponds to the transposition $(i\ j)$. In particular, the Coxeter generators $(i\ i+1)$ of \mathfrak{S}_n correspond to reflections in the hyperplanes that border the fundamental chamber C . We also note that two chambers that share a face on hyperplane $x_i = x_j$ differ only by swapping i and j in their associated permutations (that is, replacing σ with $(i\ j)$.)

The Coxeter presentation of \mathfrak{S}_n has particularly nice properties which it shares with other Coxeter groups. Namely, let $S = \{\sigma_1, \dots, \sigma_{n-1}\}$ be the generating set, and let $w = s_1 s_2 \dots s_l$ be a word over S , with $s_i \in S$, representing an element $\sigma \in \mathfrak{S}_n$. We say that l is the length of w , and we say w is reduced if any other word representing the same permutation σ has length at least w . We define the length of a permutation σ to be the length of any reduced word representing it. We can represent the word w as a walk through the chambers of the braid arrangement, travelling through $C, s_1 C, s_1 s_2 C, \dots$ and finally to wC . Each two consecutive chambers in this walk share an $n - 1$ -dimensional face. Then the length of w is equal to the shortest possible walk from C to σC , which is also equal to the number of pairs $\{i, j\}$ such that $i < j$ but $\sigma(i) > \sigma(j)$. This provides an easy test for words being reduced: just calculate the associate permutation and count how many pairs are out of order.

Next, we give a more interesting result. Suppose that the word w is not reduced. Then one can shorten w by writing

$$s_1 \dots, \hat{s}_i \dots \hat{s}_j \dots s_n,$$

where \hat{s}_i indicates that the permutation s_i has been dropped from w . This result is quite deep and is special to Coxeter groups. In particular, it implies that the length of any word w representing σ is of the same parity, showing that the alternating group is well-defined. Finally, if w and w' are two different reduced words representing σ , then w can be converted into w' by the two operations of replacing $\sigma_i \sigma_j$ with $\sigma_j \sigma_i$ (for $|i - j| > 1$) and replacing $\sigma_i \sigma_{i+1} \sigma_i$ with $\sigma_{i+1} \sigma_i \sigma_{i+1}$. A form of this result is proved for all Coxeter groups in Theorem 2.33 of [7]. This is a strong statement, because normally, checking that two words are equal in a group presentation might require both lengthening and shortening operations. In Coxeter groups, one never needs to make words longer in this procedure, and we'll find this fact very useful later.

4.2.2 Hecke algebras from double cosets

We now introduce Hecke algebras. There are many different objects called Hecke algebras, most of which are related by being some sort of endomorphism algebra of an induced module. We will give an overview of a few different definitions, chosen for their relevance to \mathfrak{S}_n and its representation theory.

First of all we give a simple definition in terms of double cosets. A reference for the basics is [35]. Let G be a group and let H be a commensurated subgroup, which we recall means that $gHg^{-1} \cap H$ is of finite index in H and gHg^{-1} , for all $g \in G$. Let \mathbf{k} be a field. We define $L(G, H)$ to be the \mathbf{k} -vector space with basis the set of left cosets gH of H in G . To make this into an algebra, we would like to introduce the multiplication $g_1H \cdot g_2H = g_1g_2H$, but this is not well-defined in general (when the subgroup H is not normal). To remedy this, we restrict to the set $L(G, H)^G$ of G -invariants of the left coset space. That is,

$$L(G, H)^G = \{X \in L(G, H) : g \cdot X = X \text{ for all } g \in G\}.$$

Then there is a well-defined multiplication $L(G, H) \times L(G, H)^G \rightarrow L(G, H)$, linearly extended from:

$$gH \cdot \left(\sum_{i=1}^k \lambda_i g_i H \right) = \sum_{i=1}^k \lambda_i g g_i H,$$

for $\lambda_i \in \mathbf{k}$ and g_i, G in G . This restricts to a multiplication $L(G, H)^G \times L(G, H)^G \rightarrow L(G, H)^G$. We write $\mathcal{H}(G, H)$ for $L(G, H)^G$ with this multiplication, and call it the *Hecke algebra* of G , over the subgroup H .

Before giving some examples, we make the remark that if HgH is a double coset, then we can write

$$HgH = gH \sqcup g_2H \sqcup g_3H \sqcup \dots \sqcup g_mH,$$

as a disjoint union of left cosets. The union is finite when H is commensurated by G . Put $g_1 = g$. Then $\sum_{i=1}^m g_i H$ is an element of $\mathcal{H}(G, H)$, and it is easy to check that these elements form a basis for $\mathcal{H}(G, H)$ in the case that the index of H in G is finite. We now give a couple of basic examples before going on to the most important example, that will inspire our work on Higman-Thompson groups.

First, suppose that H is a normal subgroup of G . Then $L(G, H)^G = L(G, H)$

since every coset is G -invariant. Thus we see that $\mathcal{H}(G, H)$ is isomorphic to the group algebra of the quotient group, $\mathbf{k}G/H$. So Hecke algebras can be viewed as a way of making quotient groups work - as near as possible - for subgroups that are not normal.

Second, let $G = \mathfrak{S}_n$ and $H = \mathfrak{S}_{n-1} = \text{Stab}(n)$, in the usual action of \mathfrak{S}_n on $\{1, 2, \dots, n\}$. Left cosets of H are given by sets:

$$H_i = \{g \in G : g(n) = i\},$$

with $H = H_n$, and we can choose as coset representatives the transpositions $\tau_i = (i \ n)$. Then one can verify that H has two double cosets, H itself and $\tau_1 H \cup \dots \cup \tau_{n-1} H$ - this is equivalent to saying that H acts transitively on the set $\{1, 2, \dots, n-1\}$. Thus the algebra $\mathcal{H}(G, H)$ is 2-dimensional, spanned by H and X , where

$$X = \sum_{i=1}^{n-1} \tau_i H.$$

To find the multiplication on the Hecke algebra, we remark that H is the identity of the algebra (which always happens), so it suffices to find X^2 . This is:

$$\begin{aligned} X^2 &= \sum_{i,j=1}^{n-1} \tau_i \tau_j H \\ &= \sum_{i=1}^{n-1} \left(H + \sum_{j=1, j \neq i}^{n-1} (j \ i \ n) H \right) \\ &= \sum_{i=1}^{n-1} \left(H + \sum_{j=1, j \neq i}^{n-1} (j \ n) H \right) \\ &= (n-1)H + (n-2)X \end{aligned}$$

This completely determines the Hecke algebra. The multiplication is easy to compute, in part because there's a natural interpretation of H as the stabilizer of some object.

Before giving the most important example, we remark that Hecke algebras can be understood as endomorphism algebras of an induced module. We work over a field \mathbf{k} . Then $L(G, H)$ with its natural G -action is the induced module $\text{Ind}_H^G \mathbf{k}$, where \mathbf{k} here is seen as the trivial $\mathbf{k}H$ -module. A G -endomorphism θ of $L(G, H)$ is entirely determined by $\theta(H)$, which must satisfy $h\theta(H) = \theta(hH) =$

$\theta(H)$ for all $h \in H$. Moreover, any H -invariant choice of $\theta(H)$ extends to a G -endomorphism. Thus $\text{End}_{\mathbf{k}G} \text{Ind}_H^G \mathbf{k}$ is in bijection with $L(G, H)^H$, and one can check that the multiplication on endomorphisms is compatible with the Hecke algebra multiplication. So we obtain:

$$(\text{End}_{\mathbf{k}G} \text{Ind}_H^G \mathbf{k}) \cong \mathcal{H}(G, H).$$

This result is the beginning of the connection between Hecke algebras and representations. Indeed, this fact makes $\text{Ind}_H^G \mathbf{k}$ into a \mathcal{H}, G -bimodule (where, say, $\mathcal{H} = \mathcal{H}(G, H)$ acts on the left and G acts on the right). If V is then any right \mathcal{H} -module, we can form

$$V \otimes_{\mathcal{H}} \text{Ind}_H^G \mathbf{k},$$

and this is a right G -module. This provides a source of G -modules from representations of \mathcal{H} , and we can study how these decompose. Often, \mathcal{H} is easier to find modules for than G , so this is a good source of representations of G .

We now give one final example of a Hecke algebra.

Endomorphisms of flag space The most important and common example of a Hecke algebra comes from the Lie group GL_n (of type A_n) over its Borel subgroup. This is called the Hecke algebra of type A_n . Frequently, the term Hecke algebra is used to mean this example specifically, or perhaps a similar example from a Lie group of another type.

Let \mathbf{k} be a field, usually \mathbb{C} , and let G be the group $GL_n(\mathbb{F}_q)$ for some $n \in \mathbb{N}$ and finite field \mathbb{F}_q (of order q). Let B be a Borel subgroup of G , which we take to be the subgroup consisting of upper triangular matrices. We will define an algebra $\mathcal{H}_{n,q}$ from the Hecke algebra $\mathcal{H}(G, B)$. We will give a presentation for $\mathcal{H}_{n,q}$, which will be a deformation of a presentation of $\mathbf{k}\mathfrak{S}_n$, the group algebra of the symmetric group.

Recall that $\mathcal{H}(G, B)$ is defined as the set of G -endomorphisms of the left coset space $L(G, B)$. We saw in the \mathfrak{S}_n example that it can be helpful to think about an action of G where B is a point stabilizer. With this in mind, define a *flag* F to be a sequence of subspaces of $W = \mathbb{F}_q^n$:

$$F : \{0\} = W_0 \subset W_1 \subset \dots \subset W_{n-1} \subset W_n = W,$$

such that W_i has codimension 1 in W_{i+1} for each $1 \leq i \leq n-1$. We write $F(i)$

for the subspace W_i . Since G acts on W linearly, it acts on the set of flags; let \mathcal{F} be the $\mathbf{k}G$ -module with the set of flags as basis. Observe that if e_1, e_2, \dots, e_n is the standard basis of \mathbb{F}_q^n , then B is the stabilizer of the *standard flag* F_{st} :

$$F_{st} : \{0\} \subset \langle e_1 \rangle \subset \langle e_1, e_2 \rangle \subset \dots \subset \langle e_1, e_2, \dots, e_n \rangle = W.$$

Thus \mathcal{F} is isomorphic as a G -module to $L(G, B)$.

For $1 \leq i \leq n - 1$, we define an endomorphism Σ_i of the G -module \mathcal{F} as follows: we say flags F_1 and F_2 are i -neighbours if $F_1(i) \neq F_2(i)$, but $F_1(j) = F_2(j)$ for all $j \neq i$. Then we define $\Sigma_i(F)$ to be the sum of all flags that are i -neighbours of F . It is then easy to see that Σ_i is an endomorphism of the $\mathbf{k}G$ -module \mathcal{F} .

It can be shown that the Hecke algebra $\mathcal{H}(G, B)$ of all endomorphisms of \mathcal{F} is generated by the Σ_i . It has a presentation with generators the Σ_i and with relations:

- $\Sigma_i^2 = (q - 1)\Sigma_i + q$ for each $1 \leq i \leq n - 1$
- $\Sigma_i \Sigma_{i+1} \Sigma_i = \Sigma_{i+1} \Sigma_i \Sigma_{i+1}$ for each $1 \leq i \leq n - 2$
- $\Sigma_i \Sigma_j = \Sigma_j \Sigma_i$ otherwise.

We define $\mathcal{H}_{n,q}$ to be the \mathbf{k} -algebra presented by these generators and relations: this algebra equals $\mathcal{H}(G, B)$ when q is a prime power, but is defined for any $q \in \mathbf{k}$. In particular, if $q = 1$, then $\mathcal{H}_{n,0} \cong k\mathfrak{S}_n$, where the isomorphism sends Σ_i to the transposition $\sigma_i = (i \ i + 1)$, and the relations give a presentation for \mathfrak{S}_n as a Coxeter group. We call $\mathcal{H}_{n,q}$ the Hecke algebra of type A_n , with parameter q .

There is a natural basis for $\mathcal{H}_{n,q}$ labelled by the symmetric group. If $\sigma \in \mathfrak{S}_n$, write $\sigma = \sigma_{i_1} \dots \sigma_{i_k}$ as a reduced word in the Coxeter generators of \mathfrak{S}_n . Define $\Sigma = \Sigma(\sigma) \in \mathcal{H}_{n,q}$ to be the element $\Sigma_{i_1} \dots \Sigma_{i_k}$. This definition does not depend on the choice of reduced expression for σ , because we know that any two reduced expressions can be related by the relations $\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$ and $\sigma_i \sigma_j = \sigma_j \sigma_i$, which also hold for Σ_i . As σ varies, these elements Σ then form a basis for $\mathcal{H}_{n,q}$. In particular, $\mathcal{H}_{n,q}$ has dimension $n!$ (regardless of the parameter q). See [32] for proofs.

The connections between \mathfrak{S}_n and $\mathcal{H}_{n,q}$ go much further, and [32] is a good summary. Perhaps the most important result is that $\mathcal{H}_{n,q}$ has a family of modules called the *cell modules*, labelled by partitions λ of n . This generalizes

the familiar Specht modules of \mathfrak{S}_n . The cell modules give a complete set of nonisomorphic simple modules for $\mathcal{H}_{n,q}$ if $\mathcal{H}_{n,q}$ happens to be semisimple, which occurs whenever the characteristic of \mathbf{k} does not divide $n!$. The study of $\mathcal{H}_{n,q}$ then goes on to analyse how the cell modules decompose into simple modules in the remaining cases.

$\mathcal{H}_{n,q}$ as quotient of a braid group algebra There is another important way to define $\mathcal{H}_{n,q}$. Recall that the group presented with generators $\Sigma_i : 1 \leq i \leq n-1$ and relations $\Sigma_i \Sigma_{i+1} \Sigma_i = \Sigma_{i+1} \Sigma_i \Sigma_{i+1}$ and $\Sigma_i \Sigma_j = \Sigma_j \Sigma_i$ for $|i-j| > 1$ is called the Artin braid group B_n (on n strands). Elements of B_n can be described using arrangements of n strings in space, similar to diagrams of permutations in \mathfrak{S}_n but where we care about the orientations when two lines cross each other. There is a homomorphism from B_n to \mathfrak{S}_n , defined by forgetting whether strings pass over or under each other, and just taking the permutation of the strings given by a braid. In terms of presentations, we add the relations $\Sigma_i^2 = 1$ for all i . If instead we add the relations $\Sigma_i^2 = (q-1)\Sigma_i + q$ to $\mathbf{k}B_n$, we get the Hecke algebra. We will try to imitate this as one of the main ideas in constructing a Hecke algebra for V .

Hecke algebras, generalizations, and representation theory There are many variants of Hecke algebras available for different groups and subgroups. For example, one can define a Hecke algebra for a compact open subgroup P of a locally compact group G , or for a closed subgroup P of a profinite group G . In these cases, the Hecke algebra looks at (P, P) -invariant functions on G with some finiteness property, under a convolution product. We're interested in V as an analogue to the group \mathfrak{S}_n , so we will not concern ourselves as much with the topology.

Another interesting generalization replaces the trivial module with an arbitrary H -module. We sketch this from chapter 12 of [24] (which is a good introduction to the use of Hecke algebras in representation theory). Work over the field \mathbb{C} here, so that we can talk about characters of G . Let G be a group and H a subgroup. Let ψ be a character of H , and let e be an idempotent of the group algebra $\mathbb{C}H$ such that $\mathbb{C}He$ is an irreducible $\mathbb{C}H$ -module with character ψ (such an e exists: $\mathbb{C}H$ certainly contains an irreducible submodule with any given character as a direct summand, and then we use the fact that an ideal of a semisimple ring is generated by an idempotent). Then the Hecke algebra $\mathcal{H}(G, H, \psi)$ is defined to be the subalgebra $e\mathbb{C}Ge$ of $\mathbb{C}G$. It is shown that $\mathbb{C}Ge$

affords the induced character $\text{Ind}_H^G \psi$, and that $\text{End}_{\mathbb{C}G} \mathbb{C}Ge$ is isomorphic to (the opposite algebra of) $\mathcal{H}(G, H, \psi)$. This generalizes what we saw for Hecke algebras with trivial characters (albeit with the opposite algebra here, because of different side conventions).

Now we state a result given as Theorem 11.25 of [24]. Let χ be an irreducible character of G . Then since we've given $\mathcal{H} = \mathcal{H}(G, H, \psi)$ as a subalgebra of $\mathbb{C}G$, we can consider the restriction of χ to \mathcal{H} . The theorem says that the restriction is non-zero if and only if χ appears as a component of $\text{Ind}_H^G \psi$, with restriction setting up a bijection between irreducible characters of G meeting $\text{Ind}_H^G \psi$ and irreducible characters of the algebra \mathcal{H} . Moreover, the inner product $\langle \chi, \text{Ind}_H^G \psi \rangle$ gives the degree of the character $\chi|_{\mathcal{H}}$.

This result has a nice consequence when \mathcal{H} is abelian, which is often easy to check. Then all characters of \mathcal{H} are linear, and so every irreducible character of G appears at most once in $\text{Ind}_H^G \psi$. Thus, we can learn about the representations of G and H by studying the algebras \mathcal{H} .

4.3 Obstructions to building a Hecke algebra

We are eventually going to construct an object $\mathcal{H}_{V,q}$ that we will call a Hecke algebra for the Higman-Thompson groups $V = V_{n,d}$. The aim will be that, in some sense, $\mathcal{H}_{V,q}$ is to V as the Hecke algebra $\mathcal{H}_{n,q}$ of type A_n is to \mathfrak{S}_n . This is purposefully vague, because there are many connections between Hecke algebras and \mathfrak{S}_n which cannot be maintained. In order to motivate its definition, which will seem quite far from classical Hecke algebras, we first explain why simpler approaches cannot work. This means that this section will read pessimistically, since we spend our time beginning several natural approaches and then finding that they fail. However, the objects we describe in this section are still of interest.

We point out that the classical Iwahori-Hecke algebra is a feature in the representation theory of $GL_n(F)$ (it's an algebra of endomorphisms of a particular induced module) and is not as significant in the representation theory of \mathfrak{S}_n . So building a Hecke algebra here (by deforming V) should tell us not so much about V but about some larger 'general linear' group. However, the irreducible representations of \mathfrak{S}_n do generalize to irreducible representations of its Hecke algebra, and studying the class of representations of V that can be similarly deformed could be interesting. We don't study (double coset) Hecke algebras

for V over its subgroups here; this might be possible, but the lack of commensuration makes this task difficult. Essentially, though, our aim is to produce an algebra that deforms $\mathbf{k}V$.

We initially suggest four reasonable approaches to the task of building a Hecke algebra for $V_{n,d}$. We list them here.

1. **Deform a presentation for the Thompson group.** We know that the algebra $\mathcal{H}_{n,q}$ can be defined by starting with the Coxeter presentation of \mathfrak{S}_n and changing the quadratic relations. Like the symmetric group, Thompson's groups are generated by transpositions which satisfy many of the same relations. Perhaps we could take a presentation for $V_{n,d}$, with generators the transpositions, and just rewrite the quadratic relations to get a new algebra.
2. **Define an equivalent to the general linear group, and study its action on some version of flags.** The symmetric group can be identified with the permutation matrices of GL_n , which are the unitary matrices of GL_n whose entries are all 0 or 1. Moreover, \mathfrak{S}_n forms a set of coset representatives for GL_n over its Borel subgroup. We know that $V_{n,d}$ can be constructed as a group of particular unitary elements of a Leavitt path algebra, which have coefficients 0 or 1. This is strikingly similar to the symmetric group case. This motivates defining a general linear group inside the Leavitt path algebra, and trying to find endomorphisms of its action on flags.
3. **Fit together copies of $\mathcal{H}_{n,q}$.** One of the main reasons that Higman-Thompson groups 'feel like' symmetric groups is that they contain many copies of finite symmetric groups (including one for each leaf set). We already know how to deform \mathfrak{S}_n into a Hecke algebra. If we can find appropriate maps between the finite $\mathcal{H}_{n,q}$, we will have built an analogue to V .
4. **Add relations to a braided Thompson group.** Hecke algebras can be defined as quotients of a braid group. Braided versions of Thompson's group have been studied (where one braids the pictures of permutations in diagrams like Figure 3.1.1). Adding quadratic relations to a braided Thompson group could give an interesting quotient.

In the rest of this section, we study these ideas one at a time to show what progress can be made with them. We will prove some results that are

interesting in their own right but don't ultimately give a satisfactory Hecke algebra. Instead, we'll learn the compromises that have to be made in the definitions later. We will make some preliminary definitions but then state where they stop working. The algebra we end up defining can be understood in terms of any of these approaches, although it doesn't fit any of them perfectly.

For simplicity, we will work with Thompson's original group V in this section (which you recall is $V_{2,1}$ as a Higman-Thompson group). We return to the generality of $V_{n,d}$ to define an algebra in the next section.

4.3.1 Deforming a presentation for V

First of all, we try to find a presentation for V and deform it. We will talk about V acting on the ends of the infinite binary tree \mathcal{T} , whose vertices we identify with X^* , for $X = \{a, b\}$. Moreover, if $\mathcal{L}_1, \mathcal{L}_2$ are leaf sets of \mathcal{T} and $\phi : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ is a bijection, then we'll talk about $\bar{\phi} \in V$ acting on all vertices of \mathcal{T} below \mathcal{L}_1 . The action is defined for sufficiently deep $v \in \{a, b\}^*$ by $\bar{\phi}(v) = \phi(l)v'$, where l is the unique element of \mathcal{L}_1 such that v can be written $v = lv'$. This gives a partial action of V on \mathcal{T} . The usual action on $\partial\mathcal{T}$ can be thought of as a limit of this action.

We seek a presentation in terms of transpositions, and we want the presentation to resemble the Coxeter presentation for \mathfrak{S}_n as much as possible. Thus, most of the short finite presentations for V won't actually be helpful here, because their relations are complicated, and it's not clear how to deform them.

A good candidate for a Coxeter-type presentation for V comes from the following result, due to Collin Bleak and Martyn Quick ([14]). Take a symbol $s_{v,w}$ for each (unordered) pair of incomparable leaves v and w . Then V is isomorphic to the group generated by the symbols $s_{v,w}$ with the relations, for any leaves v, w :

1. $s_{v,w}^2 = 1$
2. $s_{x,y}^{-1} s_{v,w} s_{x,y} = s_{(x\ y)\cdot v, (x\ y)\cdot w}$ whenever $(x\ y)\cdot v$ and $(x\ y)\cdot w$ are defined.
3. $s_{v,w} = s_{va,wa} s_{vb,wb}$

This presentation has a quadratic relation and a conjugation relation, which generalize the relations between transpositions in \mathfrak{S}_n . It also has an expansion relation, where a transposition in one layer is written as a product of two transpositions in the next. This is a phenomenon peculiar to V , and is one of the

main reasons why forming a Hecke algebra will be difficult. The presentation given still has too many generators for our purposes though - in order to get a Coxeter-like presentation, we only want to take transpositions of adjacent elements (in some sense). So we will refine this to a presentation whose generators are symbols $t_{v,w}$ where v and w are *adjacent* leaves of some leaf set. If $v, w \in X^*$ and v is to the left of w , it is possible for v and w to be adjacent leaves if and only if w is formed from v by replacing a terminal string ab^k with ba^l , for some $k, l \in \mathbb{N}_0$.

Proposition 4.3.1. *V has a presentation as follows. The generating set has symbols $t_{v,w}$ for each pair $v, w \in X^*$, where v is immediately to the left of w in some leaf set. We write this relation as $v \sim_L w$. The relations on such symbols are:*

$$I: t_{v,w}^2 = 1 \text{ for all } v, w \text{ with } v \sim_L w.$$

II: $t_{v,w}$ and $t_{x,y}$ commute whenever $v \sim_L w, x \sim_L y$, and $v, w, x, y \in X^*$ are all incomparable.

$$III: t_{v,w}t_{w,x}t_{v,w} = t_{w,x}t_{v,w}t_{w,x} \text{ for all } v, w, x \in X^* \text{ with } v \sim_L w \sim_L x.$$

$$IV: t_{v,w} = t_{vb,wa}t_{va,vb}t_{wa,wb}t_{vb,wa} \text{ for all } v, w \in X^* \text{ with } v \sim_L w.$$

$$V: t_{vb,wa}t_{va,vb}t_{vb,wa} = t_{v,w}t_{v,wa}t_{v,w} \text{ and } t_{v,w}t_{vb,wa}t_{v,w} = t_{v,wa}t_{wa,wb}t_{v,wa} \text{ for all } v, w \in X^* \text{ with } v \sim_L w.$$

$$VI: t_{w,x}t_{v,wa}t_{w,x} = t_{w,xa}t_{v,w}t_{w,xa} \text{ for all } v, w, x \in X^* \text{ with } v \sim_L w \sim_L x.$$

Whenever we write down a term such as $t_{vb,wa}$, it requires that $vb \sim_L wa$. It is easy to check that this relation is a consequence of $v \sim_L w$ (and similar facts hold elsewhere).

The proof of this will occupy most of the rest of this subsection on presentations. At the end, we'll discuss this presentation and say why it is not ultimately productive for defining an algebra. Nevertheless, it's an interesting presentation for V .

We call the group specified by this presentation V_t , and claim it is isomorphic to V (where $t_{v,w}$ represents the transposition swapping v and w). First we show that $t_{v,w} \mapsto s_{v,w}$ defines a homomorphism from V_t to Bleak and Quick's presentation of V . This is done by showing that the relations among the $t_{v,w}$ hold in V . Relation I is clear, and relations II, III, V, and VI are a special case

of Bleak and Quick's relation 2. Our relation IV is a rewritten version of their relation 2, where we observe:

$$\begin{aligned} s_{v,w} &= s_{va,wa} s_{vb,wb} = (s_{vb,wa} s_{va,wa} s_{vb,wa}) (s_{vb,wa} s_{wa,wb} s_{vb,wa}) \\ &= s_{vb,wa} s_{va,vb} s_{wa,wb} s_{vb,wa}. \end{aligned}$$

Informally, our presentation has a quadratic relation, a commutativity relation and a braid relation (I-III) - just like for \mathfrak{S}_n - as well as a relation (IV) that expands transpositions at one layer to the next. It also has two less obvious relations V and VI, which seem to be necessary to deal with transpositions at different layers interacting.

To prove Proposition 4.3.1 we need to show that the homomorphism $t_{v,w} \mapsto s_{v,w}$ has an inverse. In other words, we will show that all the relations in Bleak and Quick's presentation can be derived from relations I - VI. We do this now. First we extend the definition of $t_{v,w}$ to define more elements of V_t .

Definition 4.3.2. *We define the symbol $t_{v,w} \in V_t$ for general leaves $v, w \in X^*$ (with v left of w) as follows: take any leaf set containing v and w , and let the leaves from v to w (left to right) be $v_0 = v, v_1, v_2, \dots, v_n = w$. Then we define:*

$$t_{v,w} = t_{v_{n-1},v_n} \cdots t_{v_2,v_3} t_{v_1,v_2} t_{v_0,v_1} t_{v_1,v_2} t_{v_2,v_3} \cdots t_{v_{n-1},v_n}.$$

We will prove that all of Bleak and Quick's relations hold for these symbols $t_{v,w}$, so that $s_{v,w} \mapsto t_{v,w}$ extends to a homomorphism.

Lemma 4.3.3. *The symbol $t_{v,w}$ is well-defined and is independent of the choice of v_1, v_2, \dots, v_n .*

Proof. Given two different choices of v_i , we can find a third choice of v_i that is an expansion of both of them. Thus it's sufficient to establish that two choices of v_i differing by a simple expansion define the same $t_{v,w}$. So it's enough to show the following two elements of V_t are equal:

$$t_{v_{k-1},v_k} \cdots t_{v_2,v_3} t_{v_1,v_2} t_{v_0,v_1} t_{v_1,v_2} t_{v_2,v_3} \cdots t_{v_{k-1},v_k}$$

and

$$\begin{aligned} t_{v_{k-1},v_k} t_{v_{k-1},v_k} a t_{v_{k-2},v_{k-1}} b t_{v_{k-2},v_{k-1}} a \cdots t_{v_1,v_2} t_{v_0,v_1} t_{v_1,v_2} \cdots \\ t_{v_{k-2},v_{k-1}} a t_{v_{k-1},v_k} b t_{v_{k-1},v_k}. \end{aligned}$$

We do this by induction on n . We'll do the inductive step first. Notice that

$$t_{v_1, v_2} t_{v_0, v_1} t_{v_1, v_2} = t_{v_0, v_1} t_{v_1, v_2} t_{v_0, v_1}$$

and then t_{v_0, v_1} commutes with all other terms. So for $n > 3$, we can commute t_{v_0, v_1} to the start and the end on both sides and then cancel, and so reduce to the case of one fewer term. So it suffices to prove the $n = 3$ case, where the two expressions we have to prove are equal can be written:

$$t_{x, y} t_{w, x} t_{v, w} t_{w, x} t_{x, y}$$

and

$$t_{x, y} t_{x, a} t_{x, b} t_{w, x} t_{v, w} t_{x, a} t_{x, a} t_{x, b} t_{x, y}$$

(where we have renamed v_0, v_1, v_2, v_3 as v, w, x, y to avoid subscripts. This differs from the earlier use of v and w , but this shouldn't cause confusion as we won't refer back to them.) As in the $n > 3$ case, we can rewrite and remove $t_{v, w}$ from the start and end of each expression and reduce this to showing equality of

$$t_{x, y} t_{w, x} t_{x, y}$$

and

$$t_{x, y} t_{x, a} t_{x, b} t_{w, x} t_{x, a} t_{x, b} t_{w, x}.$$

The second of these expressions is equal, by relation V, to

$$t_{x, y} t_{x, y} t_{x, a} t_{x, y} t_{x, b} t_{x, y} t_{w, x} t_{x, a} t_{x, b} t_{x, y} t_{x, y} t_{x, a} t_{x, y}$$

so it remains to show

$$t_{x, y} t_{x, y} t_{x, a} t_{x, y} t_{x, b} t_{x, y} t_{w, x} t_{x, a} t_{x, b} t_{x, y} t_{x, y} t_{x, a} = t_{w, x}.$$

Now $t_{x, b, y}$ commutes with $t_{w, x, a}$, so we can rearrange to needing to show

$$t_{x, y} t_{w, x, a} t_{x, y} = t_{x, y} t_{w, x} t_{x, y}.$$

This holds by relation VI. □

Now we just need to show that these symbols $t_{v, w}$ satisfy the relations that Bleak and Quick's $s_{v, w}$ satisfy (so that $s_{v, w} \mapsto t_{v, w}$ extends to a homomor-

phism). First of all we remark that if $\mathcal{L} = \{\ell_1, \ell_2, \dots, \ell_n\}$ is any leaf set, then the symbols $t_{\ell_i, \ell_{i+1}}$ satisfy the Coxeter relations of $\mathfrak{S}_{\mathcal{L}}$, presented by $(\ell_i \ \ell_{i+1})$. So the group presented by the $t_{v,w}$ contains a quotient of $\mathfrak{S}_{\mathcal{L}}$, and we can use relations we know from $\mathfrak{S}_{\mathcal{L}}$ to understand it.

We will prove Bleak and Quick's relations one at a time.

Lemma 4.3.4. *Each $t_{v,w}$ is of order 2.*

This is clear, because they are defined as conjugates of order 2 elements.

Lemma 4.3.5. *We have the expansion rule:*

$$t_{x,y} = t_{xa,ya} t_{xb,yb}.$$

Proof. Let $x = v_0, v_1, \dots, v_n = y$ be a sequence of adjacent leaves connecting x to y . We induct on n . For $n = 1$, we have:

$$\begin{aligned} t_{x,y} &= t_{xb,ya} t_{xa,xb} t_{ya,yb} t_{xb,ya} \\ &= t_{xb,ya} t_{xa,xb} t_{xb,ya} t_{xb,ya} t_{ya,yb} t_{xb,ya} \\ &= t_{xb,ya} t_{xa,xb} t_{xb,ya} \cdot t_{ya,yb} t_{xb,ya} t_{ya,yb} \\ &= t_{xa,ya} t_{xb,yb} \end{aligned}$$

For $n > 1$, we let \mathcal{L} be any leaf set containing $v_0a, v_0b, v_1a, v_1b, \dots, v_na, v_nb$. By definition:

$$t_{x,y} = t_{v_{n-1},v_n} \cdots t_{v_1,v_2} t_{v_0,v_1} t_{v_1,v_2} \cdots t_{v_{n-1},v_n}$$

which, by the $n = 1$ case, equals

$$\begin{aligned} &t_{v_{n-1}a,v_na} t_{v_{n-1}b,v_nb} \cdots t_{v_1a,v_2a} t_{v_1b,v_2b} \cdot t_{v_0a,v_1a} t_{v_0b,v_1b} \cdots \\ &\quad t_{v_1a,v_2a} t_{v_1b,v_2b} \cdots t_{v_{n-1}a,v_na} t_{v_{n-1}b,v_nb} \end{aligned}$$

All the $t_{x,y}$ in this expression have $x, y \in \mathcal{L}$. So by the remarks preceding the proof, we can evaluate this as a product of transpositions in $\mathfrak{S}_{\mathcal{L}}$. Any term $t_{v_ka, v_{k+1}a}$ commutes with any $t_{v_l b, v_{l+1} b}$. So the product in $\mathfrak{S}_{\mathcal{L}}$ gives $t_{v_0a, v_na} t_{v_0b, v_nb}$ as required. \square

Lemma 4.3.6. *$t_{x,y} t_{v,w} t_{x,y} = t_{(x \ y) \cdot v, (x \ y) \cdot w}$ whenever $(x \ y) \cdot v$ and $(x \ y) \cdot w$ are defined (that is, neither v nor w lies above x or y).*

Proof. This will be done in various cases.

- Case 1: any two of v, w, x, y are either incomparable or equal.
Choose a leaf set \mathcal{L} containing all of v, w, x, y ; let $\mathcal{L} = \{\ell_1, \ell_2, \dots, \ell_k\}$ written left-to-right. By the usual argument, the relations we seek hold for transpositions in $\mathfrak{S}_{\mathcal{L}}$ so they hold for the symbols $t_{v,w}$ also.
- Case 2: v lies below x , and w is incomparable to x and y .
Let $v = xz$. We want to establish:

$$t_{x,y}t_{xz,w}t_{x,y} = t_{yz,w}.$$

We will do this by induction on the length of z ; the case where z is the empty word has been done. We already know that $t_{x,y} = t_{xa,ya}t_{xb,yb}$. Now z either begins with a or b; without loss of generality write $z = az'$. Then $t_{xb,yb}$ commutes with $t_{xz,w}$ so it remains to establish:

$$t_{xa,ya}t_{xaz',w}t_{xa,ya} = t_{yaz',w}.$$

This is true by induction.

The other cases where exactly one of v, w lies below exactly one of x, y are similar. The remaining cases are more difficult and we will return to them after proving some auxiliary lemmas. \square

Lemma 4.3.7 (Generalization of relation V). *Let v, w be any incomparable vertices. Then:*

$$t_{v,w}t_{v,wa}t_{v,w} = t_{va,w} \text{ and } t_{v,w}t_{v,wb}t_{v,w} = t_{vb,w}.$$

Proof. First let v and w be neighbouring leaves. Then relation V tells us:

$$t_{va,w} := t_{vb,w}t_{va,wa}t_{vb,w} = t_{v,w}t_{v,wa}t_{v,w}$$

and

$$t_{v,wb} := t_{v,wa}t_{wa,wb}t_{v,wa} = t_{v,w}t_{v,wb}t_{v,w}.$$

These are the results we need to show.

Now suppose we want the result for v and some general x , to the right of v and not a neighbour of v in any leaf set. Choose w immediately to the right of v in some leaf set, and incomparable with x . We have:

$$t_{v,w}t_{v,wa}t_{v,w} = t_{va,w},$$

and we conjugate both sides by $t_{w,x} = t_{wa,xa}t_{wb,xb}$. By cases we've already done, this gives:

$$t_{v,x}t_{v,xa}t_{v,x} = t_{va,x}$$

as required; the other part of the result is similar. \square

Lemma 4.3.8 (Generalization of relation VI). *Let v, w, x be any incomparable vertices (in left-to-right order). Then:*

$$t_{w,x}t_{v,wa}t_{w,x} = t_{v,xa} \text{ and } t_{w,x}t_{v,wb}t_{w,x} = t_{v,xb}$$

Proof. Consider

$$t_{w,x}t_{v,xa}t_{w,x}.$$

Writing $t_{w,x}$ as $t_{wa,xa}t_{wb,xb}$, this evaluates to $t_{v,wa}$. This gives the result after conjugating both sides by $t_{w,x}$. \square

We return to the final cases we left earlier.

Proof. • Case 3: v and w both lie below x .

Writing $v = xv', w = xw'$, we need to show that

$$t_{x,y}t_{xv',xw'}t_{x,y} = t_{yv',yw'}.$$

We work inductively on $|v'|+|w'|$; notice that v' and w' must be non-empty in order for xv' and xw' to be incomparable.

Replace $t_{x,y}$ by $t_{xa,ya}t_{xb,yb}$. If both v' and w' begin with a (say), then the equation to prove becomes:

$$t_{xa,ya}t_{xv',xw'}t_{xa,ya} = t_{yv',yw'},$$

which can be considered in the same case, but with a removed from the front of v', w' . So we're done inductively. Otherwise we have that v' begins

with a and w' begins with b . So we can use Case 2 of Lemma 4.3.6 (which we've already proved), to give:

$$\begin{aligned} t_{xb,yb}t_{xa,ya}t_{xav''},xbw''t_{xa,ya}t_{xb,yb} &= t_{xb,yb}t_{yav''},xbw''t_{xb,yb} \\ &= t_{yav''},ybw'' \\ &= t_{yv'},yw' \end{aligned}$$

This completes the proof of this case.

- Case 4: v lies below x , and w lies below y .

Let $v = xv'$, $w = yw'$, so that we seek to show:

$$t_{x,y}t_{xv'},yw't_{x,y} = t_{yv'},xw'.$$

First assume that v' is empty (the case of w' being empty is similar). By Lemma 4.3.7, we already have the special cases:

$$\begin{aligned} t_{x,y}t_{x,ya}t_{x,y} &= t_{xa,y} \\ t_{x,y}t_{x,yb}t_{x,y} &= t_{xb,y} \\ t_{x,y}t_{xa,y}t_{x,y} &= t_{x,ya} \\ t_{x,y}t_{xb,y}t_{x,y} &= t_{x,yb} \end{aligned}$$

for any x and y . We will establish:

$$t_{x,y}t_{x,yw'}t_{x,y} = t_{xw'},y.$$

So now we consider $w' = aw''$ (for non-empty w'' ; the case of w' beginning with b is the same) and look at $t_{x,y}t_{x,yaw''}t_{x,y}$. Then we calculate:

$$\begin{aligned} t_{x,y}t_{x,yaw''}t_{x,y} &= t_{x,y}t_{x,ya}t_{x,ya}t_{x,yaw''}t_{x,ya}t_{x,ya}t_{x,y} \\ &= t_{x,y}t_{x,ya}t_{xw''},ya}t_{w,ya}t_{x,y} \text{ inductively on length of } w'' \\ &= t_{xa,y}t_{x,y}t_{xw''},ya}t_{x,y}t_{xa,y} \text{ by the first special case} \end{aligned}$$

Write $t_{x,y}$ as $t_{xa,ya}t_{xb,yb}$. If w'' begins with a , write $w'' = aw'''$. Then

$t_{xb,yb}$ commutes with $t_{xw'',ya}$ and we get:

$$\begin{aligned}
t_{xa,y}t_{x,y}t_{xw'',ya}t_{x,y}t_{xa,y} &= t_{xa,y}t_{xa,ya}t_{xaw''',ya}t_{xa,ya}t_{xa,by} \\
&= t_{xa,y}t_{xa,ya}w'''t_{xa,y} \text{ inductively} \\
&= t_{xaw''',y} \text{ inductively} \\
&= t_{xw',y} \text{ as required.}
\end{aligned}$$

Suppose now instead w'' begins with b, and write $w'' = bw'''$. We do the same calculations, but use the already-proved part 2 of Lemma 4.3.6 instead of induction:

$$\begin{aligned}
t_{xa,y}t_{x,y}t_{xw'',ya}t_{x,y}t_{xa,y} &= t_{xa,y}t_{xa,ya}t_{xb,yb}t_{xbw''',ya}t_{xb,yb}t_{xa,ya}t_{xa,y} \\
&= t_{xa,y}t_{xa,ya}t_{ya,ybw'''}t_{xa,ya}t_{xa,y} \\
&= t_{xa,y}t_{xa,ybw'''}t_{xa,y} \\
&= t_{xabw''',y} \text{ inductively} \\
&= t_{xw',y}.
\end{aligned}$$

This completes the proof in the case of v' empty (or w' empty). This will serve as the base case for an induction on $\min(|v'|, |w'|)$. If both v' and w' are non-empty and begin with the same digit (WLOG a), write:

$$\begin{aligned}
t_{x,y}t_{xv',yw'}t_{x,y} &= t_{xa,ya}t_{xb,yb}t_{xav''',yav'''}t_{xb,yb}t_{xa,ya} \\
&= t_{xa,ya}t_{xav''',yav'''}t_{xa,ya} \\
&= t_{xav''',yav'''} \text{ inductively} \\
&= t_{xw',yv'}.
\end{aligned}$$

Otherwise, write $v' = av''$, $w' = bw''$, say, and calculate:

$$\begin{aligned}
t_{x,y}t_{xv',yw'}t_{x,y} &= t_{xa,ya}t_{xb,yb}t_{xav''',ybw'''}t_{xb,yb}t_{xa,ya} \\
&= t_{xa,ya}t_{xav''',xbw'''}t_{xa,ya} \text{ by Case 2} \\
&= t_{xbw''',yav'''} \text{ by Case 2} \\
&= t_{xw',yv'}.
\end{aligned}$$

This completes the proof in all cases and we're done. □

Now we discuss the presentation we have created. Every relation in it is of the form $A^2 = 1$, $AB = BA$ or at worst, $ABA = CDC$ (where each letter is some symbol $t_{v,w}$), except for the expansion rule which is of the form $T = BACB$. This is nicely similar to the case of \mathfrak{S}_n , where every relation is $A^2 = 1$, $AB = BA$ or at worst $ABA = BAB$ (so a bit nicer than for V). If we're copying the deformation of \mathfrak{S}_n , we would keep all the relations for V but replace $A^2 = 1$ by a quadratic relation.

This does define an algebra, but there are a few problems. First, without anything for the algebra to act on, it's very unclear what the resulting algebra looks like. For example, we would want to know whether distinct elements of V remain distinct under the deformation of $\mathbf{k}V$, and this kind of question is difficult to answer for a general presentation. So even if rewriting this presentation does give us an interesting algebra, it's not clear what that algebra would be.

Later, we're going to discuss fitting together finite Hecke algebras, and we get a sense of this now. We ask whether $\mathcal{H}_{n,q}$ has an embedding into $\mathcal{H}_{2n,q}$ (corresponding to the embedding $\mathfrak{S}_n \hookrightarrow \mathfrak{S}_{2n}$ by expanding a leaf set). We'll later see evidence that these embeddings do not exist. So if we were to change the quadratic relations in this presentation of V , we would not get embeddings of Hecke algebras. Instead, the deformed algebra would have lots of extra relations added to the Hecke algebras at each level. So this doesn't yet give us a nice object.

Nevertheless, this seems an interesting presentation for V : it is as close as possible to a Coxeter presentation, where all the relations can be understood in terms of conjugating adjacent transpositions. So it seems worth establishing even though it doesn't do what we want at the moment.

4.3.2 Defining a generalized matrix group.

In this approach, we look at the group V as a subset of the Leavitt path algebra $L_K(\Gamma)$, where Γ is the directed graph with one vertex v and two loops at v , labelled a and b . Finite paths through Γ are then labelled by X^* and infinite paths by X^ω , for $X = \{a, b\}$. $L_K(\Gamma)$ then has the presentation:

$$L_K(\Gamma) = \langle S_a, S_b, S_a^*, S_b^* : S_a^* S_a = S_b^* S_b = S_a S_a^* + S_b S_b^* = 1 \rangle.$$

We choose K to be a finite field \mathbb{F} with the discrete topology. $L_{\mathbb{F}}(\Gamma)$ acts faithfully on the space $C(X^\omega, \mathbb{F})$ of continuous (that is, locally constant) \mathbb{F} -valued functions on X^ω , via:

$$S_a(f)(x) = \begin{cases} f(x') & \text{if } x = ax' \\ 0 & \text{if } x \text{ begins in } b. \end{cases}$$

and

$$S^*(f)(x) = f(ax).$$

Similar equations exist for the action of S_b, S_b^* . Alternatively, $L_{\mathbb{F}}(\Gamma)$ can act on the space $\mathbb{F}X^\omega$ whose basis is X^ω , and both these actions are useful.

Now we define a group of elements of $L_K(\Gamma)$ which we will work with as a general linear group.

Definition 4.3.9. *Let \mathcal{L} be a leaf set in X^* , where $\mathcal{L} = \{v_1, v_2, \dots, v_n\}$. We say that an \mathcal{L} -matrix is an element M of $L_{\mathbb{F}}(G)$ of the form:*

$$M = \sum_{i,j=1}^N \lambda_{i,j} S_{v_i} S_{v_j}^*,$$

for scalars $\lambda_{i,j} \in \mathbb{F}$. We write $M(\mathcal{L}, \mathbb{F})$ for the set all \mathcal{L} -matrices, and write $GL(\mathcal{L}, \mathbb{F})$ for the set of all invertible \mathcal{L} -matrices.

Notice that since $S_{v_j}^* S_{v_i} = \delta_{i,j}$, the algebra $M(\mathcal{L}, \mathbb{F})$ is isomorphic to the algebra $M_n(\mathbb{F})$ of $n \times n$ matrices over \mathbb{F} .

Definition 4.3.10. *The Leavitt general linear group for the graph Γ , $GL(\Gamma, \mathbb{F})$, is defined to be the group generated by all groups $GL(\mathcal{L}, \mathbb{F})$ as the leaf set \mathcal{L} varies (with the multiplication equal to the multiplication in $L_{\mathbb{F}}(G)$).*

Notice that not every element of $GL(\Gamma, \mathbb{F})$ is an \mathcal{L} -matrix. We only ask that the generators are \mathcal{L} -matrices. General elements are of the form $\sum_{i=1}^n \lambda_i S_{v_i} S_{w_i}^*$, with no conditions on the words $v_i, w_i \in \Gamma^*$. However, it is true in general that every element of $GL(\Gamma, \mathbb{F})$ can be written in the form:

$$M = \sum_{i=1}^n \left(\sum_{j=1}^{m_i} \lambda_{i,j} S_{w_{i,j}} \right) S_{v_i}^*,$$

where the v_i form a leaf set. This can be done by expanding, replacing monomial

$S_w S_v^*$ with $S_w(S_a S_a^* + S_b S_b^*) S_v^*$. Similarly, we could expand and write:

$$M = \sum_{j=1}^m S_{w_j} \left(\sum_{i=1}^{n_j} \lambda_{i,j} S_{v_{i,j}}^* \right),$$

where the w_j form a leaf set. What can't be guaranteed is that both the v_i and the w_j form a leaf set at the same time.

We remark that the group $GL(\Gamma, \mathbb{F})$ has some similar structure to V . For example, if \mathcal{L} is a leaf set, recall there is a leaf set \mathcal{L}^+ formed by expanding every vertex of \mathcal{L} . If M is an \mathcal{L} -matrix, then we can rewrite M to be an \mathcal{L}^+ -matrix by replacing each monomial $S_v S_w^*$ with $S_v(S_a S_a^* + S_b S_b^*) S_w^*$. This gives an embedding of $GL(\mathcal{L}, \mathbb{F})$ into $GL(\mathcal{L}^+, \mathbb{F})$, which in terms of matrices can be written $M \mapsto M \otimes I_2$ (a Kronecker product with a 2×2 identity matrix). We will see this kind of map more later.

We now define a Borel subgroup.

Definition 4.3.11. *Let $M \in GL(\Gamma, \mathbb{F})$, and write M in the form:*

$$M = \sum_{i=1}^n \left(\sum_{j=1}^{m=m(i)} \lambda_{i,j} S_{w_{i,j}} \right) S_{v_i}^*,$$

where the v_i form a leaf set. We say that M is upper-triangular if $w_{i,1} = v_i$ for all i , and $w_{i,j}$ lies strictly to the left of v_i for all $i > 1$. We define the Borel subgroup $B(\Gamma, \mathbb{F})$ to be the subgroup of $GL(\Gamma, \mathbb{F})$ consisting of upper triangular elements.

It's perhaps easiest to think about $B(\Gamma, \mathbb{F})$ in terms of the action on an end $x \in X^\omega$: if M is upper-triangular, then Mx is a sum

$$Mx = \lambda x + \lambda_1 x_1 + \dots + \lambda_k x_k,$$

where the ends x_i lie strictly to the left of x . In other words, $B(\Gamma, \mathbb{F})$ acts on $\mathbb{F}X^\omega$ in a manner that preserves rightmost elements of linear combinations of ends. This description makes it clear that $B(\Gamma, \mathbb{F})$ is a group.

We point out that there is more than one sensible way to define a general linear group and a Borel subgroup in the Leavitt path algebra. Alternatives would be to define $GL(\Gamma, \mathbb{F})$ as the set of all invertible elements of $L_{\mathbb{F}}(G)$, or to define $B(\Gamma, \mathbb{F})$ as the subgroup of $GL(\Gamma, \mathbb{F})$ generated by upper-triangular \mathcal{L} -matrices. It doesn't seem clear whether these definitions would give the same

groups. We've chosen the definitions we have to make the theory work as much as possible, and won't pursue these questions, but it seems interesting to study these groups in their own right.

We now define flags, one of which the Borel subgroup fixes.

Definition 4.3.12. *The standard flag for V is the following function F_{st} from X^ω plus a single point x_0 to subspaces of $\mathbb{F}X^\omega$: for $x \in X^*$, $F_{st}(x)$ is the space $1_{[0,x]}$ of all functions that vanish to the right of x , and $F_{st}(x_0) = \{0\}$ (so think of x_0 as being to the left of all of X^ω). Since $GL(\Gamma, \mathbb{F})$ acts on $\mathbb{F}X^\omega$, it acts on functions from $X^\omega \cup \{x_0\}$ to subspaces of $\mathbb{F}X^\omega$, and we define a flag to be a point in the $GL(\Gamma, \mathbb{F})$ -orbit of F_{st} .*

We should check that the stabilizer of F_{st} is the Borel subgroup. Indeed, let $x \in X^\omega$, and let $M \in GL(\Gamma, \mathbb{F})$. Write $Mx = \sum_{i=1}^n \lambda_i x_i$ where $\lambda_i \in \mathbb{F}$ and $x_i \in X^\omega$, with x_n rightmost. Notice that for $y \in X^\omega$, $1_x \in F_{st}(y)$ if and only if x does not lie to the right of y , and $Mx \in F_{st}(y)$ if and only if 1_{x_n} does not lie to the right of y . If M is in the stabilizer of F_{st} , these conditions must be the same, and so we must have $x_n = x$. Since this holds for all x , we see that the stabilizer of F_{st} is contained in the Borel subgroup. It's easy to see that the Borel subgroup stabilizes F_{st} so we have equality.

Now we relate $B(\Gamma, \mathbb{F})$ and the group V . For $x \in X^\omega$, we will continue writing $[0, x]$ for the set of ends to the left of or equal to x . If S is a set of ends, we write $1_S \subset \mathbb{F}X^\omega$ for the subspace of $\mathbb{F}X^\omega$ of functions supported on S .

Proposition 4.3.13. *Elements of Thompson's group V lie in distinct double cosets for $B(\Gamma, \mathbb{F})$ in $GL(\Gamma, \mathbb{F})$.*

Proof. Write B for $B(\Gamma, \mathbb{F})$. Suppose that $X_1, X_2 \in V$ with $BX_1B = BX_2B$. It follows that $X_1(F_{st}) = bX_2(F_{st})$ for some $b \in B$. Consider $bX_2(F_{st})(x)$ for $x \in X^\omega$. $X_2(F_{st})(x)$ is the image of $1_{[0,x]}$ under X_2 , which is $1_{X_2([0,x])}$. This set contains $X_2(x)$ (viewed as an element of $\mathbb{F}X^\omega$), and so $bX_2(F_{st})(x)$ contains an element of $\mathbb{F}X^\omega$ whose rightmost end is $X_2(x)$. On the other hand, for y lying left of x , $X_2(F_{st})(y) = 1_{X_2([0,y])}$ contains no function supported at $X_2(x)$, so $bX_2(F_{st})(y)$ cannot contain an element whose rightmost end is $X_2(x)$. By the same reasoning for X_1 , we see that $X_1(F_{st})(x)$ contains elements of $\mathbb{F}X^\omega$ whose rightmost end is $X_2(x)$, whilst $X_1(F_{st})(y)$ does not contain them for y left of x . This is only possible if $X_1(x) = X_2(x)$, and so $X_1 = X_2$. \square

To make a Hecke algebra from this setup, we would want to look at endomorphisms of the $GL(\Gamma, \mathbb{F})$ -module of flags. However, one can check that B is

not a commensurated subgroup, so we cannot define this easily. It also doesn't seem that there's a topology to put on flags that will help here. So we leave this idea for now.

4.3.3 Fitting together finite Hecke algebras

The last ideas will be dealt with more briefly. We won't be creating much new theory here, and we've already identified most of the obstructions to the existence of a Hecke algebra.

Here we consider trying to fit together finite Hecke algebras $\mathcal{H}_{n,q}$. We would like an algebra formed from deforming copies of \mathfrak{S}_n inside V into Hecke algebras $\mathcal{H}_{n,q}$, preserving the embeddings between them. In particular, we would like to deform the embedding $\mathfrak{S}_2 \hookrightarrow \mathfrak{S}_4$ into an embedding $\mathcal{H}_{2,q} \hookrightarrow \mathcal{H}_{4,r}$. In the symmetric groups (permuting the layers X^1 and X^2), this embedding is

$$(a \ b) \mapsto (aa \ ba)(ab \ bb) = (ab \ ba)(aa \ ab)(ba \ bb)(ab \ ba),$$

where we have written the element of \mathfrak{S}_4 as a product of Coxeter generators. Let the generators of $\mathcal{H}_{n,q}$ be $\Sigma_{n,1}, \dots, \Sigma_{n,n-1}$. To deform this, we would ideally want an embedding mapping

$$\Sigma_{2,1} \mapsto \Sigma_{4,2}\Sigma_{4,1}\Sigma_{4,3}\Sigma_{4,2}.$$

But it is easy to verify that $\Sigma_{4,2}\Sigma_{4,1}\Sigma_{4,3}\Sigma_{4,2}$ does not satisfy any quadratic relations and so cannot be the image of $\Sigma_{2,1}$ under an algebra embedding (in fact, its minimal polynomial is of degree 7). One could try mapping $\Sigma_{2,1}$ to a linear combination of monomials of the Hecke algebra, of which the longest is $\Sigma_{4,2}\Sigma_{4,1}\Sigma_{4,3}\Sigma_{4,2}$. So we're looking for elements of $\mathcal{H}_{4,r}$ satisfying a particular quadratic relation. I haven't been able to find such an element with extensive calculation: they may exist, but there doesn't seem to be any natural reason for this embedding to exist.

4.3.4 Taking a quotient of a braided Thompson group algebra

We will discuss this idea more later, as it turns out to be the most useful. Braided Thompson groups BV were constructed by Matthew Brin in [15] and [16]. These groups contain a copy of the braid group B_n on n strands whenever

V contains a copy of \mathfrak{S}_n . In particular, BV has a presentation in terms of generators σ , to which relations $\sigma^2 = 1$ can be added to give a presentation of Thompson's group V . Our aim would be to add different quadratic relations $\sigma^2 = (q-1)\sigma + q$ to $\mathbb{F}BV$. But we now run into the same problem as the previous section. Indeed, performing this quotient would turn $B_2 \hookrightarrow B_4$ into an embedding $\mathcal{H}_{2,q} \hookrightarrow \mathcal{H}_{4,q}$. But this would require $\Sigma_{4,2}\Sigma_{4,1}\Sigma_{4,3}\Sigma_{4,2}$ to satisfy a quadratic relation in the algebra $\mathcal{H}_{4,q}$, which it does not.

4.3.5 Consequences and concessions

The result of this work is that we have learnt some things that we are going to have to sacrifice if we want to define a Hecke algebra for V . In particular, we are not going to be able to have quadratic relations. Instead, the best we can do is to define an algebra of operators on a suitable flag space \mathcal{F} and impose *local* quadratic relations: a deformed transposition Σ of the Hecke algebra \mathcal{H} will have $\Sigma^2 \cdot f = (q-1)\Sigma \cdot f + qf$ for some $f \in \mathcal{F}$, but not all of them. Also, we saw that our attempts to write down a useful presentation of V still had too many transpositions in the generating set to give a useful presentation. Eventually, we will write down a smaller set of transpositions which don't generate V , but do generate when F is added to the generating set (recall F is the smallest of Thompson's groups, whose elements are given by order-preserving bijections between leaf sets). Our plan will be to deform the transpositions but not deform F , and this will be more achievable. This means that we end up giving an action of the braided Thompson group. Subgroups of the braided Thompson group will act on subspaces via a quotient to $\mathcal{H}_{n,q}$. This seems to be the best that can be done.

4.4 The braided Thompson group, partial actions, and the Hecke algebra

In this section we introduce the braided Thompson group, and describe how to take a quotient of its group algebra to form an object which we will call a Hecke algebra for V . In fact the construction is a bit more complicated than this overview suggests: we won't just take a quotient of the group algebra, but will have to equip the braid group algebra with a partial action, and take a quotient of that. The result of this is that we form an algebra with partial

action, reflecting the binary tree structure of V . We make this precise later.

In this section, we will draw elements of the braid group B_n as diagrams of braided strands, and will represent elements of the symmetric group with the same sort of figures. An example, in B_3 , is shown in Figure 4.1. This means that we have to view our elements of the symmetric group as permuting ‘objects in places’: so for example, the transposition $(2\ 3)$ will act by swapping the second and third objects in any permutation of $\{1, 2, \dots, n\}$ rather than interchanging the numbers 2 and 3. The element of B_3 in Figure 4.1 can thus be written as a product of generators as $\Sigma_1\Sigma_2\Sigma_1\Sigma_2$ (since when reading from top to bottom, it swaps the second pair of strings, then the first, then the second, then the first again, each time moving the left strand over the right strand). This should establish our conventions on braid groups.

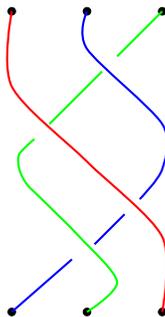


Figure 4.1: Example of a braid group element

4.4.1 The braided Thompson group

A braided version of Thompson’s group V was described in [15], and the construction generalizes in an obvious manner to give braided versions of Higman-Thompson groups $V_{n,d}$, which we denote $B_{V,n,d}$ or just B_V . We describe it now. As with Higman-Thompson groups, elements of $B_{V,n,d}$ can be understood as permutations of an end space DX^ω which arise from a bijection between leaf sets. However, in the braided group we care not just about the permutation (or bijection between leaf sets), but the braiding that arises when one imagines moving the points in space to carry out the bijection. This is easiest to explain with pictures, similar to Figure 3.1.1. So we show a typical element of B_V below (for $V = V_{2,1}$, the usual Thompson group):

The black lines in Figure 4.2 show two full subtrees, whose leaf sets are put

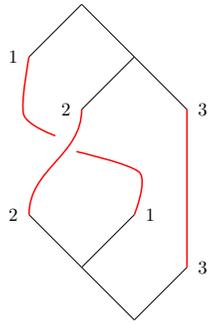


Figure 4.2: An element X of B_V

in bijection by the braided red lines.

These pictures are multiplied using expansions, in the normal fashion for Thompson groups. In a simple expansion, we replace a pair of leaves and the connecting strand between them with all leaves immediately below that pair, joined by parallel strands with the same braiding. We demonstrate this with a simple expansion of X as above, expanding at the leaf labelled 1. The result is shown in Figure 4.3.

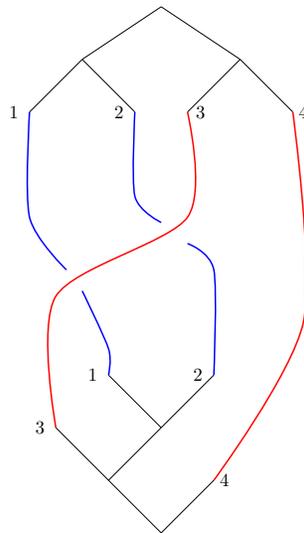


Figure 4.3: An expanded diagram for X

To multiply two elements X_1 and X_2 , we take expansions such that the domain leaf set of X_2 is the source leaf set of X_1 , then cancel the common leaf

sets and multiply the braids. This is explained more fully in [15] Section 1.2; we give an example of calculating X^2 (where X is as shown in Figure 4.2). The top and bottom of Figure 4.4 both consist of an expansion of X , with the expanded lines shown in blue. The product is shown on the right, which is found by replacing the identical subtrees with the dotted green connecting lines.

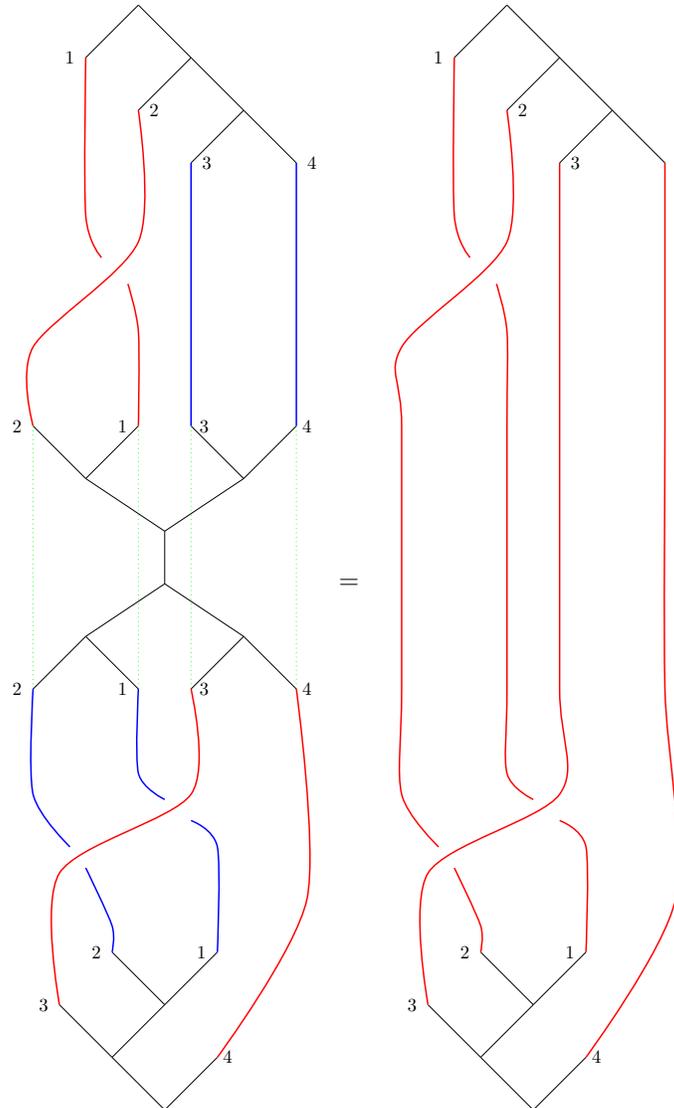


Figure 4.4: The product X^2

We make some easy remarks about B_V . First notice that $F = F_{n,d}$ embeds

into B_V as the subgroup with no braiding. The group V is not a subgroup of B_V , but it is a quotient (via the homomorphism that forgets the braiding, and just keeps the permutation).

We will build a Hecke algebra for V out of B_V , but equipped with additional information about its partial action on leaf sets. First of all we find a useful factorization of V and of B_V .

4.4.2 The homogeneous subgroup

Let $V = V_{n,d}$ be a Higman-Thompson group. We define an important subgroup of V and of B_V . Fix sets D, X of size d, n respectively (for the Higman-Thompson group $V = V_{n,d}$ and its braided version B_V).

Definition 4.4.1. *Let $V = V_{n,d}$ be a Higman-Thompson group. We define the subgroup $\mathfrak{S}_{\text{hom}} = \mathfrak{S}_{\text{hom},n,d}$ of V to be the subgroup generated by all $\bar{\sigma} \in V$ where σ is a depth-preserving bijection between leaf sets. Similarly, define $B_{\text{hom}} = B_{\text{hom},n,d}$ to be the subgroup of B to be the subgroup generated by all elements that can be represented by a braided depth-preserving bijection between leaf sets.*

The notation $\mathfrak{S}_{\text{hom}}$ is intended to suggest that these are homogeneous elements: they preserve depth in their partial action on DX^* . We give a structure theorem for $\mathfrak{S}_{\text{hom}}$.

Proposition 4.4.2. *$\mathfrak{S}_{\text{hom}}$ is isomorphic to a direct limit of finite symmetric groups,*

$$\mathfrak{S}_d \hookrightarrow \mathfrak{S}_{dn} \hookrightarrow \mathfrak{S}_{dn^2} \hookrightarrow \mathfrak{S}_{dn^3} \hookrightarrow \dots$$

The embeddings are as follows. If σ is an element of \mathfrak{S}_m (for $m = dn^k$), and $1 \leq j \leq nm$, then write $j = qm + r$ for $0 \leq q \leq n-1$ and $1 \leq r \leq m$. Then the image of σ in \mathfrak{S}_{nm} sends j to $(\sigma(q+1) - 1)n + r$.

Similarly, B_{hom} is a direct limit of braid groups on finitely many strands,

$$B_d \hookrightarrow B_{dn} \hookrightarrow B_{dn^2} \hookrightarrow B_{dn^3} \hookrightarrow \dots,$$

by a series of embeddings compatible with the embeddings for $\mathfrak{S}_{\text{hom}}$.

In particular, for every $\sigma \in \mathfrak{S}_{\text{hom}}$ (or $\Sigma \in B_{\text{hom}}$), σ can be written as a permutation of the leaf set DX^k , for each sufficiently large k .

Proof. The proof is the same for both groups. Let $X \in \mathfrak{S}_{\text{hom}}$ (or B_{hom}). Then by expanding, we can assume that the domain leaf set of X is some DX^k . Since X is level-preserving, the range leaf set of X must also be DX^k . Then X is just specified by a (braided) permutation of the dn^k vertices of DX^k , so lies in \mathfrak{S}_{dn^k} (or B_{dn^k}). It is then easy to check that expanding DX^k to DX^{k+1} gives the claimed embedding from \mathfrak{S}_{dn^k} to $\mathfrak{S}_{dn^{k+1}}$. Finally, the constructions for B_{hom} and $\mathfrak{S}_{\text{hom}}$ were parallel, so clearly the two direct limit structures are compatible with the quotient $B_V \rightarrow V$. \square

Corollary 4.4.3. *For $m = dn^k$ (each $k \in \mathbb{N}_0$), let the group \mathfrak{S}_m have Coxeter generators $\sigma_{k,1}, \sigma_{k,2}, \dots, \sigma_{k,m-1}$ (so that $\sigma_{k,i}$ is the transposition $(i \ i+1)$.) Then $\mathfrak{S}_{\text{hom}}$ is generated by the various $\sigma_{k,i}$, and has a presentation with these generators and the relations:*

1. $\sigma_{k,i}^2 = 1$, where $k \in \mathbb{N}_0$, $m = dn^k$ and $1 \leq i \leq m-1$.
2. $\sigma_{k,i}\sigma_{k,i+1}\sigma_{k,i} = \sigma_{k,i+1}\sigma_{k,i}\sigma_{k,i+1}$, whenever $k \in \mathbb{N}_0$, $m = dn^k$ and $1 \leq i \leq m-2$
3. $\sigma_{k,i}$ and $\sigma_{k,j}$ commute whenever $|i-j| > 1$.
4. For each $k \in \mathbb{N}_0$, $m = dn^k$ and $1 \leq i \leq m-1$, then

$$\sigma_{k,i} = c_{k+1,ni}c_{k+1,ni+1} \cdots c_{k+1,ni+n-1},$$

where

$$c_{k+1,y} = \sigma_{k+1,y}\sigma_{k+1,y-1} \cdots \sigma_{k+1,y-n+1}.$$

If we let B_m have standard generators $\Sigma_{k,1}, \Sigma_{k,2}, \dots, \Sigma_{k,m-1}$, then the various $\Sigma_{k,i}$ generate B_{hom} , which has a presentation with these generators satisfying relations 2-4.

Proof. This follows from the description of the groups $\mathfrak{S}_{\text{hom}}$ and B_{hom} as a direct limit. For fixed m , the first three relations give a presentation of \mathfrak{S}_m (or B_{hom}), and the fourth relation gives the embedding of \mathfrak{S}_m into \mathfrak{S}_{mn} (or B_m into B_{mn}). We just need to verify that $c_{k+1,ni} \cdots c_{k+1,ni+n-1}$ describes the same permutation as the embedding of Proposition 4.4.2. To do this, we draw a picture for the braid group in Figure 4.5. The result clearly swaps the leaves $(i-1)n+1, \dots, in$ (which lie below the i th vertex of level k) with $in+1, \dots, in+n$ (which lie below the $i+1$ th). This means it is the correct expansion of a transposition.

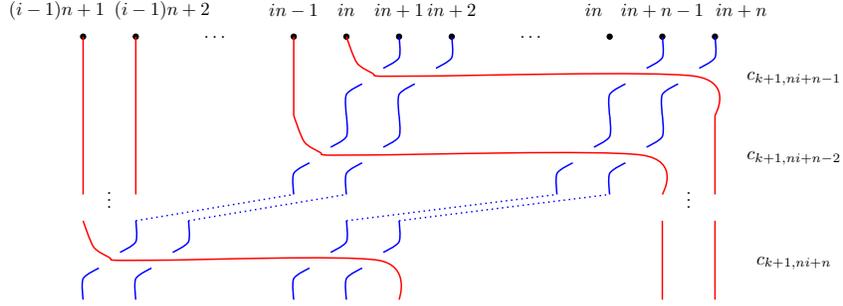


Figure 4.5: Expanding a transposition in the braid group

□

We also state a result about the expansion of $\sigma_{k,i}$ to higher levels that will be useful later.

Proposition 4.4.4. *Let $\sigma = \sigma_{k,i_1} \sigma_{k,i_2} \dots \sigma_{k,i_r}$ be a reduced word in the Coxeter generators of \mathfrak{S}_{dn^k} . Fix $l > k$, and use relation 4 of Corollary 4.4.3 repeatedly, to write σ as a product of terms $\sigma_{l,j}$. Then the resulting expression*

$$\sigma = \sigma_{l,j_1} \sigma_{l,j_2} \dots \sigma_{l,j_s}$$

is reduced, as a word in the Coxeter generators of \mathfrak{S}_{dn^l} .

Proof. We induct on $l - k$. Suppose first that $l = k + 1$. First we do the case when $r = 1$, so that $\sigma = \sigma_{k,i}$, and we need to show the expression

$$C_{k+1,ni} C_{k+1,ni+1} \dots C_{k+1,ni+n-1}$$

is reduced, when each $c_{k+1,j}$ is written as a product of permutations. Let the elements of DX^k , left-to-right, be $v_{k,1}, v_{k,2}, \dots, v_{k,m}$ and let the elements of DX^{k+1} be $v_{k+1,1}, \dots, v_{k+1,mn}$. Consider the action of σ on the sequence of $2n$ leaves

$$v_{k+1,(i-1)n+1}, \dots, v_{k+1,ni}, v_{k+1,ni+1}, \dots, v_{k+1,ni+n-1}.$$

We drew out a picture for this expansion in Figure 4.5. In that picture, one can easily see that no two strands cross twice. This is a necessary and sufficient condition for a word in the generators of $\mathfrak{S}_{dn^{k+1}}$ to be reduced. Moreover, we notice that every permutation swaps a red strand with a blue strand in Figure

4.5 - that is, it swaps a leaf below the i th vertex at level k with one below the $i + 1$ st. This will be useful later.

Now suppose $l = k + 1$ but that σ is arbitrary. We argue that after expanding, $\sigma_{l,j_1} \dots \sigma_{l,j_s}$ does not swap any pair of leaves twice, noting that $\sigma_{k,i_1} \dots \sigma_{k,i_s}$ has this property in its action on level k vertices. Consider two leaves $v_{l,a}$ and $v_{l,b}$ that lie below leaves $v_{k,a'}$ and $v_{k,b'}$ at level k . Our discussion of the $r = 1$ case shows that $v_{l,a}$ and $v_{l,b}$ are swapped by some transposition (which is unique) in the expansion of $\sigma_{k,i}$ if and only if $\sigma_{k,i}$ swaps $v_{k,a'}$ and $v_{k,b'}$. There is at most one such $\sigma_{k,i}$ so we are done.

The case of general l now results by an easy induction on $l - k$, since we know σ expands to a reduced word in $\mathfrak{S}_{dn^{k+1}}$. \square

Remark 4.4.5. *One could generalize this theorem as follows: let $\sigma \in \mathfrak{S}_n$. Write σ as a reduced word in the Coxeter generators, and draw this reduced word as a diagram of crossing strands (as in Figure 4.1, but not worrying about the braiding). Suppose we replace the i th strand with k_i parallel strands, for each $1 \leq i \leq n$. Then the resulting diagram corresponds to a reduced word for an element of $\mathfrak{S}_{k_1+k_2+\dots+k_n}$. This defines a function from $\mathfrak{S}_n \rightarrow \mathfrak{S}_{k_1+\dots+k_n}$ which is not a group homomorphism in general, but is still useful on occasion. It is proved by observing that if no two strands cross twice in the original diagram, the same is true for the expanded diagram.*

We now relate V to F and $\mathfrak{S}_{\text{hom}}$. This will be the critical step in our construction of a Hecke algebra. The proposition relating them is probably well-known, although it doesn't seem to be included in the standard sources on V . We reprove it here anyway, as the proof is instructive.

Proposition 4.4.6. *V is generated by its subgroups F and $\mathfrak{S}_{\text{hom}}$. Moreover, every element X of V can be uniquely written as $X = f\sigma$, for $f \in F$ and $\sigma \in \mathfrak{S}_{\text{hom}}$. Similarly every element of B_V can be uniquely written as $X = f\Sigma$ for $f \in F$ and $\Sigma \in B_{\text{hom}}$.*

Proof. The proof is the same for both groups. Given $X \in V$ (or B), we can choose trees representing X where the domain leaf set is DX^k , for some k . There exists a unique permutation $\sigma \in \mathfrak{S}_{dn^k}$ which permutes the elements of DX^k into the same order as in the range leaf set of X (and with the same braiding). Then $X = f\sigma$ for some order-preserving f , and we're done. This is illustrated in Figure 4.6, for an element of $V = V_{1,2}$, the usual Thompson's group V .

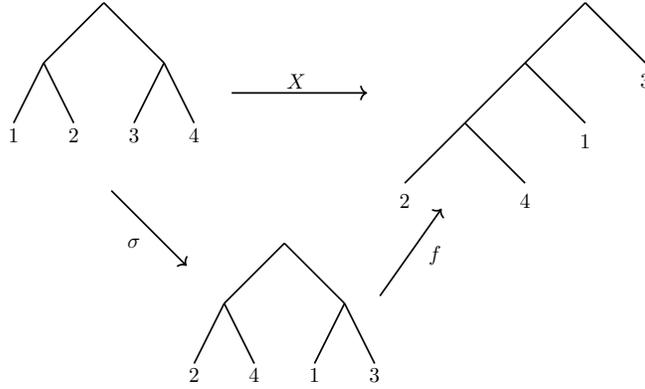


Figure 4.6: Writing $X = f\sigma$

Uniqueness comes from the fact that F and $\mathfrak{S}_{\text{hom}}$ (or B_{hom}) have trivial intersection: indeed an element in their intersection can be represented as an order-preserving (unbraided) permutation of DX^k for some k , which must be the identity. □

Corollary 4.4.7. *Every element X of V (or B) can be written uniquely as $X = \sigma f$ for $f \in F$ and $\sigma \in \mathfrak{S}_{\text{hom}}$ (or B_{hom}).*

Proof. Apply the previous corollary to X^{-1} . □

We record one more similar fact here too. Suppose that L, L' are leaf sets of size dn^k , so that there exist $f, f' \in F$ which map DX^k to L, L' respectively. Let $X \in V$ (or B_V) be represented by a (braided) bijection between L and L' . Then $(f')^{-1}Xf$ defines a (braided) bijection from DX^k to itself, which lies in $\mathfrak{S}_{dn^k} \subset \mathfrak{S}_{\text{hom}}$ (or $B_{dn^k} \subset B_{\text{hom}}$). Put another way:

Proposition 4.4.8. *Let $X \in V$ (or B_V) be defined by a (braided) bijection between two leaf sets of size dn^k , which we write as $f(DX^k)$ and $f'(DX^k)$ for $f, f' \in F$. Then $X = f'\phi f^{-1}$ for $\phi \in \mathfrak{S}_{\text{hom}}$ (or B_{hom}).*

The advantage of this over Proposition 4.4.6 is that no expansion is necessary, but the resulting expression isn't unique. We now use these factorizations to make a useful presentation for V .

Proposition 4.4.9. *V has a presentation with generators $\sigma_{m,i}$ (for each $m = dn^k, k \in \mathbb{N}_0$ and $1 \leq i \leq m - 1$) and X_f (for each $f \in F$), with relations:*

- $\sigma_{k,i}^2 = 1$, all $m = dn^k$ and $1 \leq i \leq m - 1$.
- $\sigma_{k,i}\sigma_{k,i+1}\sigma_{k,i} = \sigma_{k,i+1}\sigma_{k,i}\sigma_{k,i+1}$, each $m = dn^k$ and $1 \leq i \leq m - 2$
- $\sigma_{k,i}$ and $\sigma_{k,j}$ commute whenever $|i - j| > 1$.
- For each $m = dn^k$ and $1 \leq i \leq m - 1$, then

$$\sigma_{k,i} = c_{k+1,ni}c_{k+1,ni+1} \cdots c_{k+1,ni+n-1},$$

where

$$c_{k+1,y} = \sigma_{k+1,y}\sigma_{k+1,y-1} \cdots \sigma_{k+1,y-n+1}.$$

- Whenever $f, g \in F$, then $X_f X_g = X_{fg}$.
- Whenever $\sigma_{k,i} f = g \sigma$ in V , for $\sigma \in \mathfrak{S}_{\text{hom}}$ and $f, g \in F$, then we have a relation:

$$\sigma_{k,i} X_f = X_g \sigma_{k_1,i_1} \sigma_{k_2,i_2} \cdots \sigma_{k_n,i_n},$$

where $\sigma = \sigma_{k_1,i_1} \sigma_{k_2,i_2} \cdots \sigma_{k_n,i_n}$ as a product of generators of $\mathfrak{S}_{\text{hom}}$.

Dropping the relations $\sigma_{k,i}^2 = 1$ gives a presentation for B_V .

Proof. For V , let this presentation define a group V_1 . We've seen in Corollary 4.4.3 that the first four relations give a presentation for $\mathfrak{S}_{\text{hom}}$. All the relations hold in V for $\sigma_{k,i}$ and f , so there's a homomorphism from V_1 to V sending $\sigma_{k,i}$ to $\sigma_{k,i}$ and X_f to f , which is surjective because its image contains the subgroups F and $\mathfrak{S}_{\text{hom}}$ which generate V . It remains to prove this homomorphism is injective. To do that, it's enough to see that every element of V_1 can be written in the form $X_f \sigma$ (for σ in the group generated by the $\sigma_{k,i}$) because we know that all these expressions represent different elements of V . It's enough to prove this for elements of the form σX_f (since a general element of V_1 can be written as a product $\sigma_1 f_1 \cdots \sigma_p f_p$ for σ_i in the subgroup generated by the $\sigma_{m,i}$, and $f_i \in F$. Then we can put it in the desired form by moving all σ_i to the right and all the f_j to the left). This is achieved by writing σ as a product of generators $\sigma_{m,i}$ and using the final relation repeatedly.

The argument for B_V is identical except that we use the fact that the second to fourth relations give a presentation for B_{hom} . \square

We will use this presentation of B_V to define a Hecke algebra, but first we will add to the group algebra $\mathbf{k}B_V$ the information of its partial action on DX^* .

4.4.3 Groups and algebras with partial actions

Here we make some general definitions about what a group with partial action should mean. It's important to do this in generality because we will define the general Hecke algebra as the quotient of a braid group algebra with a partial action, and so we need to understand these things as algebraic objects.

Definition 4.4.10. *Let \mathcal{C} be a small category whose hom-sets are R -modules for some (commutative, unital) ring R . We say that a group with partial \mathcal{C} -action is a group G equipped with function $S : g \mapsto S(g)$, where $S(g)$ is a subset of the morphisms of \mathcal{C} , such that:*

- *The domain and range maps are injective on $S(g)$ for each $g \in G$. In other words, $S(g)$ is a set of morphisms no two of which have the same domain or the same range.*
- *$S(1)$ is the set of all identity maps in \mathcal{C} .*
- *If $\phi : V_1 \rightarrow V_2 \in S(g)$, and $\psi : V_2 \rightarrow V_3 \in S(h)$, then $\psi\phi : V_1 \rightarrow V_3 \in S(hg)$.*
- *Each morphism of $S(g)$ is invertible, and $S(g^{-1}) = S(g)^{-1}$, the pointwise inverse of the set $S(g)$.*

If some morphism of $S(g)$ has domain V_1 , then we write $g(V_1) = \phi$, where ϕ is the unique such morphism.

Definition 4.4.11. *For \mathcal{C} as above, we say that an R -algebra A with partial \mathcal{C} -action is a unital R -algebra A with function $S : x \mapsto S(x)$, a subset of the morphisms of \mathcal{C} , such that:*

- *The domain and range maps are injective on $S(x)$ for each $x \in G$.*
- *$S(1)$ is the set of all identity maps in \mathcal{C} .*
- *If $\phi : V_1 \rightarrow V_2 \in S(x)$, then $\lambda\phi : V_1 \rightarrow V_2 \in S(\lambda x)$ for each $\lambda \in R$.*
- *If $\phi : V_1 \rightarrow V_2 \in S(x)$, and $\psi : V_2 \rightarrow V_3 \in S(y)$, then $\psi\phi : V_1 \rightarrow V_3 \in S(yx)$.*
- *If $g \in A$ is invertible, then each morphism of $S(g)$ is invertible, and $S(g^{-1}) = S(g)^{-1}$, the pointwise inverse of the set $S(g)$.*

- If $x, y \in A$, and $\phi : V_1 \rightarrow V_2$ lies in $S(x)$ and $\psi : V_1 \rightarrow V_2$ lies in $S(y)$, then $S(x + y)$ contains $\phi + \psi$.

The most common examples of R -algebras with partial \mathcal{C} -action will come from linearly extending a group with partial \mathcal{C} -action into its group algebra.

All the examples we work with will be closely related to Thompson's group and have $\text{Ob}(\mathcal{C})$ some collection of leaf sets. We illustrate with some important examples.

Example 1: (A trivial example, from a group action). Let G be a group acting on a set X . Let \mathcal{C} be the category whose set of objects is X and where there is a 1-dimensional R -space of morphisms $m_{x,y} \cdot R$ from x to y for each $x, y \in X$, with $m_{y,z}m_{x,y} = m_{x,z}$. Define $S(g)$ by:

$$S(g) = \{m_{x,g \cdot x} : x \in X\}.$$

Then G becomes a group with partial \mathcal{C} -action.

Example 2: Let \mathbf{k} be a field. Let \mathcal{C} be the category whose objects are \mathbf{k} -vector spaces F_L for each leaf set L of DX^* , where F_L has basis $\{v_l : l \in L\}$. The set of morphisms from F_L to $F_{L'}$ is the space $M(F_L, F_{L'})$ of all linear maps from F_L to $F_{L'}$. If $X \in V$, then we define $S(X)$ as follows: suppose that L is a leaf set on which X is defined, such that $X \cdot L = L'$. Write $X(l_i) = l'_i$ for $l_i \in L$. Then we define $\theta_X \in M(F_L, F_{L'})$ by

$$\theta_X(v_{l_i}) = v_{l'_i}.$$

We take $S(X)$ to be the set of all such θ_X . It is then easy to check that this makes V into a group with partial \mathcal{C} -action.

Example 3: We find a partial action of B_V . Let \mathcal{C} have objects labelled by leaf sets L . For morphisms, take all linear combinations of diagrams such as in Figure 4.7. Each of these diagrams has points at the top and bottom labelled left-to-right by leaf sets of the same size, which are put in bijection by braided lines. The domain of this diagram is the top leaf set and the range is the bottom; multiplication is by stacking the diagrams (and simplifying the braids if possible)

As in the case of V , whenever $X \in B_V$ can be defined by a picture with domain leaf set L , then a picture as in Figure 4.7 appears in a diagram

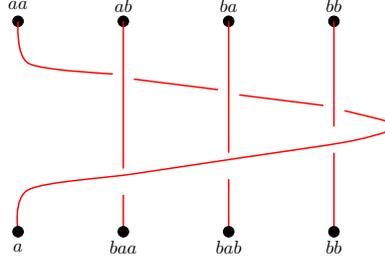


Figure 4.7: A typical morphism for the partial action of B_V

representing X , and we call this morphism $\theta_{X,L}$. We then take $S(X)$ to be the set of all such $\theta_{X,L}$ as the leaf set L varies (across all sufficiently deep leaf sets). Again, this gives a group with partial \mathcal{C} -action.

We remark that the partial actions we have put on V and B_V both extend linearly to their group algebras, which are algebras with partial \mathcal{C} -action. Also, we could restrict to a full subcategory of \mathcal{C} and get a new group with partial action. For example, we could restrict to all leaf sets of size dn^k , for some $k \in \mathbb{N}$. In particular, we define $\mathcal{C}_{n,d,R}$ to be the category whose morphisms are R -linear combinations of diagrams as in Figure 4.7, where both leaf sets have size dn^k . We call this the *standard braided category* (over R).

We now show an important structure theorem for the partial action of B_V .

Proposition 4.4.12. *Let $V = V_{n,d}$ be a Higman-Thompson group. Let B_V be given its partial action on the standard braided category, where the morphisms are (R -linear combinations of) diagrams as in Figure 4.7. For each leaf set L of size dn^k , let f_L be the unique element of $F \subset B_V$ mapping DX^k to L . Then the algebra of morphisms $\text{Hom}(DX^k, DX^k)$ is isomorphic to the group algebra $\mathbf{k}B_{dn^k}$, for each $k \in \mathbb{N}$. Moreover, for any leaf sets L, L' of size dn^k , we have*

$$\text{Hom}(L, L') = f_{L'} \text{Hom}(DX^k, DX^k) f_L^{-1}.$$

In particular, $\text{Hom}(L, L)$ is isomorphic to $\mathbf{k}B_{dn^k}$ for all L of size dn^k .

Proof. This is clear from the description of $\mathcal{C}_{n,d,R}$. Indeed morphisms from DX^k to itself are represented by braided arrangements of dn^k strings joining two sets of vertices labelled by DX^k , which is also the definition of the braid group. For the second part, the map $\theta \mapsto f_{L'}^{-1} \theta f_L$ clearly maps $\text{Hom}(L, L')$ to

$\text{Hom}(DX^k, DX^k)$, and is invertible with inverse $\phi \mapsto f_{L'}\phi f_L^{-1}$. This proves the equality required, and for $L = L'$, this is an algebra isomorphism. \square

We will define the Hecke algebra as a quotient of B_V with this partial action on $\mathcal{C}_{n,d,R}$.

Definition 4.4.13. *Let G be a group with a partial \mathcal{C} -action, given by function S . We define a quotient partial action of G to be the group G with its partial action on a quotient \mathcal{D} of \mathcal{C} . That is, for each morphism space $\text{Hom}(V, W)$ of \mathcal{C} , let $R_{V,W}$ be a subspace, satisfying the closure properties:*

$$R_{W,X} \circ \text{Hom}(V, W) \subset R_{V,X}; \quad \text{Hom}(W, X) \circ R_{V,W} \subset R_{V,X}.$$

Then there exists a category \mathcal{C}/R whose object space is $\text{Ob}(\mathcal{C})$ and whose morphism spaces are $\text{Hom}(V, W)/R(V, W)$. Moreover, G has a partial action on this space, via a function S/R that maps $g \in G$ to $S(g)/R$. This gives G a partial action on the category \mathcal{C}/R , which we call the quotient partial action.

This will let us define a Hecke algebra.

Definition 4.4.14. *Let \mathbf{k} be a field and let q be a variable. Let R be the ring $\mathbf{k}[q, q^{-1}]$. Let $V = V_{n,d}$ be a Higman-Thompson group, and let B_V be its braid group, with its partial action on the standard braided category. By Proposition 4.4.12, every hom space of \mathcal{C} has the form*

$$f_{L'} \cdot RB_{dn^k} f_L^{-1}.$$

Let π be the quotient map $\pi : B_{dn^k} \rightarrow \mathcal{H}_{dn^k, q}$ defined by adding relations $\Sigma^2 = (q-1)\Sigma + q$, whenever Σ is a standard generator of B_{dn^k} . Then there is a quotient of \mathcal{C} by π , where the hom spaces are of the form $f_{L'}\mathcal{H}_{dn^k, q}f_L^{-1}$. We define the quotient of RB_V under this partial action to be the Hecke algebra $\mathcal{H}_{V, q}$.

The quotient π is formed in this definition by quotienting out by the subspaces $R_{L, L'} = f_{L'}I_k f_L^{-1}$, where I_k is the two-sided ideal of RB_{dn^k} generated by $\Sigma_i^2 - (q-1)\Sigma_i - q$. These spaces $R_{L, L'}$ clearly have the necessary closure property for Definition 4.4.13 to apply.

This is our definition of the Hecke algebra. We will now go on to understand it from other points of view and discuss how it can be used, but fundamentally, it is just a partial action of $\mathbf{k}B_V$. The usual Hecke algebra $\mathcal{H}_{n, q}$ is a quotient of

$\mathbf{k}B_n$; here we weren't able to quotient the entire algebra, so just had to quotient its action on leaf sets instead.

$\mathcal{H}_{V,q}$ as a representation for the braided Thompson group

We first make a few remarks about what $\mathcal{H}_{V,q}$ tells us about representations. We will discuss this more later on. The algebra $\mathcal{H}_{V,q}$ that we have produced is just a group algebra of B_V , but its partial action is new and interesting. We remark that this gives a representation of B_V . Indeed, if $X \in B_V$, then we can study its partial action by viewing $S(X)$ as an element in the $\mathbf{k}[q, q^{-1}]$ -space

$$\mathrm{Hom}(\mathcal{C}) = \prod_{\mathcal{L}, \mathcal{L}'} \mathrm{Hom}(\mathcal{L}, \mathcal{L}').$$

The product here is taken over all pairs of leaf sets $\mathcal{L}, \mathcal{L}'$ where both have the same size, of the form dn^k . For $v \in \mathrm{Hom}(\mathcal{C})$, we write $v_{\mathcal{L}, \mathcal{L}'}$ for the component of v in $\mathrm{Hom}(\mathcal{L}, \mathcal{L}')$.

Notice that $S(X)$ lies in the subspace M_0 of $\mathrm{Hom}(\mathcal{C})$ on which the domain and range maps are injective. This M_0 has a multiplication where we multiply two elements of M_0 by composing homomorphisms wherever possible: it's easy to see that this preserves the fact that the domain and range are injective maps. This means that $X \mapsto S(X)$ almost gives a representation of B_V in the space M_0 , but this fails, because it's possible that $S(XY)$ contains morphisms that aren't in the product $S(X)S(Y)$.

To make this a representation, we notice that $S(X), S(Y)$ are defined on all sufficiently deep leaf sets (all leaf sets containing \mathcal{M} , for some leaf set \mathcal{M}), and so $S(XY)$ agrees with $S(X)S(Y)$ on all sufficiently deep leaf sets. So we will define a space M by taking a quotient of M_0 by a certain subspace N_0 . Here N_0 is defined as the set of all $v \in M_0$ such that there exists leaf set \mathcal{M} where $v_{\mathcal{L}, \mathcal{L}'}$ vanishes for all \mathcal{L} below \mathcal{M} . Then $M = M_0/N_0$ is an algebra, on which the quotient of S defines a representation of B_V . This seems to be an interesting new representation of B_V .

This completes our definition of \mathcal{H}_q , which we have defined as a quotient of $\mathbf{k}B_V$ with a partial action. We will now relate it to the other ways one can define finite-dimensional Iwahori-Hecke algebras. In particular, we will also try to understand $\mathcal{H}_{V,q}$ as an algebra of endomorphisms of a space of flags.

4.5 The Hecke algebra as collection of endomorphisms

In this section we describe $\mathcal{H}_{V,q}$ as an algebra of endomorphisms, in the case where q is the order of a finite field \mathbb{F} . We will relate $\mathcal{H}_{V,q}$ to constructions like the group $GL(\Gamma, \mathbb{F})$ which we tried to define a Hecke algebra with earlier. We will also see how to build $\mathcal{H}_{V,q}$ from finite Hecke algebras $\mathcal{H}_{m,q}$, which are endomorphism algebras of flag spaces. We will extend the action of endomorphisms in $\mathcal{H}_{m,q}$ to higher flag spaces to allow us to fit these finite Hecke algebras together.

Here we'll again work with the subgroup $\mathfrak{S}_{\text{hom}}$ of V . In fact we'll define an algebra for $\mathfrak{S}_{\text{hom}}$ first and then add in Thompson's group F (essentially by inducing the representation) to form an algebra for all of V .

4.5.1 The $\mathfrak{S}_{\text{hom}}$ Hecke algebra

In this section, we define an algebra by deforming just the subalgebra $\mathbf{k}\mathfrak{S}_{\text{hom}}$ of $\mathbf{k}V$. Put $G = \mathfrak{S}_{\text{hom}}$ here. As in the \mathfrak{S}_n case, the Hecke algebra will arise from endomorphisms of a G -module spanned by flags, and will satisfy relations that come from deforming a presentation of G .

To begin the construction, we need an equivalent of GL_n for this situation.

Definition 4.5.1. *Let \mathbb{F} be a finite field. We define the homogeneous general linear group, $GL_{\text{hom}}(\mathbb{F})$, as a direct limit of finite general linear groups:*

$$GL_d(\mathbb{F}) \hookrightarrow GL_{nd}(\mathbb{F}) \hookrightarrow GL_{n^2d}(\mathbb{F}) \hookrightarrow \dots$$

The embedding from GL_m to GL_{nm} sends X to $X \otimes Id_n$, the tensor product (Kronecker product) of X and an $n \times n$ identity matrix.

In particular, one can verify that the permutation matrices of GL_{hom} form a copy of $\mathfrak{S}_{\text{hom}}$. We also point out that GL_{hom} is a subgroup of $GL(\Gamma, \mathbb{F})$ of Section 4.3.1 (at least in the case of $V_{2,1}$, and we can generalize at least to $V_{n,1}$ by making Γ into n loops). Indeed $GL_{n^k d}$ is the group of all \mathcal{L} -matrices of $L_{\mathbb{F}}(\Gamma)$ where \mathcal{L} is the leaf set of all depth k leaves, and the embeddings are compatible.

We now introduce flags for this group.

Definition 4.5.2. *Let \mathbf{k} be a field (which we will usually think of as \mathbb{C} , but which is unrestricted). We define a $\mathbf{k}GL_{\text{hom}}$ module \mathcal{F} as follows. Let W be the*

vector space:

$$W = \mathbb{F}^d \otimes \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n \otimes \dots,$$

with infinitely many copies of \mathbb{F}^n . We fix a basis $e_{i,1}, e_{i,2}, \dots, e_{i,n}$ for the i th copy of \mathbb{F}^n . Then each $k \in \mathbb{N}_0$, define the dn^k -dimensional subspace W_k as:

$$W_k = \mathbb{F}^d \otimes \underbrace{\mathbb{F}^n \otimes \dots \otimes \mathbb{F}^n}_{k \text{ factors}} \otimes \langle e_{k+1,1} \rangle \otimes \langle e_{k+2,1} \rangle \otimes \dots \leq W.$$

For each $k \in \mathbb{N}$, the group $GL_{n^k d}(\mathbb{F})$ acts on W by acting on the first $k+1$ spaces. This action preserves the subspace W_k . Moreover, these actions are compatible with the embeddings $GL_{n^k d} \hookrightarrow GL_{n^{k+1} d}$, by definition of the Kronecker product. Thus, GL_{hom} acts on W .

For each $k \in \mathbb{N}$, we say a **level k flag** is a chain of subspaces of the form:

$$F_k : \{0\} = W_{k,0} \subset W_{k,1} \subset W_{k,2} \subset \dots \subset W_{k,n^k d} = W_k,$$

where each $W_{k,i}$ is a subspace of W_k , of codimension 1 in $W_{k,i+1}$. We write $F_k(i) = W_{k,i}$. Let \mathcal{F}_k be the $\mathbf{k}GL_{dn^k}$ module with basis the level k flags. Then GL_{dn^k} also acts on $\mathcal{F}_{k'}$ for each $k' > k$ via the embedding defining GL_{hom} .

Finally, we let \mathcal{F}_\times be the direct product of all the spaces \mathcal{F}_k , and let \mathcal{F}_+ be their direct sum. Take \mathcal{F} to be the quotient space $\mathcal{F}_\times / \mathcal{F}_+$ (that is, elements of \mathcal{F} are sequences of linear combinations of flags, one combination for each d , modulo sequences of finite support). Then \mathcal{F} is a $\mathbf{k}GL_{\text{hom}}$ module, because each element of GL_{hom} acts on all but finitely many of the modules in the direct product.

Now we define endomorphisms of the GL_{hom} -module \mathcal{F} , which will be defined by giving an endomorphism of \mathcal{F}_k for all sufficiently large k .

Definition 4.5.3. For each $k \in \mathbb{N}$ and $1 \leq i \leq dn^k - 1$, we define an endomorphism $\Sigma_{k,i}$ of the GL_{dn^k} -module \mathcal{F}_k in the usual manner for GL_{dn^k} : say two flags $F_{k,1}$ and $F_{k,2}$ are i -neighbours if $F_{k,1}(i) \neq F_{k,2}(i)$, but $F_{k,1}(j) = F_{k,2}(j)$ for all $j \neq i$. Then $\Sigma_{k,i}$ acts on the basis of \mathcal{F}_k by sending a flag to the sum of all its i -neighbours. This gives a GL_{dn^k} -endomorphism, as we saw in Section 4.2.2.

Next we will extend the action of the symbol $\Sigma_{k,i}$ to define it on spaces $\mathcal{F}_{k'}$, for $k' > k$. Let $k' = k + r$. For $F \in \mathcal{F}_{k'}$, we define:

$$\begin{aligned} \Sigma_{k,i}(F) = & (\Sigma_{k',n^r i} \cdots \Sigma_{k',n^r(i-1)+1}) (\Sigma_{k',n^r i+1} \cdots \Sigma_{k',n^r(i-1)+2}) \cdots \\ & \cdots (\Sigma_{k',n^r i+n-1} \cdots \Sigma_{k',n^r i})(F). \end{aligned}$$

In this formula, each bracket consists of terms $\Sigma_{k',a} \Sigma_{k',a-1} \Sigma_{k',a-2} \cdots \Sigma_{k',a-n^r+1}$. This is an endomorphism of the $GL_{n^{k'}d}$ -module $\mathcal{F}_{k'}$, because it's a product of terms $\Sigma_{k',i}$ which are endomorphisms. Thus, with these definitions, Σ_i defines an endomorphism of the GL_{hom} -module \mathcal{F} .

Overall, $\Sigma_{k,i}$ acts as the Hecke generators of type A_n of Section 4.2.2 in its action on \mathcal{F}_k . We extend the action of this symbol to all $\mathcal{F}_{k'}$ by an expansion rule, using a formula that gives the correct expansion of transpositions in \mathfrak{S}_{dn^k} .

We now discuss invertibility. When just viewed as endomorphisms of \mathcal{F}_k , the operators $\Sigma_{k,i}$ are invertible: indeed, $\Sigma_{k,i}$ satisfies the quadratic relation $\Sigma_{k,i}^2 = (q-1)\Sigma_i + q$, so has inverse

$$\Sigma_{k,i}^{-1} = \frac{1}{q}(\Sigma_{k,i} - (q-1)).$$

When extended to act on $\mathcal{F}_{k'}$ for $k' > k$, this formula no longer defines an inverse for $\Sigma_{k,i}$. However, the action of $\Sigma_{k,i}$ is still invertible, since it acts as a product of invertible elements $\Sigma_{k',j}$. So we can define a map $\Sigma_{k,i}^{-1}$, which is a $GL_{\text{hom}}(\mathbb{F}_q)$ -endomorphism of \mathcal{F} inverse to $\Sigma_{k,i}$.

We define $\mathcal{H}_{\text{hom},q}^0$ to be the k -algebra generated by the endomorphisms $\Sigma_{k,i}$ and $\Sigma_{k,i}^{-1}$ for $GL_{\text{hom}}(\mathbb{F}_q)$. We use the notation \mathcal{H}^0 to distinguish the algebras we construct here from the Hecke algebra $\mathcal{H}_{V,q}$ we defined earlier.

Next, we show some relations that hold in $\mathcal{H}_{\text{hom},q}^0$.

Proposition 4.5.4. *The endomorphisms $\Sigma_{k,i}$ satisfy the following relations:*

- $\Sigma_{k,i} \Sigma_{k,i+1} \Sigma_{k,i} = \Sigma_{k,i+1} \Sigma_{k,i} \Sigma_{k,i+1}$ for each $1 \leq i \leq dn^k - 2$
- $\Sigma_{k,i} \Sigma_{k,j} = \Sigma_{k,j} \Sigma_{k,i}$ if $|i - j| > 1$.
- For each k and $1 \leq i \leq dn^k - 1$, then

$$\Sigma_{k,i} = C_{k+1,ni} C_{k+1,ni+1} \cdots C_{k+1,ni+n-1},$$

where

$$C_{k+1,y} = \Sigma_{k+1,y} \Sigma_{k+1,y-1} \cdots \Sigma_{k+1,y-n+1}.$$

These relations parallel the defining relations of $\mathfrak{S}_{\text{hom}}$, but we don't have a quadratic relation in this case. Note that here we're just claiming that the $\Sigma_{k,i}$ satisfy the relations: we're not saying anything about whether it is a presentation.

Proof. Since the $\Sigma_{k,i}$ are defined as endomorphisms of \mathcal{F} , we must verify that these relations hold for the action of $\Sigma_{k,i}$ on flags of any level. First we consider the action on level k flags. Restricted to this subspace, the elements $\Sigma_{k,i}$ act as generators of the finite-dimensional Hecke algebra $\mathcal{H}_{nd^k,q}$ so satisfy the first two relations. The action of $\Sigma_{k,i}$ on level $k+1$ flags agrees with the third relation by definition.

Now consider the action on flags of level k' deeper than k . Since the action of $\Sigma_{k,i}$ on level k' flags is by a product of endomorphisms $\Sigma_{k',j}$, each of the three relations then becomes the assertion that two products of terms $\Sigma_{k',j}$ are equal. Looking at the corresponding relations with $\sigma_{k',j}$, we get a relation that is true in $\mathfrak{S}_{\text{hom}}$, given by writing a word in $\mathfrak{S}_{dn^{k'}}$ as a product of Coxeter generators in two different ways, and by Proposition 4.4.4 both these products are reduced, since they're formed by expanding reduced words in \mathfrak{S}_{dn^k} or $\mathfrak{S}_{dn^{k+1}}$. We know that any such relation can be demonstrated using only braid relations $\sigma_{k',i}\sigma_{k',i+1}\sigma_{k',i} = \sigma_{k',i+1}\sigma_{k',i}\sigma_{k',i+1}$ and $\sigma_{k',i}\sigma_{k',j} = \sigma_{k',j}\sigma_{k',i}$. Since these relations hold for the $\Sigma_{k',i}$ as well, we are done. \square

This completes the definition of a deformed version of $\mathfrak{S}_{\text{hom}}$. We remark that in defining $\mathcal{H}_{\text{hom},q}^0$ we lose the quadratic relations $\Sigma_{k,i}^2 = (q-1)\Sigma_{k,i} + q$ which hold in $\mathcal{H}_{dn^k,q}$. This happens because when we extend $\Sigma_{k,i}$ to act on spaces $\mathcal{F}_{k'}$ for $k' > k$, this relation no longer holds. This means that $\Sigma_{k,i}^2 - (q-1)\Sigma_{k,i} - q$ is a non-zero element of $\mathcal{H}_{\text{hom},q}^0$. This sort of fact is one of the reasons why we defined the Hecke algebra $\mathcal{H}_{V,q}$ to be an algebra with a partial action: we want the quadratic relation to hold locally (in the action of $\Sigma_{k,i}$ on \mathcal{F}_k), even though it cannot hold globally.

Next we extend $\mathcal{H}_{\text{hom},q}^0$ to an algebra for all of V , which we will relate back to $\mathcal{H}_{V,q}$.

4.5.2 Extending $\mathcal{H}_{\text{hom},q}^0$ to V

In this section we create a larger space than \mathcal{F} so that the Higman-Thompson group F acts on it by GL_{hom} -automorphisms. We will then define an algebra $\mathcal{H}_{V,q}^0$ as a product of $\mathcal{H}_{\text{hom},q}^0$ and F . The construction will be similar to $\mathcal{H}_{V,q}^0$,

except that we will have many more flag spaces: one flag space for every leaf set of size dn^k instead of just one flag space for each number dn^k . Essentially, we're just inducing from the representation of $\mathcal{H}_{\text{hom},q}^0$ on \mathcal{F} up to a larger algebra $\mathcal{H}_{V,q}^0$; the construction looks complicated because we're also defining the algebra as we go.

Definition 4.5.5. *Let \mathbf{k} be a field. Let $k \in \mathbb{N}_0$, and let \mathcal{L} be any leaf set of DX^* with dn^k leaves. Let F_k be a level k flag, so that*

$$F_k : \{0\} = W_{k,0} \subset W_{k,1} \subset W_{k,2} \subset \dots \subset W_{k,dn^k} = W_k.$$

Label the leaves of \mathcal{L} from left to right with the spaces $W_{k,1}, W_{k,2}, \dots, W_{k,dn^k}$. Let $\mathcal{F}_{\mathcal{L}}$ be the $\mathbf{k}GL_{dn^k}$ module with basis all such labellings of \mathcal{L} as the flag F_k varies - the action of $\mathbf{k}GL_{dn^k}$ is on the flags, changing the labels and leaving \mathcal{L} the same. As before, $\mathbf{k}GL_{dn^k}$ also acts on $\mathcal{F}_{L'}$ whenever L' is a leaf set of size $dn^{k'}$, and $k' \geq k$ (via the embedding of group algebras, $\mathbf{k}GL_{dn^k} \hookrightarrow \mathbf{k}GL_{dn^{k'}}$).

Now we introduce an action of Thompson's group F . Let $f \in F$. Then $f(L)$ is defined whenever L is a leaf set lying below the domain leaf set of f . Thus, we can define an action of f sending \mathcal{F}_L to $\mathcal{F}_{f(L)}$, where the action changes the leaf set \mathcal{L} to $f(\mathcal{L})$ but preserves the labels.

Putting this together, we'll define $\mathcal{F}_{V,\times}$ to be the direct product of all the \mathbf{k} -vector spaces $\mathcal{F}_{\mathcal{L}}$ as the leaf set \mathcal{L} varies. If $\mathbf{v} \in \mathcal{F}_{V,\times}$, we write $\mathbf{v}(\mathcal{L})$ for the coordinate of \mathbf{v} in the space $\mathcal{F}_{\mathcal{L}}$, which will be a linear combination of labellings of \mathcal{L} by flags. For M a leaf set of size nd^k , we say that \mathbf{v} is M -null if $\mathbf{v}(\mathcal{L}) = 0$ whenever \mathcal{L} lies below M . Notice that if \mathbf{v} is M -null, it is also M' -null for all M' below M . Then the set of all M -null vectors, as M varies, forms a subspace of $\mathcal{F}_{V,\times}$, which we will call $\mathcal{F}_{V,+}$. Finally we define space \mathcal{F}_V as the quotient of $\mathcal{F}_{V,\times}$ by $\mathcal{F}_{V,+}$.

The idea of this definition is that we have a flag space for every leaf set of the same size as some set DX^k . Each element of our algebra will act on all sufficiently deep flag spaces, in the same way that each element of V acts on sufficiently deep leaf sets.

Proposition 4.5.6. *The actions of GL_{dn^k} and Thompson's group F on sufficiently deep spaces $\mathcal{F}_{\mathcal{L}}$ make \mathcal{F}_V into a $\mathbf{k}GL_{\text{hom}}$ module and a $\mathbf{k}F$ -module.*

Proof. Let $X \in GL_{\text{hom}}$ and let $\mathbf{v} \in \mathcal{F}_V$. We wish to define $X(\mathbf{v})$. Suppose $X \in GL_{dn^k}$. Choose a representative $\tilde{\mathbf{v}} \in \mathcal{F}_{V,\times}$ for \mathbf{v} such that $\tilde{\mathbf{v}}(\mathcal{L}) = 0$

whenever \mathcal{L} has fewer than dn^k leaves, which is possible by definition of $\mathcal{F}_{V,+}$. Then the action of X on $\tilde{\mathbf{v}}(\mathcal{L})$ is defined for all leaf sets \mathcal{L} where $\tilde{\mathbf{v}}(\mathcal{L})$ is non-zero. If we choose a different representative $\hat{\mathbf{v}}$ for \mathbf{v} (which differs from $\tilde{\mathbf{v}}$ by an M -null vector), then $X(\tilde{\mathbf{v}})$ and $X(\hat{\mathbf{v}})$ still differ by an element of $\mathcal{F}_{V,+}$ (which is still M -null). Similarly if we considered X as an element of GL_{dn^k} instead, we have seen that the action of X on all sufficiently deep flags is the same regardless of choice of k . This shows that the action of GL_{hom} is well-defined.

To show $f \in F$ acts on \mathcal{F}_V the proof is the same, except that we instead choose a representative $\tilde{\mathbf{v}}$ for \mathbf{v} that vanishes on any leaf set not below a domain leaf set for f . Then again, $f(\mathcal{L})$ is defined everywhere $\tilde{\mathbf{v}}(\mathcal{L})$ is non-zero, which lets us define $f(\mathbf{v})$. Changing the domain leaf set for f only changes $f(\mathbf{v})$ by an element of $\mathcal{F}_{V,+}$, so we get a well-defined action in this case also. \square

Proposition 4.5.7. *F acts on \mathcal{F}_V as a group of $\mathbf{k}GL_{\text{hom}}$ -module automorphisms.*

Proof. We just need to check that the action of F commutes with the action of GL_{hom} . But this is immediate, because F permutes the leaf sets \mathcal{L} independently of the labels, whilst GL_{hom} fixes each leaf set and just changes the (linear combinations of) flags that label them. \square

Proposition 4.5.8. *The $\mathbf{k}GL_{\text{hom}}$ -module \mathcal{F} embeds into \mathcal{F}_V by the map θ identifying F_k with $F_{\mathcal{L}}$, for $\mathcal{L} = DX^k$.*

Proof. This is true because our constructions were parallel. For $\mathcal{L} = DX^k$, the space $\mathcal{F}_{\mathcal{L}}$ is isomorphic to the subspace \mathcal{F}_k of \mathcal{F} , with isomorphic GL_{dn^l} actions (for each $l \leq k$). This means that we can define $\theta : \mathcal{F}_{\times} \rightarrow \mathcal{F}_{V,\times}$, and we need to show that the image of \mathcal{F}_+ under θ lies in $\mathcal{F}_{V,+}$. If $\mathbf{v} \in \mathcal{F}_{\times}$, then $\theta(\mathbf{v})$ is null if and only if $\mathbf{v}(k)$ vanishes for all sufficiently large k ; that is, if and only if $\mathbf{v} \in \mathcal{F}_+$. So the quotients taken in each case are compatible. \square

We will identify \mathcal{F} as a submodule of \mathcal{F}_V from now on to simplify the notation.

4.5.3 Endomorphisms

The purpose of this section is to finish construction of a Hecke algebra for V by extending the endomorphisms $\Sigma_{k,i}$ from \mathcal{F} to all of \mathcal{F}_V . Recall that $\mathfrak{S}_{\text{hom}}$ is generated by elements $\sigma_{k,i}$ satisfying relations:

- $\sigma_{k,i}^2 = 1$, all k and i .
- $\sigma_{k,i}\sigma_{k,i+1}\sigma_{k,i} = \sigma_{k,i+1}\sigma_{k,i}\sigma_{k,i+1}$, each k and $1 \leq i \leq dn^k - 2$.
- $\sigma_{k,i}$ and $\sigma_{k,j}$ commute whenever $|i - j| > 1$.
- For each $k \in \mathbb{N}_0$, $m = dn^k$ and $1 \leq i \leq m - 1$, then

$$\sigma_{k,i} = c_{k+1,ni}c_{k+1,ni+1} \cdots c_{k+1,ni+n-1},$$

where

$$c_{k+1,y} = \sigma_{k+1,y}\sigma_{k+1,y-1} \cdots \sigma_{k+1,y-n+1}.$$

If we add to these generators all elements $f \in F$, and add in a relation for each $f \in F$ and $k \in \mathbb{N}$, $1 \leq i \leq dn^k - 1$, saying:

- $\sigma_{k,i}f = g\sigma$ where $g \in F$, $\sigma = \sigma_{k_1,i_1} \cdots \sigma_{k_r,i_r} \in \mathfrak{S}_{\text{hom}}$ are the unique elements of those groups such that this relation holds in V .

we get a presentation for V by Proposition 4.4.9. We will produce an algebra $\mathcal{H}_{V,q}^0$ based on this presentation for $\mathbf{k}V$. First of all we show that there is a well-defined Hecke version of any $\sigma \in \mathfrak{S}_{\text{hom}}$.

Proposition 4.5.9. *Let $\sigma \in \mathfrak{S}_{\text{hom}}$, and write $\sigma = \sigma_{k,i_1}\sigma_{k,i_2} \cdots \sigma_{k,i_r}$ as a product of generators (each of which lies in \mathfrak{S}_{dn^k} for fixed k). Assume that as a word in \mathfrak{S}_{dn^k} , the word $\sigma_{k,i_1}\sigma_{k,i_2} \cdots \sigma_{k,i_r}$ is reduced. Then there exists an endomorphism Σ of \mathcal{F} defined by:*

$$\Sigma = \Sigma_{k,i_1}\Sigma_{k,i_2} \cdots \Sigma_{k,i_r},$$

and this Σ does not depend on the choice of k or on the i_r .

Proof. We know that Σ is an endomorphism, so we just need to check the definition of Σ does not depend on the choice of k and the i_r . First fix k . Suppose that $\sigma_{k,i_1} \cdots \sigma_{k,i_r}$ and $\sigma_{k,j_1} \cdots \sigma_{k,j_r}$ are two reduced expressions for σ (which necessarily have the same length r). We make the usual argument: we can convert one such expression into the other by only using braid relations

$$\sigma_{k,i}\sigma_{k,i+1}\sigma_{k,i} = \sigma_{k,i+1}\sigma_{k,i}\sigma_{k,i+1}$$

and commutation relations

$$\sigma_{k,i}\sigma_{k,j} = \sigma_{k,j}\sigma_{k,i}.$$

These relations hold for $\Sigma_{k,i}$ in place of $\sigma_{k,i}$ so both expressions will yield the same element Σ .

Now we need to compare two different values of k . So suppose $\sigma_{k,i_1} \dots \sigma_{k,i_r}$ and $\sigma_{k+r,j_1} \dots \sigma_{k+r,j_s}$ are both reduced expressions for σ . Define

$$\Sigma_0 = \Sigma_{k,i_1} \dots \Sigma_{k,i_r}$$

and

$$\Sigma_1 = \Sigma_{k+r,j_1} \dots \Sigma_{k+r,j_s}.$$

We wish to show $\Sigma_0 = \Sigma_1$. First repeatedly use the relation

$$\Sigma_{k,i} = C_{k+1,ni} C_{k+1,ni+1} \dots C_{k+1,ni+n-1}$$

to write Σ_0 as a product of terms of the form $\Sigma_{k+r,i}$. The same process also writes σ as a product of Coxeter generators of $\mathfrak{S}_{dn^{k+l}}$, and Proposition 4.4.4 tells us that this expression is reduced. So now we have produced two different reduced expressions for $\sigma \in \mathfrak{S}_{dn^{k+l}}$ - one from Σ_1 and one from expanding Σ_0 . Once again, the two expressions for σ can be converted into each other with only braid and commutation relations, which also hold among the $\Sigma_{k+l,i}$, so we get $\Sigma_0 = \Sigma_1$. \square

The point of this theorem is that we can now talk about the endomorphism Σ coming from deforming a permutation $\sigma \in \mathfrak{S}_{\text{hom}}$, and this is well-defined. We now extend these endomorphisms from \mathcal{F} to all of \mathcal{F}_V .

Theorem 4.5.10. *Let \mathcal{L} be a leaf set with dn^l leaves and let f be the unique element of F sending DX^l to \mathcal{L} . Let $\sigma_{k,i}$ be a generating transposition of $\mathfrak{S}_{\text{hom}}$ such that (as an element of V) $\sigma_{k,i}$ sends leaf set \mathcal{L} to \mathcal{L}' , and let f' be the unique element of F sending DX^l to τ' . Define σ to be the element of \mathfrak{S}_{dn^l} making the diagram below commute:*

$$\begin{array}{ccc} \mathcal{L} & \xleftarrow{f} & DX^l \\ \downarrow \sigma_{k,i} & & \downarrow \sigma \\ \mathcal{L}' & \xleftarrow{f'} & DX^l \end{array}$$

Let Σ be the element of $\mathcal{H}_{\text{hom},q}^0$ corresponding to permutation σ , as in the previous proposition. Then we define the action of $\Sigma_{k,i}$ on $\mathcal{F}_{\mathcal{L}}$ to be the endomorphism making the following diagram commute:

$$\begin{array}{ccc}
\mathcal{F}_{\mathcal{L}} & \xleftarrow{f} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i} & & \downarrow \Sigma \\
\mathcal{F}_{\mathcal{L}'} & \xleftarrow{f'} & \mathcal{F}_l
\end{array}$$

With these definitions, the $\Sigma_{k,i}$ still satisfy the relations of Proposition 4.5.4

- $\Sigma_{k,i}\Sigma_{k,i+1}\Sigma_{k,i} = \Sigma_{k,i+1}\Sigma_{k,i}\Sigma_{k,i+1}$ for each $1 \leq i \leq dn^k - 2$
- $\Sigma_{k,i}\Sigma_{k,j} = \Sigma_{k,j}\Sigma_{k,i}$ if $|i - j| > 1$.
- For each $k \in \mathbb{N}_0$, $m = dn^k$ and $1 \leq i \leq m - 1$, then

$$\Sigma_{k,i} = C_{k+1,ni}C_{k+1,ni+1} \cdots C_{k+1,ni+n-1},$$

where

$$C_{k+1,y} = \Sigma_{k+1,y}\Sigma_{k+1,y-1} \cdots \Sigma_{k+1,y-n+1}.$$

In addition, whenever $\sigma_{k,i}f = g\sigma$ for $f, g \in F, \sigma \in \mathfrak{S}_{hom}$, then

$$\Sigma_{k,i}f = g\Sigma,$$

where Σ is the deformed version of σ as defined as in the previous proposition, which is still well-defined. Finally, each $\Sigma_{k,i}$ is invertible as an endomorphism of \mathcal{F}_V , so can have inverse $\Sigma_{k,i}^{-1}$ defined.

We define the algebra $\mathcal{H}_{V,q}^0$ to be the k -algebra generated by the $\Sigma_{k,i}$ and their inverses acting on \mathcal{F}_V .

Proof. We prove the statements one at a time. The proofs will be similar. In each case, we will show that both sides act on $\mathcal{F}_{\mathcal{L}}$ in the same way, for all large enough leaf sets \mathcal{L} (that is, all leaf sets below some leaf set M , where M can vary from one identity to another). This means their actions on $\mathcal{F}_{V,1}$ differ by a null vector and so they act in the same way as endomorphisms of \mathcal{F}_V .

The braid relations: Fix a sufficiently deep leaf set \mathcal{L} with dn^l leaves. Consider the definition of the product $\Sigma_{k,i}\Sigma_{k,i+1}\Sigma_{k,i}$ on $\mathcal{F}_{\mathcal{L}}$. For deep enough \mathcal{L} , there is a commutative diagram showing the V -action on finite trees:

$$\begin{array}{ccc}
\mathcal{L} = \mathcal{L}_0 & \xleftarrow{f_0} & DX^l \\
\downarrow \sigma_{k,i} & & \downarrow \sigma_1 \\
\mathcal{L}_1 & \xleftarrow{f_1} & DX^l \\
\downarrow \sigma_{k,i+1} & & \downarrow \sigma_2 \\
\mathcal{L}_2 & \xleftarrow{f_2} & DX^l \\
\downarrow \sigma_{k,i} & & \downarrow \sigma_3 \\
\mathcal{L}_3 & \xleftarrow{f_3} & DX^l
\end{array}$$

Here, the \mathcal{L}_i are leaf sets, the f_i lie in F , and the maps σ_i lie in $\mathfrak{S}_{\text{hom}}$ and are chosen to make the diagram commute. There is a corresponding commutative diagram to define the action of $\Sigma_{k,i}\Sigma_{k,i+1}\Sigma_{k,i}$ on $\mathcal{F}_{\mathcal{L}}$ (where we write $F_{\mathcal{L},i}$ instead of keeping the subscript on \mathcal{L}_i):

$$\begin{array}{ccc}
\mathcal{F}_{\mathcal{L},0} & \xleftarrow{f_0} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i} & & \downarrow \Sigma_1 \\
\mathcal{F}_{\mathcal{L},1} & \xleftarrow{f_1} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i+1} & & \downarrow \Sigma_2 \\
\mathcal{F}_{\mathcal{L},2} & \xleftarrow{f_2} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i} & & \downarrow \Sigma_3 \\
\mathcal{F}_{\mathcal{L},3} & \xleftarrow{f_3} & \mathcal{F}_l
\end{array}$$

There are similar diagrams for the other braid relation (for leaf sets \mathcal{M}_i and $\tau_i \in \mathfrak{S}_{\text{hom}}$ corresponding to $T_i \in \mathcal{H}_{\text{hom},q}$ and $g_i \in F$):

$$\begin{array}{ccc}
\mathcal{L} = \mathcal{M}_0 & \xleftarrow{g_0} & DX^l \\
\downarrow \sigma_{k,i+1} & & \downarrow \tau_1 \\
\mathcal{M}_1 & \xleftarrow{g_1} & DX^l \\
\downarrow \sigma_{k,i} & & \downarrow \tau_2 \\
\mathcal{M}_2 & \xleftarrow{g_2} & DX^l \\
\downarrow \sigma_{k,i+1} & & \downarrow \tau_3 \\
\mathcal{L}_3 = \mathcal{M}_3 & \xleftarrow{g_3} & DX^l
\end{array}
\qquad
\begin{array}{ccc}
\mathcal{F}_{\mathcal{L},0} = \mathcal{F}_{\mathcal{M},0} & \xleftarrow{g_0} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i+1} & & \downarrow T_1 \\
\mathcal{F}_{\mathcal{M},1} & \xleftarrow{g_1} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i} & & \downarrow T_2 \\
\mathcal{F}_{\mathcal{M},2} & \xleftarrow{g_2} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i+1} & & \downarrow T_3 \\
\mathcal{F}_{\mathcal{L},3} = \mathcal{F}_{\mathcal{M},3} & \xleftarrow{g_3} & \mathcal{F}_l
\end{array}$$

Since $\sigma_{d,i}\sigma_{d,i+1}\sigma_{d,i} = \sigma_{d,i+1}\sigma_{d,i}\sigma_{d,i+1}$, it follows that $\sigma_3\sigma_2\sigma_1 = \tau_3\tau_2\tau_1$. We claim that $l(\sigma_3\sigma_2\sigma_1) = l(\sigma_1) + l(\sigma_2) + l(\sigma_3)$ - that is, a reduced expression for $\sigma_3\sigma_2\sigma_1$ can be made by concatenating reduced expressions for σ_1, σ_2 and σ_3 (and also, the same is true for $\tau_3\tau_2\tau_1$). This means that $\sigma_3\sigma_2\sigma_1$ and $\tau_3\tau_2\tau_1$ both expand to reduced expressions for the same word, so one can be converted into the other using only braiding and commutation relations. Thus $\Sigma_3\Sigma_2\Sigma_1 = T_3T_2T_1$, which gives the result.

It remains to establish the length result. In the leaf set \mathcal{L} , let the leaves below the i th vertex v_i of DX^k be $v_i l_1, v_i l_2, \dots, v_i l_r$, the leaves below the $i+1$ th be $v_{i+1} l_{r+1}, \dots, v_{i+1} l_{r+s}$, and the leaves below the $i+2$ th be $v_{i+2} l_{r+s+1}, \dots, v_{i+2} l_{r+s+t}$. Then the permutations $\sigma_1, \sigma_2, \sigma_3$ permute leaves in the following way, fixing all other leaves:

$$\begin{aligned}
& \dots (v_i l_1, v_i l_2, \dots, v_i l_r, v_{i+1} l_{r+1} \dots, v_{i+1} l_{r+s}, v_{i+2} l_{r+s+1}, \dots, v_{i+2} l_{r+s+t}) \dots \\
& \xrightarrow{\sigma_{k,i}} \dots (v_{i+1} l_1, v_{i+1} l_2, \dots, v_{i+1} l_r, v_i l_{r+1} \dots, v_i l_{r+s}, v_{i+2} l_{r+s+1}, \dots, v_{i+2} l_{r+s+t}) \dots \\
& \xrightarrow{\sigma_{k,i+1}} \dots (v_{i+2} l_1, v_{i+2} l_2, \dots, v_{i+2} l_r, v_i l_{r+1} \dots, v_i l_{r+s}, v_{i+1} l_{r+s+1}, \dots, v_{i+1} l_{r+s+t}) \dots \\
& \xrightarrow{\sigma_{k,i}} \dots (v_{i+2} l_1, v_{i+2} l_2, \dots, v_{i+2} l_r, v_{i+1} l_{r+1} \dots, v_{i+1} l_{r+s}, v_i l_{r+s+1}, \dots, v_i l_{r+s+t}) \dots
\end{aligned}$$

Let $m_1, m_2, \dots, m_{r+s+t}$ be the corresponding vertices of DX^l (under a bijection with \mathcal{L}). Then the permutation σ_1 must leave DX^l in the same order that $\sigma_{k,1}$ leaves $\sigma_{k,1}(\mathcal{L})$. This means that DX^l is correspondingly permuted in the following manner:

$$\begin{aligned}
& \dots (m_1, m_2, \dots, m_r, m_{r+1} \dots, m_{r+s}, m_{r+s+1}, \dots, m_{r+s+t}) \dots \\
\begin{array}{c} \xrightarrow{\sigma_1} \\ \xrightarrow{\sigma_2} \\ \xrightarrow{\sigma_3} \end{array} & \dots (m_{s+1}, m_{s+2}, \dots, m_{r+s}, m_1, m_2, \dots, m_s, m_{r+s+1}, \dots, m_{r+s+t}) \dots \\
& \dots (m_{s+t+1}, m_{s+t+2}, \dots, m_{r+s+t}, m_1, m_2, \dots, m_s, m_{s+1}, m_{s+2}, \dots, m_{r+s}) \dots \\
& \dots (m_{s+t+1}, m_{s+t+2}, \dots, m_{r+s+t}, m_{r+1}, m_{r+2}, \dots, m_{r+s}, m_1, m_2, \dots, m_r) \dots
\end{aligned}$$

We calculate length of a permutation σ as the number of pairs i, j where m_i appears to the left of m_j , but $\sigma(m_i)$ appears to the right of $\sigma(m_j)$. One can verify that $\ell(\sigma_1) = rs$, $\ell(\sigma_2) = rt$ and $\ell(\sigma_3) = st$. Similarly, $\ell(\sigma_3\sigma_2\sigma_1) = r + s + t$ (put another way, it interchanges m_i and m_j if and only if i and j are in different sets when we partition $\{1, \dots, r + s + t\}$ into $\{1, \dots, r\}$, $\{r + 1, \dots, r + s\}$ and $\{r + s + 1, \dots, r + s + t\}$; each of $\sigma_1, \sigma_2, \sigma_3$ interchanges a different pair of these blocks). This establishes the result.

1. **The commutation relations:** This is similar to the previous case, but simpler. We have to compare:

$$\begin{array}{ccc}
\mathcal{L} = \mathcal{L}_0 & \xleftarrow{f_0} & DX^l \\
\downarrow \sigma_{k,i} & & \downarrow \sigma_1 \\
\mathcal{L}_1 & \xleftarrow{f_1} & DX^l \\
\downarrow \sigma_{k,j} & & \downarrow \sigma_2 \\
\mathcal{L}_2 & \xleftarrow{f_2} & DX^l
\end{array}
\qquad
\begin{array}{ccc}
\mathcal{F}_{\mathcal{L},0} & \xleftarrow{f_1} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i} & & \downarrow \Sigma_1 \\
\mathcal{F}_{\mathcal{L},1} & \xleftarrow{f_2} & \mathcal{F}_l \\
\downarrow \Sigma_{k,j} & & \downarrow \Sigma_2 \\
\mathcal{F}_{\mathcal{L},2} & \xleftarrow{f_3} & \mathcal{F}_l
\end{array}$$

and

$$\begin{array}{ccc}
\mathcal{L} = \mathcal{M}_0 & \xleftarrow{g_0} & DX^l \\
\downarrow \sigma_{k,i} & & \downarrow \tau_1 \\
\mathcal{M}_1 & \xleftarrow{g_1} & DX^l \\
\downarrow \sigma_{k,j} & & \downarrow \tau_2 \\
\mathcal{L}_2 = \mathcal{M}_2 & \xleftarrow{g_2} & F_n
\end{array}
\qquad
\begin{array}{ccc}
\mathcal{F}_{\mathcal{L},0} = \mathcal{F}_{\mathcal{M},0} & \xleftarrow{g_0} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i} & & \downarrow T_1 \\
\mathcal{F}_{\mathcal{M},1} & \xleftarrow{g_1} & \mathcal{F}_l \\
\downarrow \Sigma_{k,j} & & \downarrow T_2 \\
\mathcal{F}_{\mathcal{L},2} = \mathcal{F}_{\mathcal{M},2} & \xleftarrow{g_2} & \mathcal{F}_l
\end{array}$$

We have that $\tau_2\tau_1 = \sigma_2\sigma_1$. Moreover, we have that $l(\tau_2\tau_1) = l(\tau_2) + l(\tau_1)$,

since τ_2 and τ_1 are generated by disjoint sets of transpositions. So both $\tau_2\tau_1$ and $\sigma_2\sigma_1$ are reduced expressions, and after expanding into products of generators, $\tau_2\tau_1$ can be converted into $\sigma_2\sigma_1$ using just braiding and commutation relations. Thus $\Sigma_2\Sigma_1 = T_2T_1$ as before.

2. The expansion relation:

Consider the pair of diagrams defining the action of $C_{k+1,ni}C_{k+1,ni+1}\dots C_{k+1,ni+n-1}$, the Hecke version of $c_{k+1,ni}c_{k+1,ni+1}\dots c_{k+1,ni+n-1}$ (each section of this diagram could be expanded further using the definition of $c_{k+1,j}$ but we don't show this).

$$\begin{array}{ccc}
\mathcal{L} = \mathcal{L}_0 & \xleftarrow{f_0} & DX^l \\
\downarrow c_{k+1,ni+n-1} & & \downarrow \sigma_1 \\
\vdots & & \vdots \\
\downarrow c_{k+1,ni+1} & & \downarrow \sigma_{n-1} \\
\mathcal{L}_{n-1} = & \xleftarrow{f_{n-1}} & DX^l \\
\downarrow c_{k+1,ni} & & \downarrow \sigma_n \\
\mathcal{L}_n = & \xleftarrow{f_n} & DX^l
\end{array}
\qquad
\begin{array}{ccc}
\mathcal{F}_{\mathcal{L},0} & \xleftarrow{f_0} & \mathcal{F}_l \\
\downarrow C_{k+1,ni+n-1} & & \downarrow \Sigma_1 \\
\vdots & & \vdots \\
\downarrow C_{k+1,ni+1} & & \downarrow \Sigma_{n-1} \\
\mathcal{F}_{\mathcal{L},n-1} & \xleftarrow{f_{n-1}} & \mathcal{F}_l \\
\downarrow C_{k+1,ni} & & \downarrow \Sigma_n \\
\mathcal{F}_{\mathcal{L},n} & \xleftarrow{f_n} & \mathcal{F}_l
\end{array}$$

We compare this to:

$$\begin{array}{ccc}
\mathcal{L} = \mathcal{L}_0 & \xleftarrow{f_0} & DX^l \\
\downarrow \sigma_{k,i} & & \downarrow \tau_1 \\
\mathcal{L}_n & \xleftarrow{f_n} & DX^l
\end{array}
\qquad
\begin{array}{ccc}
\mathcal{F}_{\mathcal{L},0} & \xleftarrow{f_0} & \mathcal{F}_l \\
\downarrow \Sigma_{k,i} & & \downarrow T_1 \\
\mathcal{L}_n & \xleftarrow{f_n} & \mathcal{F}_l
\end{array}$$

As before, we will establish that

$$T_1 = \Sigma_n \dots \Sigma_2 \Sigma_1,$$

which we do by showing that $\sigma_n \dots \sigma_2 \sigma_1$ is reduced when each σ_i is expanded as a product of generators.

First we consider $\sigma_{k,i}$ and τ_1 . Write $v_{k,1}, v_{k,2}, \dots, v_{k,m}$ for the leaves of DX^k , where $m = dn^k$. Suppose that the leaves of \mathcal{L}_0 are l_1, l_2, \dots, l_m

listed left-to-right, with $l_a, l_{a+1}, \dots, l_{a+r-1}$ lying below $v_{k,i}$, and $l_{a+r}, l_{a+r+1}, \dots, l_{a+r+s-1}$ lying below $v_{k+1,i}$. Then acting on DX^l , the permutation τ_1 interchanges the two blocks $v_{l,a}, \dots, v_{l,a+r-1}$ and $v_{a+r}, \dots, v_{a+r+s-1}$, which means it has length rs .

Now we consider the various $c_{k+1,ni+b}$. It is useful to keep Figure 4.5 in mind for this part of the argument: the action of $\sigma_n \dots \sigma_1$ on DX^l follows that diagram except that each strand is replaced with some number of parallel strands (so that Remark 4.4.5 applies and would give the result here, but we did not prove that result formally so we do the full calculation here).

For each $1 \leq j \leq n$, let there be r_j leaves of \mathcal{L}_0 below $v_{k+1,(n-1)i+j}$ and let there be s_j leaves of \mathcal{L}_0 below $v_{k+1,ni+j}$. Then in particular, $r_1 + \dots + r_n = r$ and $s_1 + \dots + s_n = s$. We summarize this information as:

$$[r_1]_{ni-n+1}[r_2]_{ni-n+2} \dots [r_n]_{ni}[s_1]_{ni+1}[s_2]_{ni+2} \dots [s_n]_{ni+n},$$

where the notation $[x]_j$ means a block of x consecutive leaves below the leaf $v_{k+1,j}$. Each permutation $c_{k+1,j}$ then acts just by permuting the subscripts. In particular, $c_{k+1,ni+n-1}$ acts by the permutation

$$(v_{k+1,ni+n} v_{k+1,ni+n-1} \dots v_{k+1,ni})$$

on the vertices of DX^{k+1} , so sends the string of blocks above to:

$$[r_1]_{ni-n+1}[r_2]_{ni-n+2} \dots [r_{n-1}]_{ni-1}[r_n]_{ni+n}[s_1]_{ni+1}[s_2]_{ni+2} \dots [s_n]_{ni+n-1}.$$

The order of two leaves is only reversed by $c_{k+1,ni+n-1}$ if one lies in the $[r_n]$ -block and the other lies to the right, which tells us that the corresponding permutation σ_1 on DX^l has length $r_n(s_1 + s_2 + \dots + s_n)$ (and in fact, the terms $r_n s_i$ are the lengths of the permutations of DX^l corresponding to each transposition of $c_{k+1,ni+n-1}$).

Continuing in this way, $c_{k+1,ni+n-2}$ sends the sequence of blocks to:

$$[r_1]_{ni-n+1}[r_2]_{ni-n+2} \dots [r_{n-2}]_{ni-2}[r_{n-1}]_{ni+n-1}[r_n]_{ni+n}[s_1]_{ni}[s_2]_{ni+1} \dots [s_n]_{ni+n-2}.$$

This reverses the order of any pair of leaves where one lies in the $[r_{n-1}]$ -block and one lies in any $[s_i]$ -block. This shows us that the corresponding

permutation of DX^l has length $r_{n-1}(s_1 + s_2 + \dots + s_n)$.

Continuing in this way (with Figure 4.5 as our guide) we see that each permutation $c_{k+1,ni+j}$ (for $0 \leq j \leq n-1$) yields a permutation of DX^l of length $r_{j+1}(s_1 + s_2 + \dots + s_n)$. The total length of all these permutations is then $(r_1 + \dots + r_n)(s_1 + \dots + s_n) = rs$, which is also the length of τ_1 . So we see that $\sigma_n \dots \sigma_1$ and τ_1 must be two different reduced expressions for the same permutation of DX^l and so $T_1 = \Sigma_n \dots \Sigma_1$.

3. **For $\sigma \in \mathfrak{S}_{\text{hom}}$, then Σ is well-defined:** Examining the proof of Proposition 4.5.9, we see that it only uses the braiding, commutation and expansion relations among the generators $\Sigma_{k,i}$. We've already shown that these relations hold even in the action of the generators on the larger set \mathcal{F}_V , so this holds also.
4. **The action of F :** Suppose that $\sigma_{k,i}f = g\sigma$, and consider the action of $\sigma_{k,i}f$ on a sufficiently deep leaf set \mathcal{L} . We wish to show that $\Sigma_{k,i}f$ and $g\Sigma$ act on \mathcal{L} in the same way. We have the diagram:

$$\begin{array}{ccccc} \mathcal{L}_1 & \xleftarrow{f} & \mathcal{L} = \mathcal{L}_0 & \xleftarrow{f_0} & DX^k \\ \downarrow \sigma_{k,i} & & \downarrow \sigma & & \downarrow \sigma_0 \\ \mathcal{L}_3 & \xleftarrow{g} & \mathcal{L}_2 & \xleftarrow{f_1} & DX^k \end{array}$$

Observe that ff_0 and gf_1 are the unique elements of F sending DX^k to \mathcal{L}_1 and \mathcal{L}_3 respectively. Thus we have the defining diagrams for $\Sigma_{k,i}$ and for Σ :

$$\begin{array}{ccc} \mathcal{F}_{\mathcal{L},1} & \xleftarrow{ff_0} & DX^k \\ \downarrow \Sigma_{k,i} & & \downarrow \Sigma_0 \\ \mathcal{F}_{\mathcal{L},3} & \xleftarrow{gf_1} & DX^k \end{array} \qquad \begin{array}{ccc} \mathcal{F}_{\mathcal{L},0} & \xleftarrow{f_0} & DX^k \\ \downarrow \Sigma & & \downarrow \Sigma_0 \\ \mathcal{F}_{\mathcal{L},2} & \xleftarrow{f_1} & DX^k \end{array}$$

In other words, we have:

$$\Sigma_{k,i} = gf_1 \Sigma_0 f_0^{-1} f^{-1}$$

in its action on $\mathcal{F}_{\mathcal{L},1}$, and

$$\Sigma = f_1 \Sigma_0 f_0^{-1}$$

in its action on $\mathcal{F}_{\mathcal{L},0} = f^{-1} \mathcal{F}_{\mathcal{L},1}$. It follows that $\Sigma_{k,i} f = g \Sigma$ on \mathcal{L} , as required.

5. **The existence of inverses:** Consider the action of $\Sigma_{k,i}$ on a leaf set \mathcal{L} , which is defined by

$$\Sigma_{k,i} = f' \Sigma f^{-1},$$

where we have a diagram:

$$\begin{array}{ccc} \mathcal{F}_{\mathcal{L}} & \xleftarrow{f} & \mathcal{F}_l \\ \downarrow \Sigma_{k,i} & & \downarrow \Sigma \\ \mathcal{F}_{\mathcal{L}'} & \xleftarrow{f'} & \mathcal{F}_l \end{array}$$

Then Σ is invertible in its action on \mathcal{F}_l , since it is given as a product of generators of $\mathcal{H}_{l,q}$ which are invertible. So we define the action of $\Sigma_{k,i}^{-1}$ on $\mathcal{F}_{\mathcal{L}'}$ to be

$$\Sigma_{k,i}^{-1} = f \Sigma^{-1} (f')^{-1}.$$

This gives a pair of mutually inverse maps $\Sigma_{k,i} : \mathcal{F}_{\mathcal{L}} \rightarrow \mathcal{F}_{\mathcal{L}'}$ and $\Sigma_{k,i}^{-1} : \mathcal{F}_{\mathcal{L}'} \rightarrow \mathcal{F}_{\mathcal{L}}$. As \mathcal{L} varies over all possible leaf sets, so does \mathcal{L}' , so $\Sigma_{k,i}^{-1}$ is defined on all sufficiently deep leaf sets. Moreover, $\Sigma_{k,i}^{-1}$ is easily seen to map null elements of $\mathcal{F}_{V,\times}$ in its domain to other null elements (by mapping M -null elements to $\sigma_{k,i}(M)$ -null elements), so it descends to a well-defined endomorphism of \mathcal{F}_V . This completes the proof. □

This completes the construction of the algebra $\mathcal{H}_{V,q}^0$.

4.6 Comparing $\mathcal{H}_{V,q}$ and $\mathcal{H}_{V,q}^0$.

We have now created two objects called $\mathcal{H}_{V,q}$ and $\mathcal{H}_{V,q}^0$ related to V and Hecke algebras. To summarize:

- $\mathcal{H}_{V,q}$ is the group algebra $\mathbf{k}B_V$ equipped with a partial action. In this action, the hom spaces are finite dimensional, and the subspaces $\mathbf{k}B_k \subset$

$\mathbf{k}B_{\text{hom}}$ act on DX^k via the quotient to the Hecke algebra $\mathcal{H}_{n,q}$. It is defined for all q .

- $\mathcal{H}_{V,q}^0$ is formed from finite Hecke algebras $\mathcal{H}_{m,q}$ acting on flag spaces $\mathcal{F}_k = \mathcal{F}_{DX^k}$ (for $m = dn^k$). The construction extends the action of the symbols $\Sigma_{k,i}$ to all $\mathcal{F}_{\mathcal{L}}$ where \mathcal{L} is a leaf set below DX^k . This construction requires q to be the order of a finite field.

We now relate these constructions. The key to relating them is that $\mathcal{H}_{V,q}^0$ naturally has a partial action, because we defined it as an algebra of endomorphisms of all sufficiently deep spaces $\mathcal{F}_{\mathcal{L}}$. Moreover, the action of generator $\Sigma_{k,i}$ on the space $\mathcal{F}_{\mathcal{L}}$ is defined by a formula:

$$\Sigma_{k,i}|_{\mathcal{F}_{\mathcal{L}}} = f'\Sigma f^{-1},$$

where we have a commutative diagram

$$\begin{array}{ccc} \mathcal{F}_{\mathcal{L}} & \xleftarrow{f} & \mathcal{F}_l \\ \downarrow \Sigma_{k,i} & & \downarrow \Sigma \\ \mathcal{F}_{\mathcal{L}'} & \xleftarrow{f'} & \mathcal{F}_l. \end{array}$$

In this diagram, Σ acts on \mathcal{F}_l by an element of the finite-dimensional Hecke algebra $\mathcal{H}_{dn^l,q}$. Thus the action of $\Sigma_{k,i}$ on \mathcal{L} lies in $f'\mathcal{H}_{dn^k,q}f^{-1}$ (with the natural action of these morphisms). We also notice that the $\Sigma_{k,i}$ satisfy the defining relations of a braid group, by Theorem 4.5.10. These facts combine into the following theorem.

Theorem 4.6.1. *There is an algebra homomorphism θ from $\mathbf{k}B_V$ to $\mathcal{H}_{V,q}^0$, defined on the generators by sending $\Sigma_{k,i} \in \mathbf{k}B_V$ to $\Sigma_{k,i} \in \mathcal{H}_{V,q}^0$ and preserving $f \in F_{n,d}$. Moreover, if $X \in \mathbf{k}B_V$ and $S(X)$ contains $f'Yf^{-1} \in \text{Hom}(\mathcal{L}, \mathcal{L}') = f'\mathcal{H}_{dn^k,q}f^{-1}$, then $\theta(X)$ acts on \mathcal{L} by the morphism $f'Yf^{-1}$.*

Proof. We have argued that the algebra quotient exists. All we have left to check is that if $S(X)$ contains $f'\Phi f^{-1} \in \text{Hom}(\mathcal{L}, \mathcal{L}')$ (for $\Phi \in \mathcal{H}_{dn^k,q}$), then $S(\theta(X))$ contains $f'\Phi f^{-1}$ also. Suppose that $X(\mathcal{L})$ is defined for $X \in \mathbf{k}B_V$ and leaf set \mathcal{L} . Notice that X is only defined on \mathcal{L} if all its monomials are, so we can just consider X in the braid group B_V (rather than $\mathbf{k}B_V$). Then X is defined on \mathcal{L} if and only if we can write X with domain leaf set \mathcal{L} and range leaf set \mathcal{L}' say; then we can factor X in the braid group as $f'\Sigma f^{-1}$ where $Y \in B_{\text{hom}}$ and

$f, f' \in F$, mapping DX^k to $\mathcal{L}, \mathcal{L}'$ respectively. This same factorization defines an action of $\theta(X)$ on $\mathcal{F}_{\mathcal{L}}$. \square

In summary, we have constructed a Hecke algebra $\mathcal{H}_{V,q}$ which is an algebra with a partial action, and (for q the order of a finite field) an algebra $\mathcal{H}_{V,q}^0$ which as an algebra is a quotient of $\mathcal{H}_{V,q}$ by a map θ , and which has a partial action where $S(\theta(X))$ contains $S(X)$ for all $X \in \mathcal{H}_{V,q}$. It seems likely that more is true.

Conjecture 1. *The homomorphism θ is injective, so that $\mathcal{H}_{V,q}^0$ is isomorphic to a group algebra for B_V .*

We suggest why this conjecture seems to hold. If it fails, there is a non-zero $X \in \mathbf{k}B_V$ which when viewed as an element of $\mathcal{H}_{V,q}^0$ acts as the zero map on all sufficiently deep spaces $\mathcal{F}_{\mathcal{L}}$. We will focus on $X \in \mathbf{k}B_{\text{hom}}$ so that X maps leaf sets DX^k to themselves at least, and we write X as a (non-commutative) polynomial $X = p(\Sigma_{k,1}, \dots, \Sigma_{k,dn^k-1})$ in the generators $\Sigma_{k,i}$ of $\mathcal{H}_{dn^k,q}$. We need this polynomial to not only vanish in $\mathcal{H}_{k,q}$, but also vanish when expanded to any $\mathcal{H}_{k',q}$ for $k' > k$. Moreover X should contain some non-reduced word in the $\Sigma_{k,i}$ (or we could just study it in terms of braid groups) and so we can look at $\Sigma_{k,i}^2$, which appears in X somewhere. In $\mathcal{H}_{k,q}$, $\Sigma_{k,i}^2$ can be rewritten as $(q-1)\Sigma_{k,i} + q$, but in $\mathcal{H}_{k',q}$ this square becomes a sum of many more terms.

For example, consider $V = V_{2,1}$ (so $d = 1, n = 2$), and look at $\Sigma_{1,1}$ as the unique generator of the 2-dimensional Hecke algebra $\mathcal{H}_{2^1,q}$. Acting on \mathcal{F}_1 , this satisfies the relation:

$$\Sigma_{1,1}^2 = (q-1)\Sigma_{1,1} + q.$$

When we extend the action of $\Sigma_{1,1}$ to \mathcal{F}_2 , it acts by the Hecke element $\Sigma_{2,2}\Sigma_{2,1}\Sigma_{2,3}\Sigma_{2,2}$. But this element of $\mathcal{H}_{2^2,q}$ does not satisfy any quadratic relations (and in fact its minimal polynomial has degree 7). So $\Sigma_{1,1}^2 - (q-1)\Sigma_{1,1} - q$ is a non-zero element of $\mathcal{H}_{V,q}^0$. Our conjecture is saying that this sort of phenomenon always happens. This seems plausible, but difficult to prove: I have no good plan for how to consider all possible relations in any possible expansion.

4.6.1 Evaluating the construction

Recall that in Section 4.3 we identified four reasonable strategies for building a Hecke algebra - or in other words, four similarities between symmetric groups and Hecke algebras of type A_n . We now give a brief overview of evaluate how well $\mathcal{H}_{V,q}$ compares with these strategies.

Adding relations to a braid group

This is essentially what we did in defining $\mathcal{H}_{V,q}$. We did not take a quotient of the group itself, but took a quotient of a category on which it acted. Notice that if we set $q = 1$ (which would in the classical setting give a group algebra $\mathbf{k}\mathfrak{S}_n$) we still get an action of a braided group, rather than getting Thompson's group V . However, the square of each generator then acts by the identity on each leaf set where it is defined. So by quotienting out elements with a trivial action, we arrive back at Thompson's group V .

Studying endomorphisms of a general linear group module

We ask if $\mathcal{H}_{V,q}$ can be seen as endomorphisms of a module for some version of the general linear group. This was the purpose of defining $\mathcal{H}_{V,q}^0$. Elements of $\mathcal{H}_{V,q}^0$ were defined by an action on spaces $\mathcal{F}_{\mathcal{L}}$. The space $\mathcal{F}_{\mathcal{L}}$ was defined as the \mathbf{k} -space of flags in the \mathbb{F} -vector space $W_{\mathcal{L}}$ whose basis is labelled by \mathcal{L} . Thus, $\mathcal{F}_{\mathcal{L}}$ is a $GL_{\mathcal{L}}$ -module (where $GL_{\mathcal{L}}$ is the set of linear maps on $W_{\mathcal{L}}$).

The action of F permutes the spaces $\mathcal{F}_{\mathcal{L}}$, and general elements of $\mathcal{H}_{V,q}^0$ both permute them and act by endomorphisms on the spaces being permuted, as described in Theorem 4.5.10. Indeed, we write the action of $\Sigma_{k,i}$ on a leaf set $\mathcal{F}_{\mathcal{L}}$ as $f'\Sigma f^{-1}$ for $f, f' \in F$ and Σ an endomorphism of \mathcal{F}_l . It may be useful to rewrite this as $(f'\Sigma(f')^{-1})(f'f^{-1})$, where $f'\Sigma(f')^{-1}$ is a $GL_{\mathcal{L}'}$ -endomorphism of $\mathcal{F}_{\mathcal{L}'}$ and $f'f^{-1} \in F$. So the action of elements of $\mathcal{H}_{V,q}^0$ on flag spaces $\mathcal{F}_{\mathcal{L}}$ can be written as a product of an element of F and a $GL_{\mathcal{L}}$ -endomorphism of some flag space.

Overall, we can define a group $GL_{V,\times}$ as the product of all $GL_{\mathcal{L}}$ as the leaf set \mathcal{L} varies, and define a subgroup $GL_{V,+}$ as elements of $GL_{V,\times}$ that are the identity on all leaf sets below some M (the group analogue of what we called M -null vectors). Then $\mathcal{H}_{V,q}^0$ acts by $GL_{V,\times}/GL_{V,+}$ endomorphisms. This means that we can understand $\mathcal{H}_{V,q}^0$ as endomorphisms of some sort of general linear group, either in the local or the global action of $\mathcal{H}_{V,q}^0$.

Whilst this is nice to have, this is not a particularly satisfying way to describe $\mathcal{H}_{V,q}$ in terms of endomorphisms. For example, it does not seem to come close to containing every $GL_{V,\times}/GL_{V,+}$ endomorphism, so this is quite a weak description. It would be better if we somehow could take into account that if $X \in \mathcal{H}_{V,q}^0$ acts by endomorphism θ on $\mathcal{F}_{\mathcal{L}}$, then it tends to by an expanded version of θ on the expansion $\mathcal{F}_{\mathcal{L}+}$. However, this isn't always true: for example, $\Sigma^2 - (q-1)\Sigma - q$ might act as the zero map on some $\mathcal{F}_{\mathcal{L}}$ but not act as the zero

map on expanding. This sort of problem is why we have ended up working with the product of the spaces $\mathcal{F}_{\mathcal{L}}$, unable to put any more relations between them. This also means we can't work with a smaller linear group than $GL_{V,\times}/GL_{V,+}$ - we really do need to treat all the leaf sets separately.

In Section 4.3, we tried to construct a Hecke algebra (for $V = V_{2,1}$) using a group $GL(\Gamma, \mathbb{F})$. Here Γ was the graph with one vertex and two loops, and $GL(\Gamma, \mathbb{F})$ was a group of invertible elements of $L_{\mathbb{F}}(\Gamma)$, generated by matrix groups for each leaf set, and we defined flags for this group, but it wasn't possible to build a Hecke algebra from this construction. Instead, the algebra $\mathcal{H}_{V,q}$ requires separate flag spaces for different leaf sets, rather than one overarching idea of flags.

Deforming a presentation of V

This is maybe a less helpful way to think about $\mathcal{H}_{V,q}$, because we haven't defined what is meant by a presentation for a group with partial \mathcal{C} -action. But certainly, we can regard $\mathcal{H}_{V,q}$ as a deformation of V with its partial action. Here let \mathcal{C} be the category with objects labelled by leaf sets of size dn^k ; we take $\text{Hom}(DX^k, DX^k)$ to be $\mathbf{k}\mathfrak{S}_{dn^k}$, and for general L, L' leaf sets of the same size dn^k , we put $\text{Hom}(L, L') = f'\text{Hom}(DX^k, DX^k)f^{-1}$, where f, f' are the unique elements of F mapping DX^k to L, L' . Then $\mathcal{H}_{V,q}$ can be regarded as a deformation of the group algebra of V with its partial action on \mathcal{C} , where we replace each $\mathbf{k}\mathfrak{S}_{dn^k}$ with $\mathcal{H}_{dn^k,q}$.

Fitting together finite Hecke algebras

This is a reasonable description for $\mathcal{H}_{V,q}^0$. As far as we can tell, one cannot embed $\mathcal{H}_{k,q}$ in $\mathcal{H}_{nk,q}$ by expanding generators. Recall how we considered the particular case of trying to embed $\mathcal{H}_{2,q}$ into $\mathcal{H}_{4,q}$ (for $V = V_{2,1}$). One would seek $\Sigma \in \mathcal{H}_{4,q}$, probably of the form $\Sigma_{2,2}\Sigma_{2,1}\Sigma_{2,3}\Sigma_{2,2}$ plus smaller terms, satisfying the quadratic relation $\Sigma^2 = (q-1)\Sigma + q$. This Σ would be an expansion of $\Sigma_{2,1} \in \mathcal{H}_{1,q}$. After extensive calculations, it seems that such Σ do not exist, so we can't embed $\mathcal{H}_{k,q}$ into $\mathcal{H}_{k+1,q}$.

We've done the next best thing, by defining endomorphisms which are equal to a generator on one layer and an expanded product of generators on the next, which are our $\Sigma_{n,i}$. We've used these to generate the Hecke algebra. This is perhaps a less good description of $\mathcal{H}_{V,q}$, where we have a braid group algebra rather than finite-dimensional Hecke algebras.

4.7 Properties of the Hecke algebra construction

Here we show some possibly desirable properties of the Hecke algebra we've constructed. In particular, we show that it is finitely generated (as an algebra with partial action). We also describe a family of representations of V that are the $q = 0$ case of representations of $\mathcal{H}_{V,q}$. These representations will not be irreducible, but are the closest (it seems) one can come to defining representations of V in terms of the irreducible representations of each \mathfrak{S}_n .

4.7.1 A family of representations for $\mathcal{H}_{V,q}$

In this section we sketch some representations for $\mathcal{H}_{V,q}$ that in the case $q = 1$ give representations of V . Recall that \mathfrak{S}_m has \mathbf{k} representations V^λ for each $\lambda \vdash m$ (the Specht modules), which generalize to modules also called V^λ for the Hecke algebra $\mathcal{H}_{m,q}$. These modules are irreducible if the Hecke algebra is semisimple (see eg [32]).

We use these modules to define representations for $\mathcal{H}_{V,q}$ or for $V = V_{n,d}$ as follows. We look at $\mathcal{H}_{V,q}$ as just the braid group algebra, and look for representations that locally quotient through homomorphisms to the finite-dimensional Hecke algebras. For each $N = dn^k$ (for $k \in \mathbb{N}$), choose $\lambda_k \vdash N$. We take a space W_k isomorphic to V^{λ_k} for each k , which is a module for $\mathcal{H}_{dn^k,q}$. Moreover, for each leaf set \mathcal{L} of size dn^k , we take an isomorphic copy $W_{\mathcal{L}}$ of W_k , and specify that the isomorphism is by $f : W_k \rightarrow W_{\mathcal{L}}$, where f is the unique element of f mapping DX^k to \mathcal{L} . For general $X \in B_V$, if we write $X = f_2 X' f_1^{-1}$ (for $f_i : DX^k \rightarrow \mathcal{L}_i$ and $X_i \in B_{dn^k}$) then we define the action of X on $W_{\mathcal{L}}$ by this same formula, where we take X' acting through the quotient to $\mathcal{H}_{dn^k,q}$ on W_k . This makes X into a map from $W_{\mathcal{L}}$ to $W_{\mathcal{L}'}$. So each element of B_V acts on all sufficiently deep $W_{\mathcal{L}}$ (via a quotient to the Hecke algebra) and the space $\coprod W_{\mathcal{L}} / \oplus W_{\mathcal{L}}$ is a representation. We call it W_Λ , where $\Lambda = (\lambda_1, \lambda_2, \lambda_3, \dots)$. This is a representation, where we can think of each subgroup $B_{\mathcal{L}}$ as acting via a Hecke algebra quotient, as we wanted.

However, we remark that this construction almost never gives an irreducible representation W_Λ of V . Indeed, the representation is defined by choosing a representation W^λ of \mathfrak{S}_m for each $m = dn^k$. Each space $\langle W_{\mathcal{L}} : |\mathcal{L}| = dn^k \rangle$ is acted on individually by all sufficiently high elements of V . This means that given any subset S of $\{dn^k : k \in \mathbb{N}\}$, we can drop all $W_{\mathcal{L}}$ from W_Λ where

$|\mathcal{L}| \in S$. This will still give a representation of V , and moreover it will be a subrepresentation of W_Λ . It is a proper subrepresentation if any of the dropped $W_{\mathcal{L}}$ has dimension greater than 1. So W_Λ can only possibly be irreducible if it is a linear representation, and we gain nothing new here.

4.7.2 Extending the construction to colour-preserving Thompson groups

It is natural to ask if this construction will work for colour preserving Thompson groups $V_{\mathbf{C},\mathbf{S}}$. We explain why this seems to fail (and so it's even less possible to define a Hecke algebra here than in the case of V alone).

First we remark that a braided version $B_{\mathbf{C},\mathbf{S}}$ of the colour-preserving Thompson group clearly exists, where we just use colour-preserving braided diagrams. We can also define its partial action on the standard braided category exactly as before. We could also define $F_{\mathbf{C},\mathbf{S}}$ to be the group of all order-preserving elements of $V_{\mathbf{C},\mathbf{S}}$. However, this $F_{\mathbf{C},\mathbf{S}}$ is not as nice as in the single-coloured case, because there isn't always an element of $F_{\mathbf{C},\mathbf{S}}$ bijecting given two leaf sets with the same size and set of colours; the colours can appear in a different order. So we won't be able in general to get a factorization resembling $V = \mathfrak{S}_{\text{hom}}F$. Instead, we look at the other way to define a Hecke algebra, by giving $B_{\mathbf{C},\mathbf{S}}$ a partial action.

Suppose we are trying to do this, which means that we want to form quotients of $\mathbf{k}B_{\mathcal{L}}$, where $B_{\mathcal{L}}$ is the group of braided, colour-preserving permutations of some leaf set. The next problem we have is what the equivalent of a finite-dimensional Hecke algebra should be. Consider trying to add relations to a colour-preserving subgroup of a braid group. This is a finitely presented group, since it's a finite index subgroup of the braid group (there are finitely many permutations of the colours). In principle, one could find a presentation for this group. However, it's not obvious what that presentation is, and after finding it, we would have to decide on sensible extra relations to add to it to get a Hecke algebra. Since the presentation is more complicated than for the braid group, it seems very hard to define quadratic relations to be added. This could be worth investigating, to find if there is a sensible quotient here, but it doesn't seem to exist. So there doesn't seem to be a way of defining Hecke algebras in the greater generality.

4.7.3 Finite generation of $\mathcal{H}_{V,q}$

Finally we wish to argue that $\mathcal{H}_{V,q}$ is finitely generated. This will be finite generation as an algebra with partial action, which means that we must be able to deduce the action of every element of $\mathcal{H}_{V,q}$ from the action of the generators. The work here will basically be a careful reproof of the fact that V is finitely generated to show this holds.

First we show that F with its partial action on $\mathcal{C}_{n,d}$ is finitely generated. Here we don't need to worry about the full structure of the hom sets in $\mathcal{C}_{n,d}$, because only one morphism of each hom set could possibly come from an element of F . So we just consider the partial action on leaf sets.

Lemma 4.7.1. *Let $F = F_{n,d}$ be a Higman-Thompson group. Then F satisfies the following finite generation property in its action on leaf sets: there is a finite subset F_{gen} of F such that whenever $f \in F$ is defined on the leaf set \mathcal{L} , we can write $f = f_1 f_2 f_3 \dots f_m$, for $f_i \in F_{\text{gen}}$, such that f_i is defined on the leaf set $f_{i+1} \dots f_m(\mathcal{L})$ for all $1 \leq i \leq m-1$, and f_m is defined on \mathcal{L} .*

Since this proof is essentially reshewing that F is finitely generated, it is written somewhat briefly in places.

Proof. Let \mathcal{T} be the d -rooted n -regular forest, so that F acts on the ends of \mathcal{T} . We can assume (by removing the root of \mathcal{T} to form an isomorphic group) that $d > 1$. Let vertices of \mathcal{T} be represented by strings DX^* as is normal, where $D = \{r_1, r_2, \dots, r_d\}$ (the notation r because these are the roots) and where $X = \{x_1, x_2, \dots, x_n\}$, read left to right.

For each $1 \leq k \leq d$, define a leaf set

$$L_{0,k} = \{r_i : i \neq k\} \cup \{r_k x_i : 1 \leq i \leq n\}.$$

Define $X_{0,k}$ to be the element of F that maps $L_{0,k}$ to $L_{0,k+1}$, for each $1 \leq k \leq d-1$. Also define, for all $m \in \mathbb{N}$ and $1 \leq k \leq d$,

$$L_{1,k} = \{r_i : i \neq d\} \cup \{r_d x_i : i \neq k\} \cup \{r_d x_k x_i : 1 \leq i \leq n\}.$$

Define $X_{1,k}$ to be the element of F that maps $L_{1,k}$ to $L_{1,k+1}$. Take F_{gen} to be the set of all elements $X_{0,k}$ and $X_{1,k}$ as well as their inverses. Finally, for $N > 1$, define

$$L_{N,k} = \{r_i : i \neq d\} \cup \{r_d x_d^l x_i : i \neq d, 0 \leq l < N\} \setminus \{r_d x_d^{N-1} x_k\} \cup \{r_d x_d^{N-1} x_k x_i : 1 \leq i \leq n\},$$

and we define $X_{N,k}$ to be the element of F mapping $L_{N,k}$ to $L_{N,k+1}$. We draw $L_{2,3}$ for the case $d = 3, n = 4$ to give an example.

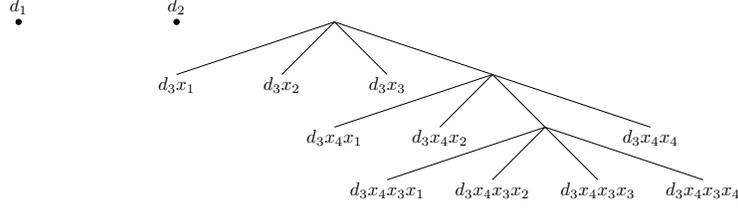


Figure 4.8: An example leaf set $L_{2,3}$

It is easy to see that $L_{N,k}$ is a minimal leaf set on which $X_{N,k}$ is defined. Moreover, for $n \geq 1$, one can verify that

$$X_{0,n-1}X_{N,k}X_{0,n-1}^{-1} = X_{N+1,k}.$$

Also, the expressions on both sides of this formula act on $L_{N+1,k}$ (and hence on all lower leaf sets) : $X_{0,n-1}^{-1}$ is defined on $L_{N+1,k}$, $X_{N,k}$ acts on $X_{0,n-1}^{-1}(L_{N+1,k})$, and $X_{0,n-1}$ acts on $X_{N,k}X_{0,n-1}^{-1}(L_{N+1,k})$. This tells us that all $X_{N,k}$ are in the subgroup generated by the $X_{N,k}$ for $N = 0, 1$, along with their partial action.

Now consider general $f \in F$, defined by a bijection $L \rightarrow L'$. Let N be the unique value such that $L_{N,n}$ is in bijection with L , and write $f = f_1f_2$ for $f_1 : L_{N,n} \rightarrow L'$ and $f_2 : L \rightarrow L_{N,n}$. We show that f_1 can be written as a product of terms $X_{m,i}$ such that the product acts on L : that is, we show that L can be converted into $L_{N,n}$ by acting on it with various $X_{m,i}$. Indeed, if $L \neq L_{N,n}$ then it lies below some $L_{m,i}$ for $m \leq N$ and $i < n$. Then acting on L by $X_{m,i}$, the leaves of L are moved rightwards. This process must terminate by eventually converting L into $L_{N,n}$. So f_1 , with all of its partial action, is in the subgroup generated by the $X_{k,i}$ for $k = 0, 1$. Similarly f_2^{-1} , with its action on L' , lies in this subgroup, so f , with its action $L \rightarrow L'$ does as well. This completes the proof. \square

We now extend this result to $\mathbf{k}B_V$.

Theorem 4.7.2. *B_V satisfies the following finite generation property in its action on the standard braided category $\mathcal{C}_{n,d}$: there is a finite subset B_{gen} of B_V such that whenever $\Sigma \in B_V$ and \mathcal{L} is a leaf set with $\Sigma(\mathcal{L})$ defined, we can write $\Sigma = \Sigma_1\Sigma_2\Sigma_3 \dots \Sigma_m$ for $\Sigma_i \in B_{gen}$. Moreover there exist leaf sets $\mathcal{L}_0, \mathcal{L}_1, \dots, \mathcal{L}_m = \mathcal{L}$ such that $S(\Sigma_i)$ contains a morphism $\mathcal{L}_i \rightarrow \mathcal{L}_{i-1}$ (so that*

$\Sigma_1 \Sigma_2 \dots \Sigma_m$ acts on \mathcal{L}). This finite generation property descends to the quotient partial action $\mathcal{H}_{V,q}$.

Again, this proof is essentially a proof that B_V is finitely generated, but with some care taken as to what leaf sets each product is defined on.

Proof. It is clear that finite generation is preserved in quotients (of groups with partial action) so we just need to study the partial action on $\mathcal{C}_{n,d}$. We already have that F is finitely generated with its partial action on all leaf sets, and it clearly stays finitely generated when we restrict to the leaf sets of size dn^k (since there are no morphisms between leaf sets of different size).

In this proof, we will take B_{gen} to consist of a generating set F_{gen} for F together with finite generating sets for the homogeneous subgroups B_d, B_{nd}, B_{n^2d} of B_{hom} (that act on depth 0, depth 1 and depth 2 leaves respectively).

Now consider general $\Sigma \in B_V$, and leaf set \mathcal{L} (of size dn^k) where $\Sigma(\mathcal{L})$ is defined. We can factor $\Sigma = f'\phi f^{-1}$ for $\phi \in B_{\text{hom}}$, and $f, f' \in F$. In this expression, f^{-1} maps \mathcal{L} to DX^k , and f' maps DX^k to \mathcal{L}' , the image of \mathcal{L} under Σ . Since we have already shown finite generation for F , it's enough to write ϕ as a product of elements of B_{gen} such that the product is defined on DX^k . That is, it's enough to show that each $B_{dn^k} \subset B_{\text{hom}}$ is in the subgroup generated by B_{gen} , with its action on DX^k . We will establish this for the standard generators of B_{dn^k} (those whose braid diagrams, as in Figure 4.1, have a single crossing, left-over-right). This is enough to imply the result.

Let τ be such a generator for B_{dn^k} . It is easy to show that there exists $f \in F$ such that $f\tau f^{-1}$ lies in B_d, B_{dn} or B_{dn^2} (we just need to define f by a bijection between leaf sets that takes the two leaves moved by τ to two leaves of level 0, 1 or 2). Then $\tau' = f\tau f^{-1}$ lies in B_{gen} and it acts on DX^i for some $i = 0, 1, 2$, so certainly τ' acts on $f(DX^k)$ which is a leaf set lying below DX^i . Thus, $\tau = f^{-1}\tau'f$ gives rise to an expression for τ as a product of generators, which is defined on DX^k and any lower leaf set - as required. \square

This finite generation property can be extended linearly to $\mathbf{k}B_V$. One could go on to study higher finiteness properties of $\mathcal{H}_{n,q}$ such as being finitely presented, after defining what this means for an algebra with partial action, but we won't pursue this any further.

This concludes our study of the Hecke algebras $\mathcal{H}_{V,q}$. Ultimately, this study has not led to much progress with the representation theory of V . One might conclude by saying that it is just good to know that deforming V to form a Hecke

algebra does not work, and so other people do not need to try this. Nevertheless, the constructions we have made along the way seem interesting.

Bibliography

- [1] G. Abrams. Leavitt Path Algebras: The First Decade. *Bull. Math. Sci.*, **5**(1):59-120 (2015).
- [2] G. Abrams, G. Aranda-Pino. The Leavitt path algebra of a graph. *J. Algebra* **293**:319-334 (2005).
- [3] G. Abrams, P. Ara, M. Siles Molina. Leavitt Path Algebras. Lecture Notes in Mathematics Series, Springer-Verlag; ebook from www.springer.com/gp/book/9781447173434
- [4] C.J.Ash, T.E.Hall. Inverse semigroups on graphs. *Semigroup Forum* **11** 140-145 (1975).
- [5] G. Aranda Pino, J. Clark, A. an Huef, I. Raeburn. Kumjian-Pask algebras of higher rank graphs. *Trans. Am. Maths. Soc.* **365**, 3613-3641 (2013)
- [6] P. Ara, K. Goodearl, E. Pardo. Nonstable K -theory for graph algebras. *Algebra. Representation Theory* **10**(2): 157-178 (2007).
- [7] P. Abramenko, K. Brown. Buildings. Theory and Applications. *Graduate Texts in Mathematics* **248** , Springer (2008).
- [8] H. Bass. Covering theory for graphs of groups. *Journal of Pure and Applied Algebra* **89** 3-47 (1993).
- [9] L. Bartholdi, R. I. Grigorchuk. On parabolic subgroups and Hecke algebras of some fractal groups. *Serdica. Math. J.* **28**, 47-90 (2002).

- [10] L. Bartholdi, R. Grigorchuk, Z. Sunik. Branch Groups. *Handbook of algebra* **3** 989-1112, North-Holland, Amsterdam (2003).
- [11] B. Bekka, M. Kalantar. Quasi-regular representations of discrete groups and associated C^* -algebras. arXiv: 1903.00202 (2019).
- [12] R. Brown. From groups to groupoids: a brief survey. *Bull. London Math. Soc.* **19**: 113-34 (1987).
- [13] N. Brownlowe, A. Munday, D. Pask, J. Spielberg, A. Thomas. C^* -algebras associated to graphs of groups. *Adv. Math.*, 316:114-186 (2017).
- [14] Collin Bleak, Martin Quick: The infinite simple group V of Richard J Thompson: presentation by permutations
- [15] Matthew G. Brin. The algebra of strand splitting. I. A braided version of Thompson's group V . *J. Group Theory* **10**(6) 757-788 (2007).
- [16] Matthew G. Brin. The algebra of strand splitting. II. A presentation for the braid group on one strand. *Internat. J. Algebra Comput.* **16**(1) 203-219 (2006).
- [17] J. Cannon, W. Floyd, W. Parry. Introductory notes on Richard Thompson's groups. *Enseign. Math.* (2) **42** no. 3-4, 215-25, (1996).
- [18] L.O. Clark, C. Edie-Michell. Uniqueness theorems for Steinberg algebras. *Algebr. Represent. Theor.* **18** (2015) 907-916.
- [19] L.O. Clark, R. Exel, E. Pardo. A generalised uniqueness theorem and the graded ideal structure of Steinberg algebras. *Forum Mathematicum* DOI:10.1515/forum-2016-0197.
- [20] L.O. Clark, C. Farthing. A. Sims. M. Tomforde. A groupoid generalisation of Leavitt path algebras. *Semigroup Forum* **89**, 501-517 (2014).
- [21] L.O. Clark, R. Hazrat. Étale groupoids and Steinberg algebras: a concise introduction. arXiv preprint, 1901.01612 (2019).
- [22] L.O. Clark, D. Martín Barquero, C. Martín González, M. Siles Molina. Using Steinberg algebras to study decomposability of Leavitt path algebras. *Forum Math.*, DOI 10.1515/forum-2016-0062 (2016).

- [23] K. Conrad. $SL_2(\mathbb{Z})$. Expository papers, on personal website <https://kconrad.math.uconn.edu>
- [24] C. W. Curtis, I. Reiner. Methods of Representation Theory with Applications to Finite Groups and Orders, volume 1. Wiley Classics Library. (1990)
- [25] J. Cuntz. Simple C^* -algebras generated by isometries. *Comm. Math. Phys.* **57**(2), 173-185 (1977)
- [26] A. Dudko, R. Grigorchuk. On irreducibility and disjointness of Koopman and quasi-regular representations of weakly branch groups. *Modern theory of dynamical systems*, Amer. Math. Soc, Providence RI, 51-66 (2017).
- [27] A. Dudko, K. Medynets. Finite factor representations of Higman-Thompson groups. *Groups Geom. Dyn.* **8**(2) 375-389 (2014).
- [28] A. Dudko, K. Medynets. On characters of inductive limits of symmetric groups. *J. Funct. Anal.* **264** no. 7, 1565-1598 (2013).
- [29] Artem Dudko: On irreducibility of Koopman representations of Higman-Thompson groups. arXiv preprint, arXiv number1512.02687 (2016).
- [30] L. Garncarek, N. Lazarovich. The Neretin groups. arXiv 1502.00991v2 (2015).
- [31] P. J. Higgins. The fundamental groupoid of a graph of groups. *J. London Math. Soc.* **13** (1976), 145-149.
- [32] Nicolas Jacon. Representations of the Symmetric Group and its Hecke algebra. *Conference Summer School in Algorithmic Mathematics*, Munich (2012).
- [33] A. Kumjian, D. Pask. Higher rank graph C^* -algebras, *New York J. Math.* **6** 1-20, 2000.
- [34] A. Kumjian, D. Pask, I. Raeburn, J. Renault. Graphs, groupoids and Cuntz-Krieger algebras. *J. Funct. Anal.* **144**(2), 505-541 (1997).
- [35] A. Krieg. Hecke algebras. *Memoirs of the American Mathematical Society*, Providence RI. Volume 87, number 435 (1980).
- [36] W. Leavitt. The module type of a ring. *Trans. Amer. Math. Soc.* **103**, 113-130 (1962).

- [37] W. Lederle. Neretin groups. *Groups Geom. Dyn.* **13** no. 2, 467-510 (2019).
- [38] G. W. Mackey. On induced representations of groups. *Amer. J. Math.* **73** 576-592 (1951) .
- [39] H. Matui. Topological full groups of one-sided shifts of finite type. *J. reine angew. Math.* **705** (35-84) (2015).
- [40] V. Nekrashevych. Cuntz-Pimsner Algebras of Group Actions. *J. Operator Theory* **52** 223-249 (2004).
- [41] V. Nekrashevych. Simple groups of dynamical origin. *Ergod. Th. and Dynam. Sys.* **39**, 707-732 (2019).
- [42] Yu. A. Neretin. On Combinatorial Analogs of the Group of Diffeomorphisms of the Circle. *Russian Academy of Sciences. Izvestiya Mathematics* **41** no. 2, 337-349 (1993).
- [43] R. Okayasu. C^* -algebras associated with the fundamental groups of graphs of groups. *Math. Scand.* **97** 49-72 (2005).
- [44] E. Pardo. The isomorphism problem for Higman-Thompson groups. *J. Algebra* textbf344 172-183 (2011).
- [45] A. Paterson. Graph inverse semigroups, groupoids and their C^* -algebras. *J. Operator Theory* **48** 645-662 (2002).
- [46] A. Paterson. Groupoids, Inverse Semigroups, and their Operator Algebras. Birkhäuser Boston Inc., Boston, MA, xvi+274 (1999).
- [47] I. Putnam. Lecture Notes on C^* -algebras. Online lecture notes, from lecturer's personal page at University of Victoria, B.C. www.math.uvic.ca/faculty/putnam/ (January 2019).
- [48] I. Raeburn. Graph Algebras. CBMS Regional Conference Series in Mathematics, **103**. Published by the American Mathematical Society, Providence, RI (2005).
- [49] J. Renault. A groupoid approach to C^* -algebras. *Lecture Notes in Mathematics*, **793**: Springer, Berlin (1980). ii+160 pp.
- [50] C. E. Röver. Constructing finitely presented simple groups that contain Grigorchuk groups. *J. Algebra* **220** (1):284-313 (1999).

- [51] J-P. Serre. *Trees*. Springer-Verlag, Berlin-Heidelberg-New York (1980).
- [52] B. Steinberg. A groupoid approach to discrete inverse semigroup algebras. *Adv. Math.* **223** 689-727 (2010).
- [53] J. Tits. *Sur le groupe des automorphismes d'un arbre*. Essays on topology and related topics, Springer-Verlag, Berlin-Heidelberg (1970).
- [54] M. Tomforde. Uniqueness theorems and ideal structure for Leavitt path algebras. *J. Algebra* **318**(1), 270-299 (2007).