Tree Actions associated with the Scale of an Automorphism of a Totally Disconnected Locally Compact Group

Daniel Horadam

BCompSc, BMath(Hons)

University of Newcastle, Australia

Thesis submitted for the degree of

Doctor of Philosophy (Mathematics)

April 2015

Statement of Originality

The thesis contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. I give consent to the final version of my thesis being made available worldwide when deposited in the University's Digital Repository, subject to the provisions of the Copyright Act 1968.

Daniel Horadam

Acknowledgements

Uncountably infinite thanks to my long-suffering supervisor George Willis, and to my colleagues, family and friends who have supported me and kept me sane throughout the inordinately long time it took to write this thesis.

Contents

Abstract

One of the key features of the structure theory of totally disconnected locally compact groups is the existence of certain compact open subgroups, called *tidy* subgroups, which are well-behaved under the action of group automorphisms. If V is a compact open subgroup that is tidy for the automorphism α then there is an associated closed subgroup V_{--} which is invariant under α. These V_{--} groups are analogous to eigenspaces for linear operators in the theory of Lie groups. There is a representation of the semi-direct product $V_{--} \rtimes \langle\alpha\rangle$ as a closed subgroup of the stabiliser of an end of a homogeneous tree, and it is this tree-representation that we aim to understand in this thesis.

First, we use the properties of the tree-representation to reduce the problem to understanding the automorphism groups of *rooted* trees that have a self-similarity property which we call $\mathcal{R}$. These groups are compact and hence profinite, which means we can understand them in terms of their finite quotients which have a corresponding property which we call $\mathcal{R}_n$. Then we use the software package `MAGMA` to perform calculations with these finite groups, generating plenty of examples and providing evidence in support of several conjectures about the behaviour of groups with property $\mathcal{R}$.

Finally, we describe two general constructions, both of which take a finite group with property $\mathcal{R}_n$ and extend it to a profinite group with property $\mathcal{R}$. One construction generates the maximal such group, which turns out to be a type of self-similar group called a *finitely constrained* group. We show that all groups with property $\mathcal{R}$ can be approximated by these finitely constrained groups. The other construction uses finite automata to produce topologically finitely generated groups with property $\mathcal{R}$.

CHAPTER 1

Introduction

1.1. Background and motivation

Every locally compact group G is an extension of the connected component of the identity G_0 by the totally disconnected group G/G_0. The study of locally compact groups can therefore be separated into the two extreme cases of connected groups and totally disconnected groups.

Connected groups are well understood through the solution to Hilbert's fifth problem. The results of Gleason [**Gle51**] and Yamabe [**Yam53**] show that they are approximated by Lie groups, in the sense that every connected, locally compact group G contains arbitrarily small normal subgroups N such that G/N is a Lie group. This result allows the powerful techniques of Lie theory to be applied to general connected groups.

On the other hand, totally disconnected groups are not as well understood. The structure theory of totally disconnected locally compact groups began with the work of van Dantzig [**vD36**] who proved that every neighbourhood of the identity in a totally disconnected locally compact group G contains a compact, open subgroup. Furthermore, if G is compact, then every neighbourhood of the identity contains a compact open *normal* subgroup N, so that G/N is finite. Hence totally disconnected compact groups are *profinite* groups (inverse limits of finite groups — see [**RZ10**]).

Further development of the structure theory by Willis in [**Wil94**] examined the behaviour of group automorphisms on these compact open subgroups. Given a totally disconnected locally compact group G and a continuous automorphism α, there exist compact open subgroups which are *tidy* for α. Such a subgroup V splits into a product $V = V_+V_-$ where α expands V_+ and shrinks V_-. The expansion factor $s(\alpha) := |\alpha(V_+) : V_+|$ does not depend on V and is called the *scale* of α. The scale of an automorphism of a totally disconnected locally compact group is analogous to the eigenvalues of a linear operator in Lie theory (more precisely, to the product of the eigenvalues with absolute value greater than 1). The analogue of an eigenspace is the closed subgroup $V_{++} = \bigcup_{n\geq 0} \alpha^n(V_+)$ which is invariant under α. The scale of α^{-1} is defined similarly and is analogous to the eigenvalues of absolute value less than 1, with the closed subgroup $V_{--} = \bigcup_{n\geq 0} \alpha^{-n}(V_-)$ playing the role of the eigenspace. These analogies are explained in more detail in [**Wil04a**] and [**Wil04b**], whereas the scale function has been calculated explicitly using Lie techniques in the case of Lie groups over local fields [**Glö98**] and in groups constructed from p-adic Lie groups [**Glö06**].

1.2. Basic definitions

Let us review some of the key definitions and background results that will be used in this thesis.

DEFINITION 1.1. A *topological group* is a group that is also a topological space, such that the group operations of multiplication and inversion are continuous functions.

A *totally disconnected locally compact group* is a topological group that is totally disconnected (connected components are all singletons) and locally compact (every point has a neighbourhood with compact closure) as a topological space.

DEFINITION 1.2 (Tidy subgroups and the scale function). Let G be a totally disconnected locally compact group and let α be a continuous automorphism of G. For each compact open subgroup V of G, define:

$$V_+ = \bigcap_{n\geq 0} \alpha^n(V) \qquad V_{++} = \bigcup_{n\geq 0} \alpha^n(V_+)$$

$$V_- = \bigcap_{n\geq 0} \alpha^{-n}(V) \qquad V_{--} = \bigcup_{n\geq 0} \alpha^{-n}(V_-)$$

$$V_0 = V_+ \cap V_-.$$

A compact open subgroup V is *tidy* for α if the following two conditions are satisfied:

(TA) $V = V_+V_-$

(TB) V_{++} (and V_{--}) is closed.

Tidy subgroups for α can be constructed from arbitrary compact open subgroups by the procedure in [**Wil01**]. Let V be tidy for α. The *scale* of α is the positive integer

$$s(\alpha) := |\alpha(V_+) : V_+| = |\alpha(V) : V \cap \alpha(V)|.$$

which is finite because $\alpha(V)$ is compact and $V \cap \alpha(V)$ is open. It is shown in [**Wil04a**] that the scale does not depend on the choice of tidy subgroup

V, and that it is equivalent to:

$$s(\alpha) = \min\{|\alpha(U) : U \cap \alpha(U)| : U \text{ is a compact open subgroup of } G\}.$$

DEFINITION 1.3 (Trees). A *tree* is a connected graph with no cycles. A tree is *homogeneous* if all vertices have the same degree. This common degree is called the *valency* of the tree. A *rooted tree* is a tree with one distinguished vertex, called the root. If v is a vertex in a rooted tree X then the *children* of v are the vertices adjacent to v that are further away from the root (equivalently, the path from the root to these vertices passes through v). The *parent* of a non-root vertex v is the unique vertex adjacent to v closer to the root (along the path joining v to the root). A rooted tree is *regular* if all vertices have the same number of children. Every vertex in a regular rooted tree has the same degree $d+1$, except for the root which has degree d. If $d = 2$ then we call it a *binary tree.*

A *path* in a tree X is a sequence of vertices such that there is an edge joining each vertex to its successor, with no 'backtracking' — that is, no edge may be traversed more than once in succession. Paths may be finite, singly infinite (of the form $(v_n)_{n=1}^{\infty}$), or doubly infinite (of the form $(v_n)_{n\in\mathbb{Z}}$). For any two vertices in a tree, there is a unique (finite) path joining them.

The *boundary* of a tree X, denoted ∂X, is the set of equivalence classes of singly infinite paths in X, where two paths ξ and η are equivalent if $\xi \cap \eta$ is infinite (equivalently, if the paths eventually coincide). These equivalence classes are called the *ends* of X. We say that every doubly infinite path $(v_n)_{n\in\mathbb{Z}}$ in X has precisely two ends, namely the ends corresponding to $(v_n)_{n\geq 0}$ and $(v_n)_{n\leq 0}$. If X is a rooted tree, then the ends of X are in bijection with the singly infinite paths starting from the root.

DEFINITION 1.4 (Automorphisms). An *automorphism* of a tree X is a bijection g on the set of vertices of X, such that $g(v)$ is adjacent to $g(w)$ if and only if v is adjacent to w. The set of all automorphisms of X is a group under composition of functions, called the *automorphism group* of X and denoted $\mathsf{Aut}(X)$. We will denote the identity in this group by e.

Tits proved in [**Tit70**] that every automorphism $g \in \mathsf{Aut}(X)$ satisfies exactly one of the following conditions: g fixes a vertex of X, g inverts an edge of X, or there is a doubly infinite path in X along which g acts as a translation. The first two types of automorphisms are called *elliptic* and the third kind is called *hyperbolic.* Note that only the first kind of automorphism can exist if X is a rooted tree, because every automorphism of X must fix the root.

The *stabiliser* of a vertex v in X is the set

$$\mathsf{st}(v) = \{g \in \mathsf{Aut}(X) : g(v) = v\}$$

which is always a subgroup of $\mathsf{Aut}(X)$. If G is a subgroup of $\mathsf{Aut}(X)$ then the stabiliser of v in G is $\mathsf{st}(v) \cap G$ which we denote $\mathsf{st}_G(v)$. It is also a subgroup of G. The stabiliser of an end $\omega \in \partial X$ is the set

$$\mathsf{st}(\omega) = \{g \in \mathsf{Aut}(X) : g(\xi) \in \omega \text{ for all } \xi \in \omega\}$$

(remembering that ω is an equivalence class of infinite paths) which is also a subgroup of $\mathsf{Aut}(X)$.

Definition 1.5 (Topology and convergence). Let X be a tree. The group $\mathsf{Aut}(X)$ can be made into a topological group by endowing it with the *compact-open topology*, defined in this case as follows. The open neighbourhoods of the identity are defined to be all unions of sets of the form $\mathsf{st}(F) = \bigcap_{v \in F} \mathsf{st}(v)$ where F is a finite set of vertices in X. In particular, $\mathsf{st}(v)$ is an open set for all $v \in X$, and since it is also a subgroup of $\mathsf{Aut}(X)$ it is closed as well. The other open sets are obtained by translating these neighbourhoods of e; thus, the general open neighbourhoods of $g \in \mathsf{Aut}(X)$ are unions of sets of the form

$$\mathcal{U}(g, F) = \{h \in \mathsf{Aut}(X) : h(v) = g(v) \text{ for all } v \in F\}$$

where F is a finite set of vertices in X.

It follows that every point in $\mathsf{Aut}(X)$ has arbitrarily small neighbourhoods which are both open and closed, which makes $\mathsf{Aut}(X)$ totally disconnected. It is locally compact if X is locally finite (each vertex has finite degree), because then each $\mathsf{st}(v)$ will be a compact set. These are the arbitrarily small compact open subgroups guaranteed by van Dantzig's theorem. Furthermore, if X is a rooted tree, then the group $\mathsf{Aut}(X)$ itself is compact, because the root is fixed by all automorphisms.

Another way of thinking of the topology on $\mathsf{Aut}(X)$ is in terms of convergence of sequences. A sequence $(g_n)_{n=1}^{\infty}$ converges to g in $\mathsf{Aut}(X)$ if and only if for every finite set of vertices F there exists an integer N_F such that $g_n(v) = g(v)$ for all $n \geq N_F$ and all $v \in F$. In other words, $g_n \to g$ if g_n eventually agrees with g on any finite set of vertices.

1.3. The tree-representation

Let G be a totally disconnected locally compact group, α a continuous automorphism of G and V a compact open subgroup tidy for α. In [**BW04**, Theorem 4.1] an action of the group $V_{--} \rtimes \langle\alpha\rangle$ on a homogeneous tree X with valency $s(\alpha^{-1})+1$ is described. The group $V_{--} \rtimes \langle\alpha\rangle$ has the structure of an *HNN extension* as defined in [**HNN49**], and this group action is a special case of the Bass-Serre construction of a graph of groups acting on a tree, as described in [**Ser80**, §5] and [**DD89**, Example 3.5(v)]. Since we are dealing only with this specific case, it is worthwhile to describe the tree X and the representation in detail:

- The vertices of X are the left cosets of V_- in $V_{--} \rtimes \langle\alpha\rangle$. The vertex $(v, \alpha^n)V_-$, where $v \in V_{--}$ and $n \in \mathbb{Z}$, will be denoted (v, n).
- There is an edge from (v, m) to (w, n) if and only if $n = m + 1$ and $w \in v\alpha^m(V_-)$. There are $s(\alpha^{-1})$ out-edges and one in-edge for every vertex. Although this defines a directed graph, we will ignore the edge directions and treat X as an undirected graph.
- The group $V_{--} \rtimes \langle\alpha\rangle$ acts on the vertices of X by left multiplication.

Let $\xi_n = (e, n)$ for each $n \in \mathbb{Z}$, where e is the identity in V_{--}. Note that $\xi = (\xi_n)_{n\in\mathbb{Z}}$ is a doubly infinite path in X. Let ∞ be the end of ξ corresponding to $(\xi_n)_{n=0}^{\infty}$ and let $-\infty$ be the other end of ξ. The tree-representation $\pi : V_{--} \rtimes \langle\alpha\rangle \to \mathsf{Aut}(X)$ has the following properties:

(a) π is continuous with respect to the compact-open topology on $\mathsf{Aut}(X)$.

(b) The image of π is a closed subgroup of $\mathsf{Aut}(X)$ that fixes the end $-\infty$ and is transitive on the other ends of X.

(c) The kernel of π is the largest compact, normal, α-stable subgroup of V_{--}.

(d) $\pi(V_{--})$ is the set of elliptic elements in the image of π.

(e) $\pi(\alpha)$ is a translation by 1 vertex along the path ξ, directed away from $-\infty$.

(f) The stabilizer of the vertex ξ_n is $\pi(\alpha^n(V_-))$, the stabiliser of the end ∞ is $\pi(V_0 \rtimes \langle\alpha\rangle)$, and the stabiliser of each end other than $-\infty$ is a conjugate of $\pi(V_0 \rtimes \langle\alpha\rangle)$.

Note that in the case where $G = \mathsf{Aut}(X)$ and α is (conjugation by) a translation along a path in X, the image of π is precisely the stabiliser in G of the repelling end for α.

In a nutshell, the tree-representation theorem shows that every totally disconnected locally compact group (provided it has automorphisms with scale $\neq 1$) has a sub-quotient, namely $\pi(V_{--} \rtimes \langle\alpha\rangle)$, that is isomorphic to a closed subgroup of the stabiliser of an end of a homogeneous tree.

The goal of this work is to understand these subgroups. As we will see in Chapter 2, it suffices to study the compact subgroup $\pi(V_-)$, which is the stabiliser of the vertex ξ_0.

1.4. Outline of thesis

In Chapter 2 we reduce the problem of understanding the groups arising from the tree-representation to the problem of understanding a family of compact groups (quotients of $\pi(V_-)$). These groups act on a rooted subtree of X, in such a way that this action completely determines the whole tree-representation. These groups possess a defining property which we call $\mathcal{R}$, which is related to self-similarity. Such groups are profinite, which reduces the problem further to finite groups and inverse limits thereof. These finite groups are required to have a property $\mathcal{R}_n$ which is a finite-depth version of the property $\mathcal{R}$.

Chapter 3 establishes some basic properties of the finite groups with property $\mathcal{R}_n$. The most significant of these results involves the existence of so-called *rigid automorphisms*, which allows us to strengthen property $\mathcal{R}_n$ — and in turn, property $\mathcal{R}$ — by choosing appropriate conjugacy class representatives for each group. It turns out that we can assume, without loss of generality, that groups with property $\mathcal{R}$ are *self-similar* (in fact, *self-replicating*), allowing us to draw upon the existing theory of self-similar groups. This has been a useful source of both examples and ideas, helping to motivate the construction of the automaton groups in Chapter 7, for example.

Chapter 4 makes things a little more concrete by describing in detail the tree representations of the p-adic numbers and the Laurent series over a finite field. These examples provide the starting point for understanding the proverbial zoo of groups satisfying property $\mathcal{R}$.

Chapter 5 describes the results of preliminary calculations (using the software package `MAGMA`) to find all groups with property $\mathcal{R}_n$ for small n in the case of the binary tree ($s(\alpha^{-1}) = 2$). The detailed results of these calculations, as well as the algorithm used to produce them, are deferred to the Appendix. Some clear patterns emerge from these results, motivating several conjectures and two general constructions (detailed in Chapters 6 and 7). We also see how to turn the family of groups satisfying $\mathcal{R}$ into a totally disconnected compact topological space, $\mathfrak{S}$, by associating it with the boundary of an infinite rooted tree. We find a countable dense subset of $\mathfrak{S}$, which might help to make a classification more feasible, and we also establish a connection between finitely generated groups (an algebraic condition) and isolated points in $\mathfrak{S}$ (a topological condition).

In Chapter 6 we describe a canonical way to extend each finite group G with property $\mathcal{R}_n$ to an infinite group $\mathcal{M}^\infty(G)$ with property $\mathcal{R}$. We show that this construction produces a maximal group in the sense that every other extension of G with property $\mathcal{R}$ is conjugate to a subgroup of $\mathcal{M}^\infty(G)$. It turns out that this construction always produces a *finitely constrained* group, establishing yet another connection with the theory of self-similar groups. We also compute the Hausdorff dimension of these groups and show that it is always strictly positive.

Chapter 7 provides a connection with the theory of automaton groups. As an alternative to the maximal groups of Chapter 6 which are usually not topologically finitely generated, we describe a procedure to take a group G with property $\mathcal{R}_n$ and construct at least one (in fact, usually a very large number) of topologically finitely generated extensions $\mathcal{G}(\mathcal{A}_G)$ with property $\mathcal{R}$. It turns out that it is possible for this construction to reproduce $\mathcal{M}^\infty(G)$, and we conjecture that these groups are precisely the isolated points in $\mathfrak{S}$ from Chapter 5.

CHAPTER 2

Reduction to rooted trees

2.1. Introduction

The groups arising from the tree representation π (see Section 1.3) act on a homogeneous tree X. In this chapter, we will see how to reduce the study of these groups to the study of a class of groups acting on a rooted regular tree T, which will be a subtree of X.

The image of π is a semidirect product $\Gamma \rtimes \mathbb{Z}$, where $\Gamma = \pi(V_{--})$ is closed in $\mathsf{Aut}(X)$ and $\mathbb{Z} \cong \langle \pi(\alpha) \rangle$. For convenience we will simply write α instead of $\pi(\alpha)$ from now on. Then $n \in \mathbb{Z}$ acts on Γ in the semidirect product by

$$n : g \mapsto \alpha^n g \alpha^{-n}.$$

Recall that α acts on X as a translation of amplitude 1 along an axis $\xi = (\xi_n)_{n\in\mathbb{Z}}$, so that $\alpha(\xi_n) = \xi_{n+1}$ for all $n \in \mathbb{Z}$. As in Section 1.3, let ∞ be the end of ξ corresponding to $(\xi_n)_{n=0}^{\infty}$ and let $-\infty$ be the other end of ξ. Then Γ fixes $-\infty$ and acts transitively on $\partial X \setminus \{-\infty\}$.

We begin by focusing our attention on the group Γ.

2.2. From homogeneous trees to rooted trees

The purpose of this section is to explain the steps involved in passing from the homogeneous tree X to the rooted tree T, and then to characterise the subgroups of $\mathsf{Aut}(T)$ that can be obtained in this way.

Let $\Gamma_n = \mathsf{st}_\Gamma(\xi_n)$ for each $n \in \mathbb{Z}$. Since Γ fixes $-\infty$, it follows that Γ_n fixes the entire path $(\xi_k)_{k\leq n}$. From the tree-representation and the semidirect product action we have:

$$\Gamma_n = \pi(\alpha^n(V_-)) = \alpha^n \pi(V_-)\alpha^{-n} = \alpha^n \Gamma_0 \alpha^{-n} \tag{2.1}$$

and therefore:

$$\Gamma = \pi(V_{--}) = \bigcup_{n=0}^{\infty} \pi(\alpha^{-n}(V_-)) = \bigcup_{n=0}^{\infty} \Gamma_{-n} = \bigcup_{n=0}^{\infty} \alpha^{-n}\Gamma_0\alpha^n \tag{2.2}$$

and this is an increasing union since $\Gamma_n \subseteq \Gamma_{n-1}$ for all n. It follows that in order to understand the group Γ, it suffices to understand Γ_0 and how it interacts with α.

Let us first examine the topology of Γ_0. Like Γ, it is closed in $\mathsf{Aut}(X)$, but even better, we have:

PROPOSITION 2.1. *Γ_n is compact for all $n \in \mathbb{Z}$.*

PROOF. Vertex stabilisers are compact sets in $\mathsf{Aut}(X)$ and Γ is closed, so $\Gamma_n = \mathsf{st}(\xi_n) \cap \Gamma$ is compact. ☠

All compact, totally disconnected groups are profinite by [**RZ10**, Theorem 1.1.12] so in theory we can express Γ_0 as an inverse limit of finite groups. However, for our purposes it is more convenient to express Γ_0 as an inverse limit of *infinite* groups, each of which is profinite and acts on a rooted subtree of X.

DEFINITION 2.2. For each $n \in \mathbb{Z}$ define the tree X_n to be the subtree of X rooted at the vertex ξ_n and containing all vertices descending from ξ_n in the direction away from $-\infty$. That is, a vertex $x \in X$ belongs to X_n if and only if the unique path from x to $-\infty$ passes through ξ_n.

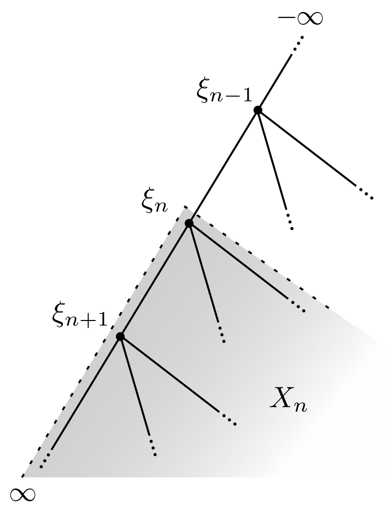

FIGURE 2.1. The rooted subtree X_n of X.

Note that the X_n are nested, with $X_n \subset X_{n-1}$ for all $n \in \mathbb{Z}$, and $\bigcup_{n\in\mathbb{Z}} X_n = X$. Because of the nesting we can do a little better and write

$$\bigcup_{n\leq 0} X_n = X. \tag{2.3}$$

Compare this to the similar formula $\bigcup_{n\leq 0} \Gamma_n = \Gamma$ from (2.2).

Since Γ_n fixes ξ_n and also fixes $-\infty$, it leaves the subtree X_n invariant. It makes sense therefore to restrict the action of Γ_n to X_n. This restriction is a continuous homomorphism from Γ_n to $\mathsf{Aut}(X_n)$, which is a compact totally disconnected group. We will see more of this group later. For now, we can say the following:

PROPOSITION 2.3. *$\Gamma_n|_{X_n}$ is a closed subgroup of $\mathsf{Aut}(X_n)$.*

PROOF. The restriction map is continuous and Γ_n is compact by Proposition 2.1, so the image in $\mathsf{Aut}(X_n)$ is compact and therefore closed. ☠

PROPOSITION 2.4. *$\Gamma_n|_{X_n}$ is transitive on ∂X_n.*

PROOF. Γ_n leaves X_n invariant so the assertion is equivalent to Γ_n being transitive on ∂X_n. Since $\infty \in \partial X_n$ it suffices to show that for every $\omega \in \partial X_n$ there exists $g \in \Gamma_n$ such that $g(\infty) = \omega$. (Indeed, suppose we have shown this. Then for any $\omega' \in \partial X_n$ there would exist $h \in \Gamma_n$ such that $h(\infty) = \omega'$ and we would have $hg^{-1}(\omega) = \omega'$).

Let $\omega \in \partial X_n$. Since Γ is transitive on $\partial X \setminus \{-\infty\}$ there exists $g \in \Gamma$ such that $g(\infty) = \omega$. We must show that $g \in \Gamma_n$; in other words, that $g(\xi_n) = \xi_n$.

Consider the unique doubly infinite path η joining $-\infty$ to ω and recall that ξ is the unique path joining $-\infty$ to ∞. Since both η and ξ share the end $-\infty$, they must eventually coincide. Therefore there exists $N \in \mathbb{Z}$ such that $\xi_k \in \eta$ for all $k \leq N$. Furthermore, since $\omega \in \partial X_n$ we must have $\xi_n \in \eta$ and hence we may take $N \geq n$.

It follows from the fact that Γ fixes $-\infty$ that $g(\xi) = \eta$. Now $g \in \Gamma$ so by (2.2) there exists $M \in \mathbb{Z}$ such that $g(\xi_k) = \xi_k$ for all $k \leq M$. If $M \geq n$ then g fixes ξ_n and we are done. Suppose that $M < n$. Since g fixes ξ_M we know that $g(\xi_n)$ is a vertex in η at a distance $n - M$ from ξ_M. There are only two such vertices, namely ξ_n and ξ_{2M-n} (note that they both belong to η since $2M - n < n \leq N$). The latter is fixed by g since $2M - n < M$, therefore $g(\xi_n) = \xi_n$ as claimed. ☠

We still need to understand how Γ_n interacts with α. We know from (2.1) that any two of the Γ_n are conjugate by a power of α. We can say more:

LEMMA 2.5. *If $k > 0$ then $\mathsf{st}_{\Gamma_n}(\xi_{n+k}) = \Gamma_{n+k} = \alpha^k \Gamma_n \alpha^{-k}$.*

PROOF. The second equality follows from (2.1). For the first:

$$\begin{aligned}\mathsf{st}_{\Gamma_n}(\xi_{n+k}) &= \mathsf{st}(\xi_{n+k}) \cap \Gamma_n \\ &= \mathsf{st}(\xi_{n+k}) \cap \mathsf{st}(\xi_n) \cap \Gamma \\ &= \Gamma_{n+k} \cap \Gamma_n \\ &= \Gamma_{n+k} \qquad (\text{since } \Gamma_{n+k} \subseteq \Gamma_n).\end{aligned}$$

☠

In other words, the stabiliser of ξ_{n+k} in Γ_n is conjugate to Γ_n itself. Furthermore, since α is a translation along ξ, conjugation by α has the effect of shifting the action of $g \in \Gamma$ by one vertex away from $-\infty$. To explain this more precisely, first define the map

$$\begin{aligned}\phi : \Gamma &\longrightarrow \Gamma \\ g &\longmapsto \alpha g \alpha^{-1}.\end{aligned}$$

Clearly ϕ is just the tree representation of the automorphism α from $V_{--} \rtimes \langle\alpha\rangle$. It is introduced here for ease of notation. By Lemma 2.5,

$$\phi(\Gamma_n) = \Gamma_{n+1} = \mathsf{st}_{\Gamma_n}(\xi_{n+1}) \text{ for all } n \in \mathbb{Z}. \tag{2.4}$$

We have a similar equation for the rooted subtrees X_n:

$$\alpha(X_n) = X_{n+1} \text{ for all } n \in \mathbb{Z}. \tag{2.5}$$

These two equations are compatible, in the sense that

$$\phi(g)(\alpha(v)) = (\alpha g \alpha^{-1})(\alpha(v)) = \alpha(g(v)) \tag{2.6}$$

for all $g \in \Gamma$ and $v \in X$. We can depict this diagramatically:

$$\begin{array}{ccc}\Gamma_n & \xrightarrow{\ \phi\ } & \Gamma_{n+1} \\ \circlearrowleft & & \circlearrowleft \\ X_n & \xrightarrow{\ \alpha\ } & X_{n+1}\end{array}$$

In particular, if $g \in \Gamma_n$ and $v \in X_n$ then restricting to the appropriate rooted subtrees and using (2.4), (2.5) and (2.6) yields the important formula:

$$\Gamma_{n+1}|_{X_{n+1}} = \mathsf{st}_{\Gamma_n}(\xi_{n+1})|_{X_{n+1}} = \phi(\Gamma_n)|_{\alpha(X_n)} = \alpha(\Gamma_n|_{X_n}). \tag{2.7}$$

In other words, when X_n is identified with X_{n+1} via the translation α, the action of Γ_n on X_n is identical to the action of $\Gamma_{n+1} = \mathsf{st}_{\Gamma_n}(\xi_{n+1})$ on X_{n+1}. Let us denote this more concisely by writing

$$\Gamma_{n+1}|_{X_{n+1}} \overset{\alpha}{=} \Gamma_n|_{X_n}. \tag{2.8}$$

For all $n \geq 0$, Γ_0 fixes ξ_{-n} so it leaves X_{-n} invariant and the restriction $\Gamma_0|_{X_{-n}}$ makes sense. By (2.3) every vertex in X belongs to X_{-n} for some $n \geq 0$, so Γ_0 is determined by the restrictions $\Gamma_0|_{X_{-n}}$ for $n \geq 0$. Using Lemma 2.5 along with (2.7), we conclude that

$$\Gamma_0|_{X_{-n}} = \mathsf{st}_{\Gamma_{-n}}(\xi_0)|_{X_{-n}} = \mathsf{st}_{\Gamma_{-n}|_{X_{-n}}}(\xi_0) = \mathsf{st}_{\alpha^{-n}(\Gamma_0|_{X_0})}(\xi_0) \tag{2.9}$$

for all $n \geq 0$. It now follows that $\Gamma_0|_{X_0}$, along with the translation α, completely determines each $\Gamma_0|_{X_{-n}}$ and hence all of Γ_0. This means that if we know how Γ_0 acts on the rooted subtree X_0 then we can reconstruct its action on the entire homogeneous tree X.

We can be more precise about this. For each $n > 0$ define the map:

$$\begin{aligned} \psi_n : \Gamma_0|_{X_{-n}} &\longrightarrow \Gamma_0|_{X_{-n+1}} \\ g &\longmapsto g|_{X_{-n+1}} \end{aligned}$$

and note that the restriction makes sense since $X_{-n+1} \subset X_{-n}$ and both subtrees are invariant under Γ_0. It is clear that each ψ_n is a surjective homomorphism. We therefore have an inverse system:

$$\Gamma_0|_{X_0} \xleftarrow{\psi_1} \Gamma_0|_{X_{-1}} \xleftarrow{\psi_2} \Gamma_0|_{X_{-2}} \xleftarrow{\psi_3} \Gamma_0|_{X_{-3}} \leftarrow \cdots$$

and so the inverse limit $\varprojlim \Gamma_0|_{X_{-n}}$ exists.

Proposition 2.6. *Γ_0 is isomorphic to $\varprojlim \Gamma_0|_{X_{-n}}$ as a topological group.*

Proof. The inverse limit is the set of sequences of the form $(g_n)_{n=0}^{\infty}$ where $g_n \in \Gamma_0|_{X_{-n}}$ and $\psi_n(g_n) = g_{n-1}$ for each $n \geq 1$, with coordinate-wise multiplication and carrying the product topology inherited from $\prod_{n=0}^{\infty} \Gamma_0|_{X_{-n}}$. Define the map

$$\begin{aligned} \theta : \Gamma_0 &\longrightarrow \varprojlim \Gamma_0|_{X_{-n}} \\ g &\longmapsto \left(g|_{X_{-n}}\right)_{n=0}^{\infty}. \end{aligned}$$

We will show that θ is an isomorphism of topological groups. It is clear that θ is a group homomorphism since $(gh)|_{X_{-n}} = (g|_{X_{-n}})(h|_{X_{-n}})$ and multiplication in the inverse limit is defined coordinate-wise. Each $g \in \ker\theta$ acts trivially on each X_{-n} and hence on all of X by (2.3). Thus $\ker\theta$ is trivial and θ is injective.

To show that θ is surjective, suppose that $(g_n)_{n=0}^{\infty} \in \varprojlim \Gamma_0|_{X_{-n}}$. Define $g \in \mathsf{Aut}(X)$ as follows. Let $v \in X$. Then $v \in X_{-k}$ for some $k \geq 0$ by (2.3), so define $g(v) = g_k(v)$. The maps ψ_n in the definition of the inverse limit ensure that g is a well-defined automorphism of T. Since each g_n fixes ξ_0 we have $g \in \mathsf{st}(\xi_0)$. To show that $g \in \Gamma_0$ we express it as a limit of a sequence of elements of Γ_0 and use the fact (from Proposition 2.1) that Γ_0 is closed in $\mathsf{Aut}(X)$. Indeed, define a sequence $(\gamma_n)_{n=0}^{\infty}$ as follows. For each $n \geq 0$ let γ_n be any element of Γ_0 such that $\gamma_n|_{X_{-n}} = g_n$. Such a γ_n must exist since $g_n \in \Gamma_0|_{X_{-n}}$. We claim that $\lim_{n\to\infty} \gamma_n = g$. Let $\mathcal{F}$ be any finite set of vertices in X. By (2.3) there exists $k \geq 0$ such that $\mathcal{F} \subset X_{-k}$ (and hence $\mathcal{F} \subset X_{-n}$ for all $n \geq k$). Then for all $n \geq k$ and all $v \in \mathcal{F}$, the definitions of γ_n and g imply that $\gamma_n(v) = g_n(v) = g(v)$. That is, the sequence $(\gamma_n)_{n=0}^{\infty}$ eventually agrees with g on $\mathcal{F}$. Since $\mathcal{F}$ was arbitrary, we conclude from the definition of the topology on $\mathsf{Aut}(X)$ that $\lim_{n\to\infty} \gamma_n = g$ as claimed. It follows immediately from the definitions that $\theta(g) = (g_n)_{n=0}^{\infty}$, completing the proof that θ is surjective.

It remains to show that θ is continuous. Since $\varprojlim \Gamma_0|_{X_{-n}}$ carries the product topology, by [**Mun00**, Theorem 19.6] it suffices to show that the composition of θ with each of the projection maps $\delta_m : \varprojlim \Gamma_0|_{X_{-n}} \to \Gamma_0|_{X_{-m}}$ is continuous. This is immediate since $(\delta_m \circ \theta)(g) = g|_{X_{-m}}$ for all $g \in \Gamma_0$, and restriction to X_{-m} is continuous. Finally, it follows from [**Mun00**, Theorem 26.6] that θ is a homeomorphism because Γ_0 is compact and $\varprojlim \Gamma_0|_{X_{-n}}$ is Hausdorff. ☠

Combining Proposition 2.6 with (2.9) means that if we know $\Gamma_0|_{X_0}$ then we can reconstruct Γ_0 and then use (2.2) to reconstruct the entire group Γ. Propositions 2.3 and 2.4, together with (2.7), give us the algebraic and topological properties that $\Gamma_0|_{X_0}$ inherits from Γ. Let us therefore turn our attention to studying the subgroups of $\mathsf{Aut}(X_0)$ with these properties.

2.3. Property $\mathcal{R}$ and further reduction to finite trees

Section 2.2 reduced the study of the group Γ that arises from the tree representation, to the study of the group $\Gamma_0|_{X_0}$ which acts on a rooted tree rather than a homogeneous tree. From now on, for simplicity of notation, let $T = X_0$, so that $\Gamma_0|_{X_0}$ is a subgroup of $\mathsf{Aut}(T)$.

2.3.1. Property $\mathcal{R}$. For each $n > 0$ define $T^{(n)}$ to be the subtree of T rooted at ξ_n and directed away from the root ξ_0. That is, $T^{(n)}$ is just the tree X_n from the Section 2.2. Note that $T^{(n)}$ is invariant under Γ_n. For convenience we will summarise the results from Section 2.2 using the updated notation:

Proposition 2.7. *Let $G = \Gamma_0|_{X_0}$. Then:*

(a) *G is closed in* $\mathsf{Aut}(T)$*;*
(b) *G is transitive on ∂T;*
(c) $\mathsf{st}_G(\xi_1)|_{T^{(1)}} \overset{\alpha}{=} G$.

PROOF.

(a) See Proposition 2.3.
(b) See Proposition 2.4.
(c) Follows from (2.7) by putting $n = 0$. ☠

DEFINITION 2.8. If a subgroup G of $\mathsf{Aut}(T)$ satisfies the three hypotheses in Proposition 2.7, we will say that G has property $\mathcal{R}$.

REMARK. The symbol $\mathcal{R}$ has been chosen for convenience and does not stand for any particular word; the reader might wish to associate it with the words rooted, restriction, representation, or replicating. Indeed, property (c) above is closely related to the definition of a *self-replicating* group — the precise connection will be established in Chapter 3.

From now on we will focus entirely on studying the subgroups of $\mathsf{Aut}(T)$ with property $\mathcal{R}$. To make this task easier, let us introduce more convenient notation for vertices, subtrees and automorphisms of T.

2.3.2. Vertices and words. The vertices of T can be labelled with (finite) words over the alphabet $\mathsf{X} := \{0, 1, \ldots, p-1\}$, where $p = s(\alpha^{-1})$. We will use the letter p although it need not be a prime number. Each word v over X is a string of symbols $v_1 \cdots v_n$ where each v_i belongs to X. This includes the empty word $\varnothing$ which represents the root of T. Words are joined by concatenation; if $v = v_1 \cdots v_n$ and $w = w_1 \cdots w_m$ then vw is the word $v_1 \cdots v_n w_1 \cdots w_m$. Using this 'multiplicative' notation, we may represent repeated symbols using indices; for example, if $a \in \mathsf{X}$, the word aaa may be written as a^3.

For any vertex in T labelled with the word v, the p children of the vertex labelled v are labelled with the words of the form va where $a \in \mathsf{X}$. If $v \neq \varnothing$ then the parent of v is found by deleting the last symbol of v (the root has no parent). We always choose a labelling such that the vertex ξ_n from the previous section always has the label 0^n for each $n > 0$. The freedom in labelling the remaining vertices can be interpreted in terms of conjugation in the group $\mathsf{Aut}(T)$ — this issue will be addressed in detail in Section 3.2.

Let $|v|$ denote the length of the word v. Define the nth level of T, denoted $\mathcal{L}_n$, to be the set of words of length n. Equivalently, $\mathcal{L}_n$ is the set of vertices of T at a distance n from the root. Each level $\mathcal{L}_n$ is invariant under all automorphisms of T (because all automorphisms of T must fix the root), so $|g(v)| = |v|$ for all $g \in \mathsf{Aut}(T)$ and all $v \in T$.

2.3.3. Vertex restrictions. For all $v \in T$, define $T^{(v)}$ to be the subtree of T consisting of all words beginning with v (along with the edges connecting them), called the subtree of T rooted at v. It is clear that $T^{(v)}$ is isomorphic to T; the isomorphism can be achieved by deleting the prefix v from every word in $T^{(v)}$.

To define vertex restrictions we need the following result:

LEMMA 2.9. *Let $g \in \mathsf{Aut}(T)$. For all $v, w \in T$,*

$$g(vw) = g(v)w'$$

for some $w' \in T$ with $|w| = |w'|$.

PROOF. Since vw belongs to the subtree $T^{(v)}$ and g is an automorphism of T, $g(vw)$ must belong to $T^{(g(v))}$. This establishes the existence of w'. Finally, since automorphisms preserve word length, $|g(vw)| = |vw| = |v| + |w| = |g(v)| + |w|$. It follows that $|w| = |w'|$ as claimed. ☠

The induced map $w \mapsto w'$ is very useful, so we give it a name:

DEFINITION 2.10. Let $g \in \mathsf{Aut}(T)$ and $v \in T$. The map $g|_v : T \to T$ defined by the equation

$$g(vw) = g(v)g|_v(w) \quad \text{for all } w \in T$$

is called the *vertex restriction* of g to v (or the *section* of g at v).

PROPOSITION 2.11. *For all $g \in \mathsf{Aut}(T)$ and all $v \in T$, the map $g|_v$ is an automorphism of T.*

PROOF. Fix $g \in \mathsf{Aut}(T)$ and $v \in T$. First we will show that $g|_v$ is a bijection, by finding its inverse. For all $w \in T$, using the definition of vertex restrictions, we have:

$$\begin{aligned} vw &= g^{-1}g(vw) \\ &= g^{-1}\left(g(v)g|_v(w)\right) \\ &= v \cdot g^{-1}|_{g(v)}\left(g|_v(w)\right) \end{aligned}$$

which implies that $g^{-1}|_{g(v)}\left(g|_v(w)\right) = w$. Therefore $g^{-1}|_{g(v)} \circ g|_v = e$. A similar calculation shows that $g|_v \circ g^{-1}|_{g(v)} = e$ as well, so $g|_v$ must be a bijection.

It remains to show that $g|_v$ preserves the edge relation in T. Suppose that w and w' are adjacent vertices of T. Without loss of generality, w is the parent of w' so $w' = wx$ for some $x \in \mathsf{X}$. We must show that $g|_v(w)$ is the parent of $g|_v(wx)$. We will do this by expressing $g(vwx)$ using vertex restrictions in two different ways. First, by definition of $g|_v$,

$$g(vwx) = g(v)g|_v(wx). \tag{2.10}$$

On the other hand,

$$g(vwx) = g(vw)g|_{vw}(x) = g(v)g|_v(w)g|_{vw}(x).$$

Equating this with (2.10), we conclude that

$$g|_v(wx) = g|_v(w)g|_{vw}(x)$$

which does indeed show that $g|_v(w)$ is the parent of $g|_v(wx)$ as required.

It is worth pointing out a technicality here. If $g(v) = v$ then g leaves $T^{(v)}$ invariant, so it is convenient to think of the map $g|_v$ as the usual restriction of g to $T^{(v)}$. However, this is not quite correct, because $g|_v$ maps T to T and $T^{(v)}$ is not strictly equal to T — rather, it is identified with T by deleting the prefix v. Furthermore, $g|_v$ is well-defined even when $g(v) \neq v$. The map $g|_v$ is better understood as a composition of three maps: first the prefix v is prepended, then g is applied, then the new prefix $g(v)$ is deleted.

Let us establish some basic properties of vertex restrictions.

Proposition 2.12. *Let $g, h \in \mathsf{Aut}(T)$ and $v, w \in T$. Then:*

- **(a)** $(g|_v)|_w = g|_{vw}$
- **(b)** $(gh)|_v = g|_{h(v)}h|_v$
- **(c)** $g^{-1}|_v = (g|_{g^{-1}(v)})^{-1}$, *or equivalently* $(g|_v)^{-1} = g^{-1}|_{g(v)}$.

Proof. **(a)** Let $g \in \mathsf{Aut}(T)$ and let $v, w \in T$. Then by definition of vertex restriction, for all $x \in T$ we have

$$g(vwx) = g(vw)g|_{vw}(x). \tag{2.11}$$

On the other hand, if we split the word vwx into v and wx first, we obtain:

$$\begin{aligned} g(vwx) &= g(v)g|_v(wx) \\ &= g(v)g|_v(w)(g|_v)|_w(x) \\ &= g(vw)(g|_v)|_w(x) \end{aligned}$$

for all $x \in T$. The result follows by equating this with (2.11).

(b) Let $g, h \in \mathsf{Aut}(T)$ and $v \in T$. First we have

$$(gh)(vw) = (gh)(v)(gh)|_v(w) \tag{2.12}$$

for all $w \in T$. On the other hand:

$$\begin{aligned}(gh)(vw) &= g(h(vw)) \\ &= g(h(v)h|_v(w)) \\ &= g(h(v))g|_{h(v)}(h|_v(w)) \\ &= (gh)(v)(g|_{h(v)}h|_v)(w)\end{aligned}$$

for all $w \in T$. Again, the result follows by equating this with (2.12).

(c) Follows easily from (b) by putting $h = g^{-1}$, and the equivalent statement is obtained by swapping g and g^{-1}.

We can restate property $\mathcal{R}$ using this more convenient notation:

DEFINITION 2.13. A subgroup G of $\mathsf{Aut}(T)$ has property $\mathcal{R}$ if:

($\mathcal{R}$1) G is closed in $\mathsf{Aut}(T)$,
($\mathcal{R}$2) G is transitive on ∂T, and
($\mathcal{R}$3) $\mathsf{st}_G(0)|_0 = G$.

As remarked earlier, the condition ($\mathcal{R}$3) is closely related to the definition of a self-replicating group, where the vertex 0 may be replaced with any vertex in T. We will see in Chapter 3 that this much stronger condition actually holds, up to conjugacy, for groups with property $\mathcal{R}$.

It turns out that the transitivity condition ($\mathcal{R}$2) can be weakened, provided that ($\mathcal{R}$1) and ($\mathcal{R}$3) hold:

PROPOSITION 2.14. *Suppose that G is a subgroup of $\mathsf{Aut}(T)$ that satisfies ($\mathcal{R}$1) and ($\mathcal{R}$3). Then the following are equivalent:*

(a) *G is transitive on $\mathcal{L}_1$;*
(b) *G is transitive on $\mathcal{L}_n$ for all n;*
(c) *G is transitive on ∂T.*

PROOF. (a) $\implies$ (b): Suppose that G satisfies ($\mathcal{R}$1) and ($\mathcal{R}$3), and is transitive on $\mathcal{L}_1$. This forms the base step for an induction argument. Suppose that G is transitive on $\mathcal{L}_n$ for some $n \geq 1$. We must show that G is transitive on $\mathcal{L}_{n+1}$.

Let v_1 and v_2 be vertices in $\mathcal{L}_{n+1}$. We must find $g \in G$ such that $g(v_1) = v_2$. Since G is transitive on $\mathcal{L}_1$ there exist h_1 and h_2 in G such that both $h_1(v_1)$ and $h_2(v_2)$ belong to the subtree $T^{(0)}$. In other words, there exist w_1 and w_2 in $\mathcal{L}_n$ such that $h_1(v_1) = 0w_1$ and $h_2(v_2) = 0w_2$. Now G is transitive on $\mathcal{L}_n$ so there exists $k \in G$ such that $k(w_1) = w_2$. Since G satisfies ($\mathcal{R}3$) there exists $g \in \mathsf{st}_G(0)$ such that $g|_0 = k$. Then the product $h_2^{-1} g h_1$ belongs to G and

$$h_2^{-1} g h_1(v_1) = h_2^{-1} g(0w_1) = h_2^{-1}(g(0)g|_0(w_1)) = h_2^{-1}(0w_2) = v_2$$

as required. Since v_1 and v_2 were arbitrary, G is transitive on $\mathcal{L}_{n+1}$ and the inductive step is complete.

(b) $\Longrightarrow$ (c): Suppose that G is transitive on $\mathcal{L}_n$ for all n. Identify ∂T with the set of singly infinite paths in T descending from the root. Let $\omega = (v_n)_{n=0}^{\infty}$ and $\varpi = (w_n)_{n=0}^{\infty}$ be two such paths, so that $v_0 = w_0$ is the root and v_n and w_n are in $\mathcal{L}_n$ for each n. We must find $g \in G$ such that $g(\omega) = \varpi$; that is, $g(v_n) = w_n$ for each n.

To do this, we will use the compactness of G and the finite intersection property. Note that G is compact because it is closed in $\mathsf{Aut}(T)$ by our assumptions and $\mathsf{Aut}(T)$ is compact. For each $n \geq 1$ define the set

$$C_n = \{g \in G : g(v_n) = w_n\}.$$

Each C_n is nonempty since G is transitive on $\mathcal{L}_n$ for each n. Since v_n is the parent of v_{n+1} and w_n is the parent of w_{n+1}, any automorphism which sends v_{n+1} to w_{n+1} must also send v_n to w_n. Hence $C_{n+1} \subseteq C_n$ for each n. It follows that the collection $\{C_n\}_{n=1}^{\infty}$ has the finite intersection property.

Now each C_n is a left coset in G of the stabiliser $\mathsf{st}_G(v_n)$. To see this, fix $n \geq 1$ and let $g \in C_n$. We will show that $C_n = g\mathsf{st}_G(v_n)$. Let $h \in C_n$. Then $(g^{-1}h)(v_n) = g^{-1}(w_n) = v_n$ so $g^{-1}h \in \mathsf{st}_G(v_n)$ and $h \in g\mathsf{st}_G(v_n)$. Hence $C_n \subseteq g\mathsf{st}_G(v_n)$. Conversely, if $k \in \mathsf{st}_G(v_n)$ then $(gk)(v_n) = g(v_n) = w_n$ so $gk \in C_n$. Hence $g\mathsf{st}_G(v_n) \subseteq C_n$ and we conclude that $C_n = g\mathsf{st}_G(v_n)$. Since vertex stabilisers are closed in $\mathsf{Aut}(T)$ and so is G, the intersection $\mathsf{st}_G(v_n) = \mathsf{st}(v_n) \cap G$ is also closed. Therefore C_n is closed as well, being a coset of $\mathsf{st}_G(v_n)$.

Since G is compact and $\{C_n\}_{n=1}^{\infty}$ is a collection of closed subsets of G with the finite intersection property, we conclude that the intersection $\bigcap_{n=1}^{\infty} C_n$ is nonempty. Choose g in this intersection. Then $g \in G$ and $g(v_n) = w_n$ for each n, hence $g(\omega) = \varpi$ as required.

(c) $\Longrightarrow$ (a): Suppose that G is transitive on ∂T and let $v, w \in \mathcal{L}_1$. Then there exist $\omega, \varpi \in \partial T$ whose corresponding infinite paths descending from the root pass through v and w respectively. By transitivity there exists $g \in G$ such that $g(\omega) = \varpi$ and hence $g(v) = w$. ☠

REMARK. Subgroups of $\mathsf{Aut}(T)$ that are transitive on $\mathcal{L}_n$ for all n are called *spherically transitive* (because $\mathcal{L}_n$ is the sphere of radius n around the root, with respect to the graph distance on T). Note that the proof of (b) $\Rightarrow$ (c) above does not actually require property ($\mathcal{R}3$). Nor does the converse, which can be proved by a similar argument to the one used to show (c) $\Rightarrow$ (a). Therefore we have actually proved the following known result:

PROPOSITION 2.15. *A closed subgroup of* $\mathsf{Aut}(T)$ *is spherically transitive if and only if it is transitive on* ∂T.

If property ($\mathcal{R}3$) is satisfied then Proposition 2.14 says that we can extend this equivalence to transitivity on $\mathcal{L}_1$.

2.3.4. Quotients and finite subtrees. We turn our attention now to finite subtrees of T, with the aim of establishing a finite-depth version of property $\mathcal{R}$.

For each $n \geq 1$, define T_n to be the subtree of T consisting of words of length at most n. That is, $T_n = \bigcup_{k=0}^{n} \mathcal{L}_k$. Since each $\mathcal{L}_n$ is invariant under $\mathsf{Aut}(T)$, T_n is invariant also. This means we can restrict elements and subgroups of $\mathsf{Aut}(T)$ to T_n:

DEFINITION 2.16. For all $g \in \mathsf{Aut}(T)$ and $n \geq 1$, define $g_{[n]}$ to be the restriction of g to T_n.

We use this notation for $g \in \mathsf{Aut}(T_n)$ as well: if $m \leq n$ then $g_{[m]}$ is the restriction of g to T_m.

LEMMA 2.17. *The map* $g \mapsto g_{[n]} : \mathsf{Aut}(T) \to \mathsf{Aut}(T_n)$ *is a homomorphism, and* $\mathsf{Aut}(T_n) = \mathsf{Aut}(T)_{[n]}$.

PROOF. Restriction of a group action to an invariant subset is always a homomorphism. The second claim amounts to saying that this homomorphism is surjective. This follows from the obvious fact that every automorphism of T_n can be extended (not uniquely) to an automorphism of T. ☠

For $n \geq 1$ define the n^{th} *level stabiliser* $\mathsf{st}(\mathcal{L}_n)$ to be the pointwise stabiliser of $\mathcal{L}_n$ in $\mathsf{Aut}(T)$.

PROPOSITION 2.18. *For all* $n \geq 1$, $\mathsf{st}(\mathcal{L}_n)$ *is a normal subgroup of* $\mathsf{Aut}(T)$ *and* $\mathsf{Aut}(T_n) \cong \mathsf{Aut}(T)/\mathsf{st}(\mathcal{L}_n)$.

PROOF. Normality follows from the fact that $\mathsf{st}(\mathcal{L}_n)$ is the kernel of the homomorphism in Lemma 2.17. The image of that homomorphism is $\mathsf{Aut}(T_n)$ so the first isomorphism theorem completes the proof. ☠

Vertex restrictions can also be defined for automorphisms of T_n, since Lemma 2.9 applies to $\mathsf{Aut}(T_n)$ as well as $\mathsf{Aut}(T)$, with the appropriate modifications. The proof is essentially the same and will not be reproduced here:

LEMMA 2.19. *Let $g \in \mathsf{Aut}(T_n)$. For all $v, w \in T_n$ such that $|vw| = n$,*

$$g(vw) = g(v)w'$$

for some $w' \in T_n$ with $|w| = |w'|$.

The definition of vertex restrictions in $\mathsf{Aut}(T_n)$ is the same as in $\mathsf{Aut}(T)$ except for the fact that if $|vw| = n$ then $|w| = n - |v|$ so the induced map $w \mapsto w'$ from Lemma 2.19 acts on $T_{n-|v|}$.

DEFINITION 2.20. Let $g \in \mathsf{Aut}(T_n)$ and $v \in T_n$. The map $g|_v : T_{n-|v|} \to T_{n-|v|}$ defined by the equation

$$g(vw) = g(v)g|_v(w) \quad \text{for all } w \in T_{n-|v|}$$

is called the vertex restriction of g to v.

The argument in Proposition 2.11 also works here to show that $g|_v$ is an automorphism of $T_{n-|v|}$. It is easy to see that the formulas in Proposition 2.12 apply (replacing T with T_n as appropriate) to vertex restrictions in $\mathsf{Aut}(T_n)$ as well.

2.3.5. The property $\mathcal{R}_n$. We have already defined the property $\mathcal{R}$ for subgroups of $\mathsf{Aut}(T)$. This condition can be adapted to subgroups of $\mathsf{Aut}(T_n)$ as follows. First, we need some more notation: for each $n \geq 1$ define $T_n^{(0)}$ to be the subtree of T_{n+1} consisting of all words beginning with 0. This tree is isomorphic to T_n (by deleting the prefix 0), and this isomorphism of trees induces a natural isomorphism between the groups $\mathsf{Aut}(T_n^{(0)})$ and $\mathsf{Aut}(T_n)$. Note that $T_n^{(0)}$ is invariant under $\mathsf{st}(0)$, so if $g \in \mathsf{Aut}(T_{n+1})$ and $g(0) = 0$ then the restriction $g|_{T_n^{(0)}}$ is well-defined (and equal to $g|_0$).

Now define the following maps:

$$\begin{aligned} \varphi_n : \mathsf{Aut}(T_n) &\longrightarrow \mathsf{Aut}(T_{n-1}) & \qquad \psi_n : \mathsf{st}_{\mathsf{Aut}(T_n)}(0) &\longrightarrow \mathsf{Aut}(T_{n-1}) \\ g &\longmapsto g_{[n-1]} & g &\longmapsto g|_0. \end{aligned}$$

The map φ_n restricts $\mathsf{Aut}(T_n)$ to the subtree T_{n-1}, while ψ_n restricts the stabiliser of the vertex 0 in $\mathsf{Aut}(T_n)$ to the subtree $T_{n-1}^{(0)}$ (which is then identified with T_{n-1} as described above).

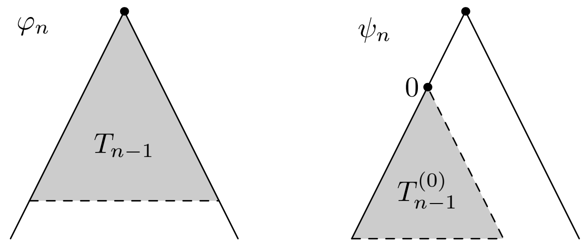

FIGURE 2.2. The restrictions φ_n and ψ_n on the tree T_n

Of course, we already have notation for these maps — we could write $g_{[n-1]}$ for $\varphi_n(g)$ and $g|_0$ for $\psi_n(g)$ — but it will often be more convenient to refer to the maps by name and use the function notation.

LEMMA 2.21. *Both φ_n and ψ_n are surjective homomorphisms for all n. They "commute" in the sense that*

$$\varphi_{n-1} \circ \psi_n = \psi_{n-1} \circ \varphi_n$$

on the domain $\mathsf{st}(0)$.

PROOF. Surjectivity is clear. They are homomorphisms because they are both the restriction of a group action to an invariant subset. To prove "commutativity", observe that both sides of the equation are equal to the restriction of $\mathsf{st}(0)$ to the subtree $T_{n-2}^{(0)}$. ☠

Vertex restrictions in $\mathsf{Aut}(T_n)$ interact with the homomorphisms φ_n and ψ_n as follows. If $g \in \mathsf{Aut}(T_n)$ and $v \in \mathcal{L}_1$ then $g|_v \in \mathsf{Aut}(T_{n-1})$ and

$$\varphi_n(g)|_v = \varphi_{n-1}(g|_v).$$

If, in addition, $g \in \mathsf{st}(0)$ then by definition $\psi_n(g) = g|_0$.

With this notation established, we can now define the property $\mathcal{R}_n$:

DEFINITION 2.22. A subgroup G of $\mathsf{Aut}(T_n)$ has property $\mathcal{R}_n$ if the following two conditions hold:

(a) G is transitive on $\mathcal{L}_1$, and
(b) $\varphi_n(G) = \psi_n(\mathsf{st}_G(0))$.

Sometimes we need to refer to the group $\varphi_n(G)$ explicitly. If G has property $\mathcal{R}_n$ and we let $H = \varphi_n(G)$, then we say that G has property $\mathcal{R}_n(H)$.

The following fact is very useful:

LEMMA 2.23. *If a subgroup G of $\mathsf{Aut}(T_n)$ has property $\mathcal{R}_n$ then $G_{[m]}$ has property $\mathcal{R}_m(G_{[m-1]})$ for $2 \leq m \leq n$.*

PROOF. Suppose that G has property $\mathcal{R}_n$. For $2 \leq m \leq n$, it follows from the definition of φ_m that $\varphi_m(G_{[m]}) = G_{[m-1]}$. The conclusion therefore holds for $m = n$, and we proceed by induction on m (where m decreases to 2). Suppose that $G_{[m]}$ has property $\mathcal{R}_m(G_{[m-1]})$ for some m where $2 \leq m \leq n$. If $m = 2$ then we are done. Otherwise, we want to show that $G_{[m-1]}$ has property $\mathcal{R}_{m-1}(G_{[m-2]})$. We already know that $\varphi_m(G_{[m]}) = G_{[m-1]}$ and $G_{[m]}$ is transitive on $\mathcal{L}_1$ so $G_{[m-1]}$ is also transitive on $\mathcal{L}_1$. Since $\varphi_{m-1}(G_{[m-1]}) = G_{[m-2]}$, it remains only to show that $\psi_{m-1}\left(\mathsf{st}_{G_{[m-1]}}(0)\right) = G_{[m-2]}$. Using Lemma 2.21 and the inductive hypothesis,

$$\begin{aligned}
\psi_{m-1}\left(\mathsf{st}_{G_{[m-1]}}(0)\right) &= \psi_{m-1}\left(\varphi_m\left(\mathsf{st}_{G_{[m]}}(0)\right)\right) \\
&= \varphi_{m-1}\left(\psi_m\left(\mathsf{st}_{G_{[m]}}(0)\right)\right) \\
&= \varphi_{m-1}(G_{[m-1]}) \\
&= G_{[m-2]}
\end{aligned}$$

as required. By induction, the conclusion holds for $2 \leq m \leq n$. ☠

When we defined property $\mathcal{R}$, the condition ($\mathcal{R}2$) required G to be transitive on ∂T. Analogously, we want groups with property $\mathcal{R}_n$ to be transitive on the ends of T_n; that is, they must be transitive on $\mathcal{L}_n$. Lemma 2.24 below (an adaptation of (a) $\Longrightarrow$ (b) from Proposition 2.14) explains why our definition of property $\mathcal{R}_n$ only requires transitivity on $\mathcal{L}_1$.

LEMMA 2.24. *If $G \leq \mathsf{Aut}(T_n)$ has property $\mathcal{R}_n$ then G is transitive on $\mathcal{L}_n$.*

PROOF. Suppose that G has property $\mathcal{R}_n$. Then $G_{[m]}$ has property $\mathcal{R}_m(G_{[m-1]})$ for $2 \leq m \leq n$, by Lemma 2.23. Obviously $G_{[1]}$ is transitive on $\mathcal{L}_1$ since G is. The proof now proceeds by induction on m.

Suppose that $1 \leq m < n$ and $G_{[m]}$ is transitive on $\mathcal{L}_m$. We know that $G_{[m+1]}$ has property $\mathcal{R}_{m+1}(G_{[m]})$, and we must show that $G_{[m+1]}$ is transitive on $\mathcal{L}_{m+1}$. Firstly, $G_{[m+1]}$ is transitive on $\mathcal{L}_m$ by the inductive hypothesis and the fact that $\varphi_{m+1}(G_{[m+1]}) = G_{[m]}$. Now let v and w be vertices in $\mathcal{L}_{m+1}$. We must find $g \in G_{[m+1]}$ such that $g(v) = w$. Let v' and w' be the parents of v and w respectively, which means they are in

$\mathcal{L}_m$. Since $G_{[m+1]}$ is transitive on $\mathcal{L}_m$, and therefore transitive on $\mathcal{L}_1$, there exist $h_1, h_2 \in G_{[m+1]}$ such that $h_1(v') \in T_m^{(0)}$ and $h_2(w') \in T_m^{(0)}$. Hence $h_1(v) \in T_m^{(0)}$ and $h_2(w) \in T_m^{(0)}$ as well. Now $G_{[m]}$ is transitive on $\mathcal{L}_m$ so $\mathsf{st}_{G_{[m+1]}}(0)\Big|_{T_m^{(0)}} = \psi_{m+1}(\mathsf{st}_{G_{[m+1]}}(0)) = G_{[m]}$ is transitive on $\mathcal{L}_m$ of $T_m^{(0)}$. Therefore there exists $k \in \mathsf{st}_{G_{[m+1]}}(0)$ such that $k(h_1(v)) = h_2(w)$. Finally, let $g = h_2^{-1} k h_1$. Then $g \in G_{[m+1]}$ since h_1, h_2 and k are, and $g(v) = w$ as required.

2.3.6. Wreath recursion. If an automorphism g in $\mathsf{Aut}(T)$ fixes $\mathcal{L}_1$ (pointwise) then it follows from the defining equation for $g|_v$ that g is completely determined by the p-tuple $(g|_0, \ldots, g|_{p-1})$, and conversely such a p-tuple uniquely defines an automorphism fixing $\mathcal{L}_1$. If we take another automorphism h that fixes $\mathcal{L}_1$ then we have another p-tuple $(h|_0, \ldots, h|_{p-1})$. It follows from Proposition 2.12 that $(gh)|_v = g|_v h|_v$ for all $v \in \mathcal{L}_1$, so multiplication in $\mathsf{st}(\mathcal{L}_1)$ is pointwise multiplication of these p-tuples. Thus, $\mathsf{st}(\mathcal{L}_1)$ is isomorphic to $\mathsf{Aut}(T) \times \cdots \times \mathsf{Aut}(T)$. By Proposition 2.18, the quotient of $\mathsf{Aut}(T)$ by $\mathsf{st}(\mathcal{L}_1)$ is $\mathsf{Aut}(T_1)$ which is isomorphic to $\mathrm{Sym}(p)$. We can identify this quotient with a subgroup of $\mathsf{Aut}(T)$ as follows.

Let g be an arbitrary element of $\mathsf{Aut}(T)$. Define σ_g to be the automorphism of T such that for all $v \in \mathcal{L}_1$, $\sigma_g(v) = g(v)$ and $\sigma_g|_v = e$. In other words,

$$\sigma_g(vw) = g(v)w \quad \text{for all } v \in \mathcal{L}_1 \text{ and all } w \in T.$$

Each σ_g is completely determined by its action on $\mathcal{L}_1$. It is now easy to see that the set $S := \{\sigma_g : g \in \mathsf{Aut}(T)\}$ is a subgroup of $\mathsf{Aut}(T)$ isomorphic to $\mathsf{Aut}(T_1)$ and hence isomorphic to $\mathrm{Sym}(p)$. Putting all this together, we can express $\mathsf{Aut}(T)$ as a semi-direct product:

PROPOSITION 2.25. $\mathsf{Aut}(T)$ *is isomorphic to* $S \ltimes (\mathsf{Aut}(T) \times \cdots \times \mathsf{Aut}(T))$, *where there are* p *factors of* $\mathsf{Aut}(T)$ *in the product.*

PROOF. We have seen (Proposition 2.18) that $\mathsf{Aut}(T) \times \cdots \times \mathsf{Aut}(T) \cong \mathsf{st}(\mathcal{L}_1)$ is normal in $\mathsf{Aut}(T)$, and it follows from the definition of S that the intersection of S with $\mathsf{st}(\mathcal{L}_1)$ is trivial. Thus, it suffices to show that every element of $\mathsf{Aut}(T)$ can be expressed as a product of an element of S and and an element of $\mathsf{st}(\mathcal{L}_1)$.

Let $g \in \mathsf{Aut}(T)$ and define $\sigma_g \in S$ as above. It follows easily from the definition that $\sigma_g^{-1} g$ fixes $\mathcal{L}_1$, and $(\sigma_g^{-1} g)|_v = \sigma_g^{-1}|_{g(v)} g|_v = g|_v$ for all $v \in \mathcal{L}_1$. Thus, using the notation above for elements of $\mathsf{st}(\mathcal{L}_1)$, we may write

$\sigma_g^{-1} g = (g|_0, \ldots, g|_{p-1})$ and hence

$$g = \sigma_g(g|_0, \ldots, g|_{p-1}) \tag{2.13}$$

which expresses g as the required product. ☠

REMARK 2.26. This kind of semi-direct product is known as a *wreath product*, denoted $S \wr \mathsf{Aut}(T)$. For multiplication in the wreath product, the group $S \cong \mathrm{Sym}(p)$ acts on the direct product of p copies of $\mathsf{Aut}(T)$ by permuting the factors. It is straightforward to show that this is precisely how multiplication in the semi-direct product for $\mathsf{Aut}(T)$ works.

A similar argument works for $\mathsf{Aut}(T_n)$ as well, but instead we get:

$$\mathsf{Aut}(T_n) \cong S \ltimes (\mathsf{Aut}(T_{n-1}) \times \cdots \times \mathsf{Aut}(T_{n-1})). \tag{2.14}$$

and (2.13) holds as well. Additionally, for $g \in \mathsf{Aut}(T_n)$ this decomposition interacts with the restriction map φ_n via the simple formula

$$\varphi_n(\sigma_g(g|_0, \ldots, g|_{p-1})) = \sigma_g(\varphi_{n-1}(g|_0), \ldots, \varphi_{n-1}(g|_{p-1})). \tag{2.15}$$

We can also use (2.14) to find a formula for $|\mathsf{Aut}(T_n)|$. Since $|S| = |\mathrm{Sym}(p)| = p!$ we have

$$|\mathsf{Aut}(T_n)| = p!\, |\mathsf{Aut}(T_{n-1})|^p\,. \tag{2.16}$$

Using the fact that $\mathsf{Aut}(T_1) \cong S$, this recurrence relation can be easily solved to get:

$$|\mathsf{Aut}(T_n)| = p!^{1+p+\cdots+p^{n-1}} = p!^{\frac{p^n-1}{p-1}}\,. \tag{2.17}$$

REMARK 2.27. Since each copy of $\mathsf{Aut}(T_{n-1})$ in (2.14) is also a semi-direct product of the same form, we get the structure of $\mathsf{Aut}(T_n)$ as an *iterated wreath product* of n copies of S, written $S \wr (S \wr (\cdots \wr S))$. For more details of iterated wreath products and their relationship with rooted trees see [**OOR04**].

It follows from (2.13) that we may define an element of $\mathsf{Aut}(T)$ by specifying an element of S and a p-tuple of elements of $\mathsf{Aut}(T)$. This method is known as *wreath recursion*. It can indeed be literally recursive, since it is possible for $g|_v$ to be equal to g for one or more v. It can also happen that an automorphism h appears in the decomposition of g while g appears in the decomposition of h, creating a kind of mutual recursion. This is very common, for example, when defining automaton groups such as the ones in Chapter 7.

2.3.7. Inverse limits. Every group G that has property $\mathcal{R}$ is profinite, since it is a closed subgroup of the profinite group $\mathsf{Aut}(T)$. We can now give an explicit description of G as an inverse limit of finite groups with property $\mathcal{R}_n$. More importantly, we can use the inverse limit to construct a group with property $\mathcal{R}$ from a sequence of groups with property $\mathcal{R}_n$. This construction will be invoked frequently in later chapters.

Definition 2.28. Let $\{G_n\}_{n=1}^{\infty}$ be a sequence of groups such that G_n is a subgroup of $\mathsf{Aut}(T_n)$ for all n, and $\varphi_n(G_n) = G_{n-1}$ for all $n \geq 2$. Define:

$$G_\infty := \{g \in \mathsf{Aut}(T) \ : \ g_{[n]} \in G_n \text{ for all } n\}.$$

Proposition 2.29. *Let $\{G_n\}_{n=1}^{\infty}$ be as above. Then G_∞ is a closed subgroup of $\mathsf{Aut}(T)$.*

Proof. First we must show that G_∞ is a group. Clearly the identity is in G_∞. Suppose $g, h \in G_\infty$. Then for all n, $(g^{-1}h)_{[n]} = (g_{[n]})^{-1}(h_{[n]})$ which belongs to G_n since G_n is a group. Therefore $g^{-1}h \in G_\infty$. It follows that G_∞ is a subgroup of $\mathsf{Aut}(T)$.

To show that G_∞ is closed in $\mathsf{Aut}(T)$, note that every finite set of vertices in T is contained in T_m for some m, so a sequence $\{g_n\}_{n=1}^{\infty}$ converges to g in $\mathsf{Aut}(T)$ if and only if for each positive integer m there exists a positive integer N_m such that $g_n(v) = g(v)$ for all $n \geq N_m$ and all $v \in T_m$. That is, $(g_n)_{[m]} = g_{[m]}$ for all $n \geq N_m$.

Let $\{g_n\}_{n=1}^{\infty}$ be a sequence in G_∞ and suppose that $\lim_{n\to\infty} g_n = g$. Fix a positive integer m. Then $(g_n)_{[m]} \in G_m$ for all n by definition of G_∞. From the definition of convergence, there exists n such that $g_{[m]} = (g_n)_{[m]}$, hence $g_{[m]} \in G_m$. Since this holds for all m, we conclude that $g \in G_\infty$. Therefore G_∞ is closed. ☠

If $\{G_n\}_{n=1}^{\infty}$ is a sequence of groups as above then the surjective maps φ_n turn the sequence into an inverse system:

$$G_1 \xleftarrow{\varphi_1} G_2 \xleftarrow{\varphi_2} G_3 \xleftarrow{\varphi_3} \cdots$$

and so the inverse limit $\varprojlim G_n$ exists. It turns out that this inverse limit is precisely the group G_∞ we have just defined.

Proposition 2.30. *Let $\{G_n\}_{n=1}^{\infty}$ be a sequence of groups as in Definition 2.28. Then G_∞ is isomorphic, as a topological group, to the inverse limit $\varprojlim G_n$.*

Proof. See the proof of Proposition 2.6. Simply replace $X = \bigcup_{n\geq 0} X_{-n}$ with $T = \bigcup_{n\geq 0} T_n$ and a similar argument carries through. Note that Proposition 2.29 implies that G_∞ is compact since it is a closed subgroup of the compact group $\mathsf{Aut}(T)$. ☠

With this result in mind, from now on we will simply use the notation $\varprojlim G_n$ to refer to the group G_∞. This has the notational advantage of explicitly referring to the sequence index n, as well as reminding us of the profinite structure of the group.

The most important fact about the group $\varprojlim G_n$ — indeed, the purpose of its construction — is that it has property $\mathcal{R}$ whenever each of the G_n has property $\mathcal{R}_n$. Before we prove this, we need to establish a number of other important properties of $\varprojlim G_n$.

Proposition 2.31. *Let $\{G_n\}_{n=1}^\infty$ be as in Definition* 2.28*. Then:*

- **(a)** $(\varprojlim G_n)_{[n]} = G_n$ *for all n;*
- **(b)** *If G is any subgroup of $\mathsf{Aut}(T)$ then $\varprojlim G_{[n]} = \overline{G}$, the topological closure of G in $\mathsf{Aut}(T)$;*
- **(c)** *If G is a closed subgroup of $\mathsf{Aut}(T)$ such that $G_{[n]} = G_n$ for all n, then $G = \varprojlim G_n$;*
- **(d)** *If $\{H_n\}_{n=1}^\infty$ is another sequence of groups as in Definition* 2.28 *and $H_n \leq G_n$ for all n, then $\varprojlim H_n \leq \varprojlim G_n$ with equality if and only if $H_n = G_n$ for all n.*

Proof. **(a)** Follows from general facts about inverse limits (see [**RZ10**, Prop. 1.1.10]) since each φ_n is surjective.

(b) and **(c)** are also general results. See [**RZ10**, Corollary 1.1.8].

(d) The inequality follows immediately from the initial definition of G_∞. The 'if' part of the equality is obvious, and the 'only if' part follows from part (a). ☠

Now we are in a position to prove the main result of this section. Recall that the sequence $\{G_n\}_{n=1}^\infty$ in Definition 2.28 was subject to the condition $\varphi_n(G_n) = G_{n-1}$ for all n. This condition is subsumed by property $\mathcal{R}_n$, so if we assume instead that G_n has property $\mathcal{R}_n(G_{n-1})$ for each n then the group $\varprojlim G_n$ is still perfectly well defined.

Theorem 2.32. *Let $\{G_n\}_{n=1}^\infty$ be a sequence of groups such that G_n has property $\mathcal{R}_n(G_{n-1})$ for each $n \geq 2$. Then $\varprojlim G_n$ has property $\mathcal{R}$, and $(\varprojlim G_n)_{[n]} = G_n$ for all n. Furthermore, it is the unique subgroup of $\mathsf{Aut}(T)$ with these properties.*

PROOF. Let $\{G_n\}_{n=1}^{\infty}$ be as stated. It follows from Proposition 2.31 that $(\varprojlim G_n)_{[n]} = G_n$ for all n, and that $\varprojlim G_n$ is the only closed subgroup of $\mathsf{Aut}(T)$ with that property. Since being closed in $\mathsf{Aut}(T)$ is part of property $\mathcal{R}$, it remains only to show that $\varprojlim G_n$ is transitive on ∂T and that $\mathsf{st}_{\varprojlim G_n}(0)|_0 = \varprojlim G_n$.

Now $(\varprojlim G_n)_{[n]} = G_n$ is transitive on $\mathcal{L}_n$ for each n because each G_n has property $\mathcal{R}_n$. Therefore $\varprojlim G_n$ is transitive on each $\mathcal{L}_n$ as well, and hence it is transitive on ∂T by Proposition 2.14. For each n we have

$$\left(\mathsf{st}_{\varprojlim G_n}(0)|_0\right)_{[n]} = \left(\mathsf{st}_{\varprojlim G_n}(0)_{[n+1]}\right)\Big|_0 = \mathsf{st}_{G_{n+1}}(0)|_0 = G_n$$

since G_{n+1} has property $\mathcal{R}_{n+1}(G_n)$. Thus, by Proposition 2.31(c), the result will follow if we can show that $\mathsf{st}_{\varprojlim G_n}(0)|_0$ is closed, or equivalently that $\mathsf{st}_{\varprojlim G_n}(0)$ is closed in $\mathsf{Aut}(T)$. We already know that $\varprojlim G_n$ is closed, as is $\mathsf{st}(0)$ by definition of the topology on $\mathsf{Aut}(T)$, so the intersection $\mathsf{st}(0) \cap \varprojlim G_n = \mathsf{st}_{\varprojlim G_n}(0)$ must be closed as well. This completes the proof. ☠

REMARK 2.33. Suppose that we only have the sequence of groups G_n for $n \geq N$ for some N. Then we can still define the inverse limit $\varprojlim G_n$ as above, simply by adding the restrictions $(G_N)_{[n]}$ for $n < N$ to the beginning of the sequence. This is necessary, for example, in Section 6.3 for the construction of the infinite maximal group $\mathcal{M}^{\infty}(G)$.

CHAPTER 3

Groups with property $\mathcal{R}$: general results

3.1. Introduction

In this chapter we establish some basic facts about groups with property $\mathcal{R}$, as well as the finite groups with property $\mathcal{R}_n$ for some n. Section 3.2 looks at conjugacy of groups in $\mathsf{Aut}(T)$. We will see how to strengthen the property $\mathcal{R}$, without loss of generality, by choosing appropriate representatives from each conjugacy class. The main result is Theorem 3.18 which characterises groups with property $\mathcal{R}$, up to conjugacy, in terms of so-called *self-replicating* groups. This theorem facilitates many of the later results in this thesis, including the constructions in Chapters 6 and 7, and connects this work with the rich theory of self-similar groups.

The finite-depth property $\mathcal{R}_n$ can be strengthened in a similar way. In section 3.3 we use this to give concrete results about the relative sizes of the finite groups with property $\mathcal{R}_n$. Then, in section 3.4, we turn our attention to the special case of p-groups with property $\mathcal{R}_n$ and pro-p groups with property $\mathcal{R}$. We provide a simple criterion for a group with property $\mathcal{R}_n$ or property $\mathcal{R}$ to be a p-group or pro-p group respectively. We also describe the lattice of subgroups of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n(G)$ when G is a p-group. This last result will help to simplify the automated search (see Appendix A), particularly in the case $p = 2$.

3.2. Conjugacy and rigid automorphisms

In Chapter 2 we identified the rooted tree T with the words over the alphabet $\{0, 1, \ldots, p-1\}$. Note that we are not assuming that p is prime. This allocates a convenient label to each vertex of T, allowing us to refer to particular vertices and describe automorphisms of T explicitly. We must remember, however, that the rooted tree T is part of a homogeneous tree X, and the subgroups of $\mathsf{Aut}(T)$ with property $\mathcal{R}$ come from larger groups acting on X that include a common translation, α. This translation determines a preferred axis, whose vertices in T we have labelled $\varnothing$, 0, 00, 000, ..., where

$\alpha(0^n) = 0^{n+1}$ (of course, α is not an automorphism of T and this equation only makes sense if we think of T as part of X).

With the labelling of the axis fixed by α, there is still freedom to label the other vertices of T however we wish, as long as the edge relation is satisfied. For example, consider the simplest case where $p = 2$ and our alphabet is $\{0, 1\}$. The two vertices on $\mathcal{L}_1$ are labelled 0 and 1 and we have no choice in the matter because 0 is determined by α. On $\mathcal{L}_2$, the vertex 00 is determined and thus 01 must be its sibling; but the other two vertices can be labelled 10 and 11 as we please. In general, α will determine which vertices we label 0^n (and when $p = 2$ this also determines the vertices of the form 0^n1) and the other vertices can be labelled arbitrarily, consistent with the edge relation in T.

The reason this is important is that if we label the tree differently, the group actions will look different. We would like to identify two groups with property $\mathcal{R}$ as equivalent if they can be identified under a relabelling of T (preserving the edge relation). Since such a relabelling is simply a bijection on T that preserves the edge relation, it can be identified with a unique automorphism of T. It turns out that, from the point of view of a group G acting on T, relabelling T is equivalent to conjugating G by this automorphism:

Proposition 3.1. *Let ϕ be a relabelling of T and let G be a subgroup of* $\mathsf{Aut}(T)$. *Let x be the automorphism of T corresponding to ϕ. Then*

$$xgx^{-1}(\phi(v)) = \phi(g(v))$$

for all $g \in G$ and all $v \in T$. In other words, the action of xGx^{-1} on T with respect to the new labelling ϕ is equal to the action of G on T with respect to the original labelling.

Proof. The assertion is an immediate consequence of identifying x with ϕ as a mapping on the vertices of T. ☠

Consequently, the equivalence relation we would like to define on the groups with property $\mathcal{R}$ is simply conjugacy in $\mathsf{Aut}(T)$. We will therefore study the groups with property $\mathcal{R}$ up to conjugacy, and hence those with property $\mathcal{R}_n$ up to conjugacy in $\mathsf{Aut}(T_n)$. This means we are free to choose representatives from each conjugacy class, which raises the question: how should we determine which representatives to choose? As the results of this section will demonstrate, we can choose a representative from each conjugacy class that satisfies a stronger version of the condition $\mathcal{R}$. Recall that the property $\mathcal{R}$ includes a kind of self-similarity condition, namely that $\mathsf{st}_G(0)|_0 = G$, where equality occurs after the canonical identification of $T^{(0)}$ with T (i.e. using α as above). What about restriction to other subtrees? For example, how does $\mathsf{st}_G(1)|_1$ compare to G when $T^{(1)}$ is identified with T? What about $\mathsf{st}_G(v)|_v$ for arbitrary $v \in T$? As it turns out, we can prove that all such restrictions are conjugate to G.

PROPOSITION 3.2. *Suppose that $G \leq \mathsf{Aut}(T)$ has property $\mathcal{R}$ and let v be any vertex in T. Then $\mathsf{st}_G(v)|_v$ is conjugate to G, under the canonical identification of $T^{(v)}$ with T.*

PROOF. Since G has property $\mathcal{R}$, we know that $\mathsf{st}_G(0)|_0 = G$. Repeatedly applying this identification we conclude that for all n, $\mathsf{st}_G(0^n)|_{0^n} = G$ when the subtree $T^{(0^n)}$ is identified with T. Now fix $v \in T$. Then $v \in \mathcal{L}_n$ for some n. Let x be any element of G such that $x(0^n) = v$. Such an x exists since $0^n \in \mathcal{L}_n$ and G is transitive on $\mathcal{L}_n$. Then a simple calculation shows that $\mathsf{st}_G(v) = x\,(\mathsf{st}_G(0^n))\,x^{-1}$. Now, using the formulae for vertex restrictions, we obtain:

$$\begin{aligned}
\mathsf{st}_G(v)|_v &= \left(x\,(\mathsf{st}_G(0^n))\,x^{-1}\right)|_v \\
&= (x\,\mathsf{st}_G(0^n))\,|_{x^{-1}(v)}\,x^{-1}|_v \\
&= (x\,\mathsf{st}_G(0^n))\,|_{0^n}\left(x|_{x^{-1}(v)}\right)^{-1} \\
&= (x|_{0^n})\,\mathsf{st}_G(0^n)|_{0^n}\,(x|_{0^n})^{-1} \\
&= (x|_{0^n})\,G\,(x|_{0^n})^{-1}
\end{aligned}$$

which is a conjugate of G, as claimed. ☠

It turns out that if we are careful, we can choose a representative from the conjugacy class of G so that all of these restrictions are actually *equal* to G. The key to this is to conjugate G by an appropriately “nice” element of $\mathsf{Aut}(T)$. By “nice” we mean an automorphism that moves a vertex v, say, but acts trivially on the subtree below v. In other words, an element $g \in \mathsf{Aut}(T)$ such that the vertex restriction $g|_v$ is trivial. We will call such automorphisms *rigid*.

DEFINITION 3.3. Let u and v be vertices of T and let g be an automorphism of T. Then g is *rigid at u* if the restriction $g|_u$ is trivial. Further, we say that g is *(u, v)-rigid* if $g|_u$ is trivial and $g(u) = v$; this means that $g(uw) = vw$ for all $w \in T$.

These definitions also apply *mutatis mutandis* to automorphisms and vertices of the finite trees T_n.

REMARK. It might be necessary to find another name for these rigid automorphisms if they are to be used in the future, to avoid potential confusion with the concept of the *rigid stabiliser* of a vertex $v \in T$ — this is the subgroup of $\mathsf{Aut}(T)$ that fixes v and all vertices *outside* the subtree $T^{(v)}$. On the other hand, an automorphism g that is rigid at v need not fix v, but fixes (in a sense, if the subtrees $T^{(g(v))}$ and $T^{(v)}$ are identified) the vertices *inside* $T^{(v)}$.

Rigid automorphisms are extremely useful tools for dealing with vertex restrictions. Conjugation by an automorphism rigid at v has the effect of moving the vertex restriction at v to another vertex:

LEMMA 3.4. *Let v be a vertex of T. Suppose that $g \in \mathsf{Aut}(T)$ fixes v and $h \in \mathsf{Aut}(T)$ is rigid at v. Then $(hgh^{-1})|_{h(v)} = g|_v$.*

PROOF. The hypotheses imply that $g(v) = v$ and $h|_v = e$. A direct calculation gives the result:

$$(hgh^{-1})|_{h(v)} = (hg)|_v \, h^{-1}|_{h(v)} = h|_{g(v)} g|_v (h|_v)^{-1} = h|_v g|_v = g|_v.$$

Our approach will be to use these rigid automorphisms to move vertex restrictions from 0 (or more generally 0^n) to other vertices of T, thereby showing that all restrictions are equal. The easiest way to do this is one level at a time; that is, we will arrive at the result for groups with property $\mathcal{R}$ by proving it for groups with property $\mathcal{R}_n$ for each n.

First, we must establish some properties of conjugacy in both $\mathsf{Aut}(T_n)$ and $\mathsf{Aut}(T)$. In particular, we need to know when conjugacy preserves property $\mathcal{R}_n$. The following result gives sufficient conditions for this:

PROPOSITION 3.5. *Suppose that G is a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$, and let $x \in \mathsf{st}(0)$ such that $\varphi_n(x) = \psi_n(x)$. Then $x^{-1}Gx$ also has property $\mathcal{R}_n$.*

Further, let $H = G_{[n-1]}$ so that G has property $\mathcal{R}_n(H)$, and suppose that $\varphi_n(x) = \psi_n(x) = e$. Then $x^{-1}Gx$ also has property $\mathcal{R}_n(H)$.

PROOF. Firstly, since G is transitive on $\mathcal{L}_1$, the same is true for $x^{-1}Gx$ for any $x \in \mathsf{Aut}(T_n)$.

Let $x \in \mathsf{st}(0)$ and suppose that $\varphi_n(x) = \psi_n(x)$. We know that $\varphi_n(G) = \psi_n(\mathsf{st}_G(0))$, and we must show that this equality holds for $x^{-1}Gx$. First, note that x fixes 0 so $x^{-1}\mathsf{st}_G(0)x = \mathsf{st}_{x^{-1}Gx}(0)$. The claim now follows immediately from the fact that both φ_n and ψ_n are homomorphisms.

Now let H be as stated and suppose that $\varphi_n(x) = \psi_n(x) = e$, i.e. that $x \in \ker\varphi_n \cap \ker\psi_n$. Since φ_n and ψ_n are homomorphisms, we obtain

$$\varphi_n(x^{-1}Gx) = \varphi_n(G) = H$$

and since $x^{-1}Gx$ has property $\mathcal{R}_n$ it follows that

$$\psi_n(\mathsf{st}_{x^{-1}Gx}(0)) = \varphi_n(x^{-1}Gx) = H$$

so $x^{-1}Gx$ has property $\mathcal{R}_n(H)$ as required. ☠

We will also need to know how conjugacy in $\mathsf{Aut}(T_n)$ for each n carries through the inverse limit to conjugacy in $\mathsf{Aut}(T)$, if indeed it does at all. Fortunately, it does — although the proof invokes the topology on $\mathsf{Aut}(T)$ and it would not work without the condition that groups with property $\mathcal{R}$ are closed.

Proposition 3.6. *Let G and G' be subgroups of $\mathsf{Aut}(T)$ with property $\mathcal{R}$ and suppose that $G_{[n]}$ is conjugate to $G'_{[n]}$ in $\mathsf{Aut}(T_n)$ for all n. Then G is conjugate to G'.*

Proof. Let G and G' be as stated. For each $n \geq 1$, define the following sets:

$$C_n = \{x \in \mathsf{Aut}(T_n) \;:\; x^{-1}G_{[n]}x = G'_{[n]}\}$$
$$\text{and} \qquad \widetilde{C}_n = \{x \in \mathsf{Aut}(T) \;:\; (x^{-1}Gx)_{[n]} = G'_{[n]}\}.$$

Our assumptions imply that each C_n is nonempty. Our conclusion will follow if we can show that the intersection $\bigcap_{n=1}^{\infty} \widetilde{C}_n$ is nonempty. To see why, suppose that there is an x in this intersection. Then $x \in \widetilde{C}_n$ for all n, so $(x^{-1}Gx)_{[n]} = G'_{[n]}$ for all n. Now G is closed in $\mathsf{Aut}(T)$, hence so is $x^{-1}Gx$. Therefore, by Proposition 2.31(c),

$$x^{-1}Gx = \varprojlim (x^{-1}Gx)_{[n]} = \varprojlim G'_{[n]} = G'$$

since G' is also closed in $\mathsf{Aut}(T)$, proving that G is conjugate to G'.

We will show that the intersection of the $\widetilde{C}_n$ is nonempty by invoking the finite intersection property in the compact group $\mathsf{Aut}(T)$. It is clear that $\widetilde{C}_{n+1} \subseteq \widetilde{C}_n$ for all n, so to show that the $\widetilde{C}_n$ have the finite intersection property we need only show that each $\widetilde{C}_n$ is nonempty. To do this, we will show that, for all n,

$$\widetilde{C}_n = \{x \in \mathsf{Aut}(T) \;:\; x_{[n]} \in C_n\}. \tag{3.1}$$

Fix $n \geq 1$. If $x \in \mathsf{Aut}(T)$ and $x_{[n]} \in C_n$ then $(x^{-1}Gx)_{[n]} = x_{[n]}^{-1}G_{[n]}x_{[n]} = G'_{[n]}$ so $x \in \widetilde{C}_n$. Conversely, if $x \in \widetilde{C}_n$ then $G'_{[n]} = (x^{-1}Gx)_{[n]} = x_{[n]}^{-1}G_{[n]}x_{[n]}$ which implies that $x_{[n]} \in C_n$. Therefore (3.1) holds. In fact, this actually shows that $(\widetilde{C}_n)_{[n]} = C_n$ for all n. We have assumed that each C_n is nonempty, so it follows that each $\widetilde{C}_n$ is also nonempty and thus the $\widetilde{C}_n$ have the finite intersection property.

Now let us show that each $\widetilde{C}_n$ is closed in $\mathsf{Aut}(T)$. Recall that $\mathsf{Aut}(T)$ is homeomorphic to the inverse limit $\varprojlim \mathsf{Aut}(T_n)$. Let θ_n be the projection onto $\mathsf{Aut}(T_n)$ for each n; that is, $\theta_n(g) = g_{[n]}$ for all $g \in \mathsf{Aut}(T)$. It follows from (3.1) that $\widetilde{C}_n = \theta_n^{-1}(C_n)$ for all n. The topology in the inverse limit means that each θ_n is continuous. Since $\mathsf{Aut}(T_n)$ is finite and therefore discrete, each C_n is closed in $\mathsf{Aut}(T_n)$, so the continuity of θ_n implies that $\widetilde{C}_n$ is closed in $\mathsf{Aut}(T)$ for each n.

In summary, $\{\widetilde{C}_n\}_{n=1}^{\infty}$ is a family of nonempty, closed sets with the finite intersection property. Since $\mathsf{Aut}(T)$ is compact, it follows that the intersection $\bigcap_{n=1}^{\infty} \widetilde{C}_n$ is nonempty as required. ☠

Having established the important properties of conjugacy, we are going to show how conjugation by appropriate rigid automorphisms allows us to strengthen the properties $\mathcal{R}$ and $\mathcal{R}_n$ without loss of generality; i.e. by choosing an appropriate conjugacy class representative. The desired representatives all contain "sufficiently many" rigid automorphisms, in the following precise sense:

Definition 3.7. Let G be a subgroup of $\mathsf{Aut}(T)$. We say that G *has sufficient rigid automorphisms* if for each pair of vertices u and v in $\mathcal{L}_1$ there exists an automorphism $g \in G$ such that g is (u, v)-rigid; that is, $g(u) = v$ and $g|_u = e$.

We can immediately rework this definition to make it easier to use, by reducing it to the case where $v = u + 1$. Recall that the vertices on $\mathcal{L}_1$ are identified with the set $\{0, 1, \ldots, p-1\}$ so the expression $u + 1$ makes sense, provided $u \leq p - 2$.

Lemma 3.8. *Let G be a subgroup of $\mathsf{Aut}(T)$. Suppose that for each vertex $v \in \{0, 1, \ldots, p-2\}$ on $\mathcal{L}_1$, there exists an automorphism $g \in G$ such that g is $(v, v+1)$-rigid. Then G has sufficient rigid automorphisms.*

Proof. Suppose that G contains a $(v, v+1)$-rigid automorphism for each $v \in \{0, 1, \ldots, p-2\}$, and let u and v be fixed vertices in $\mathcal{L}_1$. We must show that there exists $g \in G$ such that $g(u) = v$ and $g|_u = e$.

If $u = v$ then the identity in G is (u, v)-rigid so we are done. Suppose that $u < v$. Then, by our assumption, for all w such that $u \le w \le v - 1$, there exists an automorphism $g_w \in G$ such that $g_w(w) = (w + 1)$ and $g_w|_w = e$. Define $g = g_{v-1} \cdots g_u$ to be the product of all these g_w, which belongs to G since G is a group. Then it follows that $g(u) = v$ and $g|_u = g_{v-1}|_{v-1} \cdots g_u|_u = e$, which means that g is (u, v)-rigid.

Suppose now that $u > v$. By a similar argument there exists $g \in G$ that is (v, u)-rigid. Then $g^{-1}(u) = v$ and $g^{-1}|_u = (g|_{g^{-1}(u)})^{-1} = (g|_v)^{-1} = e$ so g^{-1} is (u, v)-rigid and the proof is complete. ☠

Note that Definition 3.7 and Lemma 3.8 apply also to automorphisms of the finite trees T_n and subgroups of $\mathsf{Aut}(T_n)$, since they refer only to $\mathcal{L}_1$ of the tree.

The next step is to use the above to show that if a subgroup H of $\mathsf{Aut}(T_{n-1})$ has sufficient rigid automorphisms then so does a conjugacy class representative of every group with property $\mathcal{R}_n(H)$.

Proposition 3.9. *Let $n \ge 2$ and let H be a subgroup of $\mathsf{Aut}(T_{n-1})$ that has sufficient rigid automorphisms. Suppose that G is a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n(H)$. Then there exists a conjugate of G in $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n(H)$ that has sufficient rigid automorphisms.*

Proof. Let H and G be as stated. The idea is to start with G and (carefully) construct a series of conjugates of G, producing more rigid automorphisms at each step until we can invoke Lemma 3.8. More precisely, we claim that, for each integer r such that $0 \le r \le p - 1$, there exists a subgroup $G^{(r)}$ of $\mathsf{Aut}(T_n)$ such that:

- **(a)** $G^{(r)}$ is conjugate to G in $\mathsf{Aut}(T_n)$;
- **(b)** $G^{(r)}$ has property $\mathcal{R}_n(H)$; and
- **(c)** for each $v \in \mathcal{L}_1$ such that $0 \le v \le r - 1$, there is a $(v, v + 1)$-rigid automorphism in $G^{(r)}$.

By Lemma 3.8, it suffices to prove this claim for the case $r = p - 1$. The proof is by induction on r. Starting with $r = 0$, the claim is trivially satisfied by G itself. Suppose that the claim is true for some $r \ge 0$, so that there exists a group $G^{(r)}$ with properties (a)–(c). If $r = p - 1$ then we are done so suppose further that $r \le p - 2$ (hence $r + 1$ is a valid vertex in $\mathcal{L}_1$).

We must find a group $G^{(r+1)}$ with properties (a)–(c). Any conjugate of the existing group $G^{(r)}$ will satisfy (a). If we conjugate by an $x \in \mathsf{Aut}(T_n)$ such that $\varphi_n(x) = \psi_n(x) = e$ then (b) is satisfied by Proposition 3.5. To satisfy (c) we must preserve the existing rigid automorphisms in $G^{(r)}$ while creating a new $(r, r+1)$-rigid automorphism. We have assumed that H contains such an automorphism; call it h. Then $h(r) = r+1$ and $h|_r = e$. Since $G^{(r)}$ has property $\mathcal{R}_n(H)$, there exists an automorphism $g \in G^{(r)}$ such that $\varphi_n(g) = h$. Therefore $g(r) = r+1$ and $\varphi_{n-1}(g|_r) = h|_r = e$.

Now define the automorphism $x \in \mathsf{Aut}(T_n)$ as follows:

$$x(v) = v \text{ for all } v \in \mathcal{L}_1;$$

$$x|_v = \begin{cases} g|_r & \text{if } v = r+1 \\ e & \text{if } v \neq r+1. \end{cases}$$

This means that x acts trivially everywhere except on the subtree rooted at $r+1 \in \mathcal{L}_1$, where it acts as $g|_r$. In particular, $\varphi_n(x) = \psi_n(x) = e$ so conjugating $G^{(r)}$ by x will preserve property $\mathcal{R}_n(H)$ by Proposition 3.5. With this in mind, define $G^{(r+1)} = x^{-1}G^{(r)}x$. By assumption (c) on $G^{(r)}$, for each $i = 0, 1, \ldots, r-1$ there exists $g_i \in G^{(r)}$ which is $(i, i+1)$-rigid. Define $\widetilde{g}_i = x^{-1}g_i x$ for each i, and $\widetilde{g}_r = x^{-1}gx$. Clearly each of these $\widetilde{g}_i$ belong to $G^{(r+1)}$. We claim that each $\widetilde{g}_i$ is $(i, i+1)$-rigid. First, we have $\widetilde{g}_i(i) = i+1$ since x fixes $\mathcal{L}_1$ and $g_i(i) = i+1$ for each i. For the vertex restrictions, if $0 \le i \le r-1$ we have

$$\widetilde{g}_i|_i = (x^{-1}g_i x)|_i = (x^{-1}g_i)|_i\, x|_i = (x^{-1}|_{i+1})\, g_i|_i\, x|_i = (x|_{i+1})^{-1} g_i|_i\, x|_i = e.$$

The last equality follows because $x|_i$ and $x|_{i+1}$ are both trivial (since neither i nor $i+1$ is equal to $r+1$) and $g_i|_i$ is trivial (since g_i is rigid at i). If $i = r$ a similar calculation shows that

$$\widetilde{g}_r|_r = (x^{-1}gx)|_r = (x|_{r+1})^{-1} g|_r\, x|_r = e$$

since $x|_{r+1} = g|_r$ and $x|_r = e$.

We have proven that $\widetilde{g}_i$ is $(i, i+1)$-rigid for all $0 \le i \le r$ which means $G^{(r+1)}$ has the required properties (a)–(c) and the inductive step is complete. Finally, the group $G^{(p-1)}$ is the required conjugate of G. ☠

Proposition 3.9 applies to groups with property $\mathcal{R}_n(H)$, assuming that H already has sufficient rigid automorphisms. It turns out that we can actually drop that assumption, as long as we are only interested in whether the group has property $\mathcal{R}_n$ rather than $\mathcal{R}_n(H)$ for a specific H.

Proposition 3.10. *Suppose that G is a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$. Then there exists an automorphism $x \in \mathsf{Aut}(T_n)$ where $x \in \mathsf{st}(0)$*

and $\varphi_n(x) = \psi_n(x)$, such that $x^{-1}Gx$ has property $\mathcal{R}_n$ and has sufficient rigid automorphisms.

Proof. The proof is by induction on n. The base case $n = 1$ is easy since all vertex restrictions are trivial for $g \in \mathsf{Aut}(T_1)$. Thus every element of $\mathsf{Aut}(T_1)$ is trivially rigid, so G has sufficient rigid automorphisms since it is transitive on $\mathcal{L}_1$. We may therefore take $x = e$.

Suppose that the assertion holds for all groups with property $\mathcal{R}_n$. Let G be a subgroup of $\mathsf{Aut}(T_{n+1})$ with property $\mathcal{R}_{n+1}$. Then $G_{[n]}$ has property $\mathcal{R}_n$ so the inductive hypothesis implies that there exists an $x \in \mathsf{Aut}(T_n)$ where $x \in \mathsf{st}(0)$ and $\varphi_n(x) = \psi_n(x)$, such that $x^{-1}G_{[n]}x$ has property $\mathcal{R}_n$ and has sufficient rigid automorphisms.

Let y be any element of $\mathsf{Aut}(T_{n+1})$ such that $\varphi_{n+1}(y) = \psi_{n+1}(y) = x$ (note that $\varphi_{n+1}(y) = x$ means that $y \in \mathsf{st}(0)$ so $\psi_{n+1}(y)$ is defined). We will first verify that such a y exists. The condition $\varphi_{n+1}(y) = x$ specifies the action of y on T_n, leaving us free to define its action on $\mathcal{L}_{n+1}$. Now $\psi_{n+1}(y) = y|_0$ and $\varphi_n(y|_0) = \varphi_{n+1}(y)|_0 = x|_0 = \psi_n(x) = \varphi_n(x)$ is already specified. This means that $y|_0$ agrees with x on T_{n-1}, so we may put $y|_0 = x$. This defines y on the part of $\mathcal{L}_{n+1}$ below the vertex 0. If we choose the action of y on the rest of $\mathcal{L}_{n+1}$ arbitrarily, we still have $\varphi_{n+1}(y) = x$ and now we also have $\psi_{n+1}(y) = y|_0 = x$ as desired.

Let $G' = y^{-1}Gy$. Then G' has property $\mathcal{R}_{n+1}$ by Proposition 3.5. In fact, it has property $\mathcal{R}_{n+1}(x^{-1}G_{[n]}x)$ since

$$\varphi_{n+1}(G') = \varphi_{n+1}(y)^{-1}\varphi_{n+1}(G)\varphi_{n+1}(y) = x^{-1}G_{[n]}x.$$

Since $x^{-1}G_{[n]}x$ has sufficient rigid automorphisms, Proposition 3.9 implies that there is a conjugate G'' of G' that has property $\mathcal{R}_{n+1}(x^{-1}G_{[n]}x)$ and also has sufficient rigid automorphisms. A careful reading of the proof of Proposition 3.9 shows that the conjugating element $z \in \mathsf{Aut}(T_n)$ satisfies $\varphi_{n+1}(z) = \psi_{n+1}(z) = e$. That is,

$$G'' = z^{-1}G'z = z^{-1}(y^{-1}Gy)z = (yz)^{-1}G(yz)$$

and $\varphi_{n+1}(yz) = \psi_{n+1}(yz)$ since both φ_{n+1} and ψ_{n+1} are homomorphisms. The automorphism yz has the properties we want, so the inductive step is complete and so is the proof. ☠

The upshot of these results is that we may now assume without loss of generality (that is, up to conjugacy) that every group with property $\mathcal{R}_n$ has sufficient rigid automorphisms. The benefit of this is illustrated by

Proposition 3.11 and its Corollary 3.12, which will allow us to impose a stronger self-similarity condition on groups with property $\mathcal{R}_n$.

PROPOSITION 3.11. *Suppose that G is a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n(H)$ that has sufficient rigid automorphisms. Then $g|_v \in H$ for all $g \in G$ and all $v \in \mathcal{L}_1$.*

PROOF. Let $g \in G$ and $v \in \mathcal{L}_1$. We aim to use the condition $\psi_n(\mathsf{st}_G(0)) = H$ to show that $g|_v \in H$. The first step is to find $h \in G$ such that $h(v) = v$ and $h|_v = g|_v$. If $g(v) = v$ already then simply put $h = g$. Otherwise, since G has sufficient rigid automorphisms, there exists $x \in G$ such that $x(g(v)) = v$ and $x|_{g(v)} = e$. In other words, x is a $(g(v), v)$-rigid automorphism. Now let $h = xg$ which is in G since both x and g are in G. Then $h(v) = x(g(v)) = v$ and $h|_v = x|_{g(v)}g|_v = g|_v$ as desired.

The next step is to find $k \in G$ such that $k(0) = 0$ and $k|_0 = g|_v$. If $v = 0$ then simply put $k = h$. Otherwise, since H has sufficient rigid automorphisms, there exists $y \in G$ such that $y(v) = 0$ and $y|_v = e$. Let $k = yhy^{-1}$, which is in G since both h and y are in G. Then $k(0) = yhy^{-1}(0) = yh(v) = y(v) = 0$ and Lemma 3.4 implies that $k|_0 = (yhy^{-1})|_{y(v)} = h|_v = g|_v$. Now $k \in \mathsf{st}_G(0)$ so $\psi_n(k) \in H$, but $\psi_n(k) = k|_0 = g|_v$ and thus $g|_v \in H$ as required. ☠

COROLLARY 3.12. *Suppose that G is a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$. Then there exists a conjugate G' of G in $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$ such that $g|_v \in G'_{[n-1]}$ for all $g \in G'$ and all $v \in \mathcal{L}_1$.*

Further, we may assume that the conjugating element x belongs to $\mathsf{st}(0)$ and satisfies $\varphi_n(x) = \psi_n(x)$.

PROOF. Follows immediately from Propositions 3.10 and 3.11. ☠

We may therefore, without loss of generality, strengthen property $\mathcal{R}_n$ when convenient, by adding the condition:

$$g|_v \in G_{[n-1]} \text{ for all } g \in G \text{ and all } v \in \mathcal{L}_1. \tag{3.2}$$

Note that the existing condition $\mathcal{R}_n$ implies that $g|_0 \in G_{[n-1]}$ for all $g \in \mathsf{st}_G(0)$ because $\psi_n(g) = g|_0$ for $g \in \mathsf{st}(0)$. The new condition (3.2) asserts that $g|_v \in G_{[n-1]}$ even when $g \notin \mathsf{st}(v)$. This is the critical ingredient in the construction of the maximal groups in Chapter 6 and the automata in Chapter 7.

Returning now to the infinite tree T, we can strengthen the self-similarity condition in property $\mathcal{R}$ in an analogous way to $\mathcal{R}_n$.

PROPOSITION 3.13. *Suppose that $G \leq \mathsf{Aut}(T)$ has property $\mathcal{R}$. Then there exists a group $\widehat{G}$ with property $\mathcal{R}$ that is conjugate to G in $\mathsf{Aut}(T)$, such that $g|_v \in \widehat{G}$ for all $g \in \widehat{G}$ and all $v \in T$.*

PROOF. Since G has property $\mathcal{R}$, we know that $G_{[n]}$ has property $\mathcal{R}_n$ for each n. Then Proposition 3.10 implies that for each $n \geq 1$ there exists G'_n conjugate to $G_{[n]}$ that has property $\mathcal{R}_n$ and has sufficient rigid automorphisms. The inductive step in the proof of Proposition 3.10 actually shows that we may take each G'_n to have property $\mathcal{R}_n(G'_{n-1})$ for all $n \geq 2$. This allows us to define the subgroup $\widehat{G}$ of $\mathsf{Aut}(T)$ by

$$\widehat{G} = \varprojlim G'_n.$$

Then Theorem 2.32 tells us that $\widehat{G}$ has property $\mathcal{R}$ and $\widehat{G}_{[n]} = G'_n$ for all n. This means that $\widehat{G}_{[n]}$ is conjugate to $G_{[n]}$ for all n, so we conclude from Proposition 3.6 that $\widehat{G}$ is conjugate to G in $\mathsf{Aut}(T)$.

It remains to show that $g|_v \in \widehat{G}$ for all $g \in \widehat{G}$ and all $v \in T$. Let us do this first for $v \in \mathcal{L}_1$. Fix $v \in \mathcal{L}_1$ and let $g \in \widehat{G}$. For each $n \geq 1$, $(g|_v)_{[n]} = g_{[n+1]}|_v \in \widehat{G}_{[n]}$ since $g_{[n+1]} \in \widehat{G}_{[n+1]} = G'_{n+1}$ which has sufficient rigid automorphisms, so Proposition 3.11 applies. Since this is true for all n, we conclude that $g|_v \in \varprojlim \widehat{G}_{[n]} = \widehat{G}$.

This forms the base step for an induction argument. Suppose that $g|_v \in \widehat{G}$ for all $g \in \widehat{G}$ and all $v \in \mathcal{L}_n$ for some $n \geq 1$. Now fix $v \in \mathcal{L}_{n+1}$. Then v can be expressed uniquely as $v = uw$ where $u \in \mathcal{L}_n$ and $w \in \mathcal{L}_1$. Fix $g \in \widehat{G}$. Then $g|_v = g|_{uw} = (g|_u)|_w$. By the inductive hypothesis $g|_u \in \widehat{G}$, and since $w \in \mathcal{L}_1$ it follows from the base step that $(g|_u)|_w \in \widehat{G}$ which completes the induction and the proof. ☠

The group $\widehat{G}$ in Proposition 3.13 belongs to a widely-studied class of groups called *self-similar* groups. A subset of these, called *self-replicating* groups, were foreshadowed in Chapter 2 and are the focus of the main result of this chapter. The definitions are as follows:

DEFINITION 3.14. A subgroup G of $\mathsf{Aut}(T)$ is *self-similar* if

$$G|_v \subseteq G \quad \text{for all } v \in T$$

and G is *self-replicating* if

$$\mathsf{st}_G(v)|_v = G \quad \text{for all } v \in T.$$

We can immediately rephrase Proposition 3.13 in these terms:

COROLLARY 3.15. *Suppose that G has property $\mathcal{R}$. Then G is conjugate in $\mathsf{Aut}(T)$ to a self-similar group with property $\mathcal{R}$.*

Recall that property $\mathcal{R}$ includes the condition $\mathsf{st}_G(0)|_0 = G$. The next step is to use Proposition 3.13 to replace the vertex 0 with any vertex in T without loss of generality. Then we will be able to replace 'self-similar' in Corollary 3.15 with 'self-replicating'.

PROPOSITION 3.16. *Suppose that G has property $\mathcal{R}$. Then G is conjugate in $\mathsf{Aut}(T)$ to a self-replicating group with property $\mathcal{R}$.*

PROOF. Define the same $\widehat{G}$ as in Proposition 3.13. Fix $v \in T$. Since $\widehat{G}$ has property $\mathcal{R}$, Proposition 3.2 tells us that $\mathsf{st}_{\widehat{G}}(v)|_v$ is conjugate to $\widehat{G}$. It follows that $(\mathsf{st}_{\widehat{G}}(v)|_v)_{[n]}$ is conjugate to $\widehat{G}_{[n]}$ for all n, and hence

$$\left|\left(\mathsf{st}_{\widehat{G}}(v)|_v\right)_{[n]}\right| = \left|\widehat{G}_{[n]}\right| \tag{3.3}$$

for all n. We also conclude from Proposition 3.13 that $\mathsf{st}_{\widehat{G}}(v)|_v \subseteq \widehat{G}$, whence $(\mathsf{st}_{\widehat{G}}(v)|_v)_{[n]} \subseteq \widehat{G}_{[n]}$ for all n. Thus it follows from (3.3) that $(\mathsf{st}_{\widehat{G}}(v)|_v)_{[n]} = \widehat{G}_{[n]}$ for all n. Now $\mathsf{st}_{\widehat{G}}(v)|_v$ is closed in $\mathsf{Aut}(T)$ since it is conjugate to $\widehat{G}$ which is closed in $\mathsf{Aut}(T)$. Therefore by Proposition 2.31(c),

$$\mathsf{st}_{\widehat{G}}(v)|_v = \varprojlim (\mathsf{st}_{\widehat{G}}(v)|_v)_{[n]} = \varprojlim \widehat{G}_{[n]} = \widehat{G}$$

as required. ☠

Since the self-replicating group $\widehat{G}$ we constructed in the proof of Proposition 3.16 is actually the same self-similar group $\widehat{G}$ from Proposition 3.13, we may immediately conclude:

COROLLARY 3.17. *Suppose that G is self-similar and has property $\mathcal{R}$. Then G is self-replicating.*

It follows immediately from the definitions that every closed, spherically transitive, self-replicating group has property $\mathcal{R}$. This result does not quite give us the converse of Proposition 3.16 but it does yield the following theorem — the main result of this chapter — which characterises groups with property $\mathcal{R}$ up to conjugacy:

THEOREM 3.18. *Let G be a subgroup of $\mathsf{Aut}(T)$. Then G is conjugate to a group with property $\mathcal{R}$ if and only if G is conjugate to a closed, spherically transitive, self-replicating group.*

Thus, the study of groups with property $\mathcal{R}$ is equivalent, up to conjugacy, to the study of closed, spherically transitive, self-replicating groups.

3.3. Sizes of groups with property $\mathcal{R}_n$

Let $n \geq 2$ and let H be a subgroup of $\mathsf{Aut}(T_{n-1})$ with property $\mathcal{R}_{n-1}$. The self-similarity condition in property $\mathcal{R}_n$ allows us, given any group G with property $\mathcal{R}_n(H)$, to relate $|G|$ to $|H|$. The kernels of the restriction maps φ_n and ψ_n play a crucial role in this; if we know the orders of these kernels in G then we can give an exact result for $|G|$ in terms of $|H|$. If not, then we can still use information about H to place upper and lower bounds on $|G|$.

In this section, to make the proofs easier, we will invoke Corollary 3.12 and assume (3.2) holds in addition to property $\mathcal{R}_n$. However, it is worth pointing out that the results remain valid with property $\mathcal{R}_n$ alone, since conjugacy does not affect the size of the group. As in the previous section, we are not assuming that p is prime.

PROPOSITION 3.19. *Suppose that $G \leq \mathsf{Aut}(T_n)$ satisfies $\mathcal{R}_n(H)$. Then*

- **(a)** $|G| = |\ker_G(\varphi_n)|\,|H| = p\left|\ker_{\mathsf{st}_G(0)}(\psi_n)\right|\,|H|$.
- **(b)** $\ker_G(\varphi_n)$ *is nontrivial.*
- **(c)** $|G| \geq p\,|H|$.
- **(d)** $|G| \leq |\ker_H(\varphi_{n-1})|^p\,|H| \leq \left|H_{[1]}\right|\,|H|^p$.

PROOF. Parts (b) and (c) follow from part (a). Let us prove (a). Since G satisfies $\mathcal{R}_n(H)$, we have $\varphi_n(G) = H$ which implies the first equality. Also, $\psi_n(\mathsf{st}_G(0)) = H$ implies that $|\mathsf{st}_G(0)| = \left|\ker_{\mathsf{st}_G(0)}(\psi_n)\right|\,|H|$. Since G is transitive on $\mathcal{L}_1$, it follows from the orbit-stabiliser theorem that $|G| = p\,|\mathsf{st}_G(0)|$ and the second equality in (a) follows.

There are two inequalities to prove in part (d). The first will follow from part (a) if we can show that $|\ker_G(\varphi_n)| \leq |\ker_H(\varphi_{n-1})|^p$. This can be done by finding an injection from $\ker_G(\varphi_n)$ into the product $(\ker_H(\varphi_{n-1}))^p$. Let $x \in \ker_G(\varphi_n)$. Since $\varphi_n(x)$ is trivial, x fixes $\mathcal{L}_1$ and so we can write $x = (x|_0, \ldots, x|_{p-1})$. Then $e = \varphi_n(x) = (\varphi_{n-1}(x|_0), \ldots, \varphi_{n-1}(x|_{p-1}))$ which implies that $\varphi_{n-1}(x|_v) = e$ for all $v \in \mathcal{L}_1$. Now $x|_v \in H$ for all $v \in \mathcal{L}_1$ (note that we are invoking (3.2) here) so $x|_v \in \ker_H(\varphi_{n-1})$ for all $v \in \mathcal{L}_1$. The map $x \mapsto (x|_0, \ldots, x|_{p-1})$ is therefore the required injection, since the tuple $(x|_0, \ldots, x|_{p-1})$ uniquely determines x.

For convenience, let us define $K = G_{[n-2]} = \varphi_{n-1}(H)$, so that H has property $\mathcal{R}_{n-1}(K)$. In order to prove the second inequality in (d), we need another inequality first; namely

$$|H| \leq \left|K_{[1]}\right|\,|K|^p\,. \tag{3.4}$$

Note that $K_{[1]} = H_{[1]} \cong H/\mathsf{st}_H(\mathcal{L}_1)$, so $|H| = |K_{[1]}|\,|\mathsf{st}_H(\mathcal{L}_1)|$. Hence it suffices to show that $|\mathsf{st}_H(\mathcal{L}_1)| \leq |K|^p$. Again, we will do this by finding an injection from $\mathsf{st}_H(\mathcal{L}_1)$ into K^p. Let $x = (x|_0, \ldots, x|_{p-1}) \in \mathsf{st}_H(\mathcal{L}_1)$. By (3.2), $x|_v \in K$ for all $v \in \mathcal{L}_1$ so once again the map $x \mapsto (x|_0, \ldots, x|_{p-1})$ is the required injection and thus (3.4) holds.

We can now prove the second inequality. Multiplying both sides of (3.4) by $|H|^p$ and observing that $H_{[1]} = K_{[1]}$, we obtain

$$|H|^{p+1} \leq \left|H_{[1]}\right| |H|^p \, |K|^p \, .$$

Now part (a) applied to H implies that $|H| = |\ker_H(\varphi_{n-1})|\,|K|$, hence

$$|\ker_H(\varphi_{n-1})|^p \, |H| = \frac{|H|^{p+1}}{|K|^p} \leq \left|H_{[1]}\right| |H|^p$$

as required. ☠

3.4. p-groups and the pro-p group $\mathsf{Aut}_p(T)$

DEFINITION 3.20. Suppose that p is prime. Let C_p denote the cyclic subgroup of $\mathsf{Aut}(T_1)$ generated by the p-cycle $(0\ 1\ \cdots\ p-1)$. Define:

$$\mathsf{Aut}_p(T) := \{g \in \mathsf{Aut}(T) \ : \ (g|_v)_{[1]} \in C_p \text{ for all } v \in T\}.$$

In other words, $\mathsf{Aut}_p(T)$ consists of the automorphisms of T whose action on the p children of each vertex (viewed as an element of $\mathsf{Aut}(T_1)$) is an element of C_p. It is a subgroup of $\mathsf{Aut}(T)$. We may also define the finite-depth version:

$$\mathsf{Aut}_p(T_n) := \{g \in \mathsf{Aut}(T_n) \ : \ (g|_v)_{[1]} \in C_p \text{ for all } v \in T_{n-1}\}.$$

Note that the definition specifies $v \in T_{n-1}$ rather than T_n so that there will be children of v for $g|_v$ to act on. This is a subgroup of $\mathsf{Aut}(T_n)$ and it has the structure of an iterated wreath product:

$$\mathsf{Aut}_p(T_n) \cong \underbrace{C_p \wr (C_p \wr (\cdots))}_{n \text{ factors}}$$

which is the same structure as $\mathsf{Aut}(T_n)$ (see (2.14)) but with $\mathsf{Sym}(p)$ replaced by C_p. Thus, using a similar argument to the one used to derive (2.17), we obtain the analogous formula:

$$|\mathsf{Aut}_p(T_n)| = p^{\frac{p^n-1}{p-1}} \, . \tag{3.5}$$

Hence $\mathsf{Aut}_p(T_n)$ is a p-group for all n. Then, following Proposition 2.30, we may view $\mathsf{Aut}_p(T)$ as the inverse limit of the $\mathsf{Aut}_p(T_n)$, making it a pro-p group.

It turns out that, given a group $G \leq \mathsf{Aut}(T)$ with property $\mathcal{R}$ or a finite group $G \leq \mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$, we can tell if G is contained in $\mathsf{Aut}_p(T)$ or $\mathsf{Aut}_p(T_n)$ simply by looking at its action on level 1 of the tree:

PROPOSITION 3.21. *Suppose that p is prime. Let $n \geq 2$ and suppose that $G \leq \mathsf{Aut}(T_n)$ has property $\mathcal{R}_n$ and that $G_{[1]} = C_p$. Then G is a subgroup of $\mathsf{Aut}_p(T_n)$ and is therefore a p-group.*

PROOF. Let n and G be as stated. Let $g \in G$ and $v \in T_{n-1}$. Since G has property $\mathcal{R}_n$, $(g|_v)_{[1]} \in G_{[1]} = C_p$. This holds for arbitrary v so we conclude that $g \in \mathsf{Aut}_p(T_n)$. Since g was also arbitrary, it follows that $G \leq \mathsf{Aut}_p(T_n)$. Subgroups of p-groups are also p-groups so G is a p-group as claimed. ☠

PROPOSITION 3.22. *Suppose that p is prime. Suppose that $G \leq \mathsf{Aut}(T)$ has property $\mathcal{R}$ and that $G_{[1]} = C_p$. Then G is a subgroup of $\mathsf{Aut}_p(T)$ and is therefore a pro-p group.*

PROOF. The proof is similar to Proposition 3.21. Let G be as stated and let $g \in G$ and $v \in T$. By self-similarity, $(g|_v)_{[1]} \in G_{[1]} = C_p$. Since this holds for arbitrary v, we conclude that $g \in \mathsf{Aut}_p(T)$. Since g was also arbitrary, it follows that $G \leq \mathsf{Aut}_p(T)$. Finally, since G is closed and a subgroup of the pro-p group $\mathsf{Aut}_p(T)$, it follows that G is itself a pro-p group. ☠

As a consequence of these results, the case $p = 2$ is special:

PROPOSITION 3.23. *If $p = 2$, then:*

- **(a)** *For all n, $\mathsf{Aut}_2(T_n) = \mathsf{Aut}(T_n)$;*
- **(b)** $\mathsf{Aut}_2(T) = \mathsf{Aut}(T)$*;*
- **(c)** *For all n, every group with property $\mathcal{R}_n$ is a 2-group;*
- **(d)** *Every group with property $\mathcal{R}$ is a pro-2 group.*

PROOF. If $p = 2$ then $\mathsf{Aut}(T_1) = C_2$. Parts (a) and (b) now follow from the definitions. Since the only transitive subgroup of C_2 is C_2 itself, part (c) now follows from Proposition 3.21 and part (d) follows from Proposition 3.22. ☠

When we are considering p-groups with property $\mathcal{R}_n$, as we will be in Chapter 5 when we focus on the case $p = 2$, some of the general properties of p-groups help to simplify our calculations. For example, suppose we fix a subgroup H of $\mathsf{Aut}(T_{n-1})$ and consider the possible groups $G \leq \mathsf{Aut}(T_n)$ with property $\mathcal{R}_n(H)$. We can order those groups by inclusion, setting up a lattice (see Figure 5.1 for example). We will now show that for any two groups in the lattice, every subgroup of $\mathsf{Aut}(T_n)$ between these subgroups also has property $\mathcal{R}_n(H)$ and therefore belongs to the lattice as well.

PROPOSITION 3.24. *Let $n \geq 2$ and let H be a subgroup of $\mathsf{Aut}(T_{n-1})$ with property $\mathcal{R}_{n-1}$. Suppose that G_1, G_2 and G_3 are subgroups of $\mathsf{Aut}(T_n)$ such that $G_1 \subseteq G_2 \subseteq G_3$, and that G_1 and G_3 both have property $\mathcal{R}_n(H)$. Then G_2 also has property $\mathcal{R}_n(H)$.*

PROOF. Let n, H, G_1, G_2 and G_3 be as stated. Then G_2 is transitive on $\mathcal{L}_1$ since G_2 contains G_1 which is transitive on $\mathcal{L}_1$. Now $H = \varphi_n(G_1) \subseteq \varphi_n(G_2) \subseteq \varphi_n(G_3) = H$ so $\varphi_n(G_2) = H$. Similarly $\psi_n(\mathsf{st}_{G_2}(0)) = H$ since $\mathsf{st}_{G_1}(0) \subseteq \mathsf{st}_{G_2}(0) \subseteq \mathsf{st}_{G_3}(0)$, which completes the proof. ☠

We can show that this 'filling in' of the subgroup lattice occurs to the maximum possible extent in the case of p-groups, in that every possible power of p occurs as the order of a group. This gives the lattice a particularly simple structure. In order to prove this, we need some facts about p-groups.

LEMMA 3.25. *The center of a p-group is always nontrivial.*

PROOF. Let G be a p-group. Recall the class equation for finite groups, obtained by enumerating the conjugacy classes:

$$|G| = |Z(G)| + \sum_i |G : C_G(x_i)|$$

where $Z(G)$ is the center of G, $C_G(x_i)$ is the centraliser of x_i in G, and the x_i are a full set of conjugacy class representatives not in $Z(G)$. Since G is a p-group, each $|G : C_G(x_i)|$ is divisible by p and hence so is $|Z(G)|$. Since the identity belongs to $Z(G)$, it follows that $|Z(G)| \geq p$ and thus $Z(G)$ is nontrivial. ☠

LEMMA 3.26. *If G is a p-group and H is a proper subgroup of G, then H is properly contained in its normaliser $N_G(H)$.*

PROOF. By induction on $|G|$. The result is immediate if $|G| = p$ since then G is cyclic of order p and the only subgroups of G are the trivial subgroup and G.

Suppose that the Lemma holds for all p-groups G with $|G| \leq p^n$ for some $n \geq 1$. Let G be a p-group such that $|G| = p^{n+1}$. Let H be a proper subgroup of G and suppose for a contradiction that $H = N_G(H)$ (note that H is always contained in $N_G(H)$). The center $Z(G)$ commutes with everything in G so it is normal in G and normalises H. Hence $Z(G) \lhd N_G(H)$, and so in our case $Z(G) \lhd H$.

Now H is a proper subgroup of G so we may view the quotient group $H/Z(G)$ as a proper subgroup of $G/Z(G)$. By Lemma 3.25, $Z(G)$ is nontrivial so $|G/Z(G)| < |G|$. Since $G/Z(G)$ is a p-group, we can apply the inductive hypothesis to conclude that $H/Z(G)$ is normalised by some element of $G/Z(G)$ which is not in $H/Z(G)$. Lifting to G, we conclude that H is normalised by some element of G not in H, which contradicts our assumption that $H = N_G(H)$. This completes the inductive step and the proof. ☠

LEMMA 3.27. *Let G be a p-group and let H be a proper subgroup of G. Then there exists a positive integer r and a chain of subnormal subgroups $H = H_0 \lhd H_1 \lhd \cdots \lhd H_r = G$ such that $|H_i : H_{i-1}| = p$ for $1 \leq i \leq r$.*

PROOF. Let G be a p-group. Note that if H is any proper subgroup of G then the index $|G : H| = p^m$ for some $m \geq 1$.

The proof is by induction on m. The base case $m = 1$ is trivial except for the normality condition. By Lemma 3.26, the normaliser $N_G(H)$ properly contains H. Since H has index p in G, it follows that $N_G(H) = G$ and so H is normal in G.

Now fix $m \geq 1$ and suppose that the claim holds for all subgroups of G with index p^m. Let H be a subgroup of G with index p^{m+1}. By Lemma 3.26, $N_G(H)$ properly contains H as a normal subgroup, so we can form the quotient group $N_G(H)/H$ which is a nontrivial p-group. By Cauchy's theorem, there exists an element of order p in $N_G(H)/H$. This amounts to saying that there exists $g \in N_G(H)$ such that $g \notin H$ and $g^p \in H$. Let $H_1 = \langle H, g\rangle$. Since g normalises H we have $H_1 = \langle g\rangle H = \bigcup_{i=0}^{p-1} g^i H$ and thus $|H_1| = p\,|H|$. Therefore H is a normal subgroup of H_1 with index p. It follows that H_1 has index p^m in G and so we may apply the inductive hypothesis to H_1 to complete the chain of subgroups from H to G as required. ☠

With these Lemmas at our disposal, we can now prove our result.

PROPOSITION 3.28. *Let $n \geq 2$ and let H be a subgroup of $\mathsf{Aut}_p(T_{n-1})$ with property $\mathcal{R}_{n-1}$. Suppose that G and G' are subgroups of $\mathsf{Aut}_p(T_n)$ with property $\mathcal{R}_n(H)$, and that $G \subseteq G'$. Then there exists a positive integer r and a chain of subgroups*

$$G = G_0 \lhd G_1 \lhd \cdots \lhd G_r = G',$$

all of which have property $\mathcal{R}_n(H)$, such that $|G_i : G_{i-1}| = p$ for $1 \leq i \leq r$.

Proof. Let n, H, G and G' be as stated. As subgroups of the p-group $\mathsf{Aut}_p(T_n)$, both G and G' are p-groups. Lemma 3.27 implies that there is a chain of subnormal subgroups from G to G' with index p at each step, and Proposition 3.24 implies that they all have property $\mathcal{R}_n(H)$.

If G is a p-group with property $\mathcal{R}_n$ then this result allows us to program a systematic and efficient search for all subgroups of G with property $\mathcal{R}_n$, by finding first all index-p subgroups of G with property $\mathcal{R}_n$ (taking advantage of efficient algorithms for finding maximal subgroups of a finite group) and then recursively looking inside those. This is the basis of the algorithm described in Appendix A. Because of Proposition 3.23, the restriction to p-groups is no restriction at all in the case $p = 2$, which was the focus of our calculations. We will see the fruits of these calculations in Chapter 5.

Before we enlist the help of a computer, though, some important examples of groups with property $\mathcal{R}$ (and $\mathcal{R}_n$) appear in other contexts and are already well-known. Chapter 4 covers these examples in detail.

CHAPTER 4

Basic examples

This brief chapter covers a couple of introductory examples of well-known totally disconnected locally compact groups and their tree representations. This is intended to make the ideas in Chapter 2 a little more concrete, and to provide a starting point for the investigations in Chapter 5.

EXAMPLE 4.1 (The p-adic numbers). Let $G = \mathbb{Q}_p$ be the additive group of p-adic numbers for some prime p. G is a totally disconnected locally compact group with the usual topology induced by the p-adic valuation. The subgroup $V = \mathbb{Z}_p$ of p-adic integers is a compact open subgroup of G. Define $\alpha : G \to G$ by $\alpha(x) = px$. Then α is a continuous automorphism of G. Now $\mathbb{Z}_p \subset p^{-1}\mathbb{Z}_p$ which means $V \subset \alpha^{-1}(V)$ and therefore $V_- = \bigcap_{n\geq 0} \alpha^{-n}(V) = V$. On the other hand, $V_+ = \bigcap_{n\geq 0} \alpha^{n}(V)$ is trivial and so we have $V = V_+V_-$. Finally V_{++} is obviously trivial as well, and

$$V_{--} = \bigcup_{n\geq 0} \alpha^{-n}(V) = \bigcup_{n\geq 0} p^{-n}\mathbb{Z}_p = \mathbb{Q}_p = G$$

is closed in G so it follows that V is tidy for α. The scale of α is 1 since $\alpha(V) \subset V$ (this fact alone implies that V is tidy for α) and

$$s(\alpha^{-1}) = \left|\alpha^{-1}(V) : V \cap \alpha^{-1}(V)\right| = \left|p^{-1}\mathbb{Z}_p : \mathbb{Z}_p\right| = p.$$

The tree representation $\pi : V_{--} \rtimes \langle\alpha\rangle \to \mathsf{Aut}(X)$ therefore acts on a homogeneous tree of valency $p+1$. We can describe this action in detail, illustrating the reduction from the homogeneous tree to the rooted tree described in Chapter 2. The vertices of X are the left cosets of V_- in $V_{--} \rtimes \langle\alpha\rangle$, which are the cosets of $\mathbb{Z}_p$ in $\mathbb{Q}_p \rtimes \langle\alpha\rangle$. Denote each coset $(x, \alpha^n)\mathbb{Z}_p$ by the pair (x, n), where $x \in \mathbb{Q}_p$ and $n \in \mathbb{Z}$. It follows from straightforward calculations that:

- **(a)** $(x, n) = (y, m)$ if and only if $n = m$ and $x - y \in p^n\mathbb{Z}_p$.
- **(b)** There is a unique in-edge to the vertex (x, n), from the vertex $(x, n-1)$.

(c) There are p out-edges from the vertex (x,n), to the vertices $(x+cp^n, n+1)$ where $c \in \{0,1,\ldots,p-1\}$.

The path $\xi = \{\xi_n\}_{n\in\mathbb{Z}}$ defined in Section 2.2 consists of the vertices $(0,n)$ where $n \in \mathbb{Z}$. Note that $\pi(\alpha)(0,n) = (0,n+1)$ for all n, so $\pi(\alpha)$ acts as a translation of amplitude 1 along ξ. In general the action of α on the tree is

$$\pi(\alpha)(x,n) = (px, n+1).$$

For $v \in V_{--} = \mathbb{Q}_p$, the action is

$$\pi(v)(x,n) = (v+x, n).$$

It follows from property (a) that $\pi(v)$ fixes the vertex (x,n) if and only if $v \in p^n\mathbb{Z}_p$. In other words, $\mathsf{st}(x,n) = \pi(p^n\mathbb{Z}_p)$ for each $n \in \mathbb{Z}$. We can use this to show that π is faithful, by showing that $\ker\pi$ is trivial. Indeed, suppose that $v \in \ker\pi$. Then $\pi(v)$ fixes every vertex in X, so

$$v \in \bigcap_{n\in\mathbb{Z}} p^n\mathbb{Z}_p = \{0\}$$

as claimed.

Consider the subtree T rooted at the vertex $\xi_0 = (0,0)$ as defined in section 2.3. Recall that the nth level of T, denoted $\mathcal{L}_n$, is the set of vertices of T which are a distance of n from the root ξ_0. Since $\mathcal{L}_1$ is just the set of children of ξ_0, property (c) tells us that $\mathcal{L}_1$ consists of the vertices $(0,1),(1,1),\ldots,(p-1,1)$. Repeatedly applying property (c) yields a nice description of $\mathcal{L}_n$ for each n:

$$\mathcal{L}_n = \left\{ \left(\sum_{i=0}^{n-1} c_i p^i, n \right) : c_i \in \{0,1,\ldots,p-1\} \right\}.$$

Observe that every vertex of T has p children, and $\mathcal{L}_n$ contains p^n vertices, indexed by the integers modulo p^n.

We know from Section 2.2 that $\mathsf{st}(\xi_0) = \pi(\mathbb{Z}_p)$ acts on T and that this action determines the whole representation. The action of $\pi(\mathbb{Z}_p)$ on T can be understood by considering its action on each level of T. For each $n \in \mathbb{N}$, the element $v \in \mathbb{Z}_p$ acts on $\mathcal{L}_n$ by

$$\pi(v)(x,n) = (v+x, n)$$

where $x = \sum_{i=0}^{n-1} c_i p^i$ as above. Now the vertices $(p^n + x, n)$ and (x,n) are identical by property (a), so $\pi(p^n)$ fixes every vertex of $\mathcal{L}_n$. In fact it follows easily from property (a) that the action of $\pi(\mathbb{Z}_p)$ on $\mathcal{L}_n$ is cyclic of order p^n, where

$$\pi(v)(x,n) = ((v+x) \bmod p^n, n).$$

In other words, when the vertices of $\mathcal{L}_n$ are identified with the integers modulo p^n as above, then $\pi(\mathbb{Z}_p)$ acts on $\mathcal{L}_n$ by addition modulo p^n. This is simply a reflection of the fact that $\mathsf{st}_{\pi(\mathbb{Z}_p)}(\mathcal{L}_n) = \pi(p^n\mathbb{Z}_p)$, and therefore $\pi(\mathbb{Z}_p)_{[n]} = \pi(\mathbb{Z}_p)/\pi(p^n\mathbb{Z}_p) \cong \mathbb{Z}/p^n\mathbb{Z}$.

This action of $\pi(\mathbb{Z}_p)$ on T is related to the *adding machine* or *odometer* [**Šun11**, Example 3], which is known to be self-replicating. It is also an automaton group; see Example 7.7. If we view $\mathbb{Z}$ as the cyclic subgroup of $\mathbb{Z}_p$ generated by 1, then we obtain the odometer as the action of $\pi(\mathbb{Z})$ on T. Indeed, the action of 1 on any vertex $v \in \mathcal{L}_n$ satisfies

$$\pi(1)(v, n) = ((v + 1) \bmod p^n, n).$$

This is the reason for the name *odometer* — it is generated by an automorphism that "counts" up by 1.

It is worth pointing out a subtlety that arises from the topology here. Both $\pi(\mathbb{Z})$ and $\pi(\mathbb{Z}_p)$ have the same action on $\mathcal{L}_n$ for all $n \in \mathbb{N}$; that is, $\pi(\mathbb{Z})_{[n]} = \pi(\mathbb{Z}_p)_{[n]}$. In other words, if we truncate the tree T at level n (for any n) then $\pi(\mathbb{Z}_p)$ is indistinguishable from the odometer. However, since $\pi(\mathbb{Z}_p)$ is closed, Proposition 2.31(b) tells us that it must be the closure of the odometer group in $\mathsf{Aut}(T)$. We can give an explicit example of a closure point: consider $v = \sum_{k=0}^{\infty} p^{2k} \in \mathbb{Z}_p$. The action of $\pi(v)$ on $\mathcal{L}_n$ is the same as the action of $\pi\left(\sum_{k=0}^{\lfloor (n-1)/2 \rfloor} p^{2k}\right) \in \pi(\mathbb{Z})$ since $\pi(p^{2k}) \in \mathsf{st}(\mathcal{L}_n)$ for $2k \geq n$. Yet $\pi(v) \notin \pi(\mathbb{Z})$ since v is not a finite sum nor the additive inverse of a finite sum. On the other hand we can see that $\pi(v)$ belongs to the closure of $\pi(\mathbb{Z})$ since on each level it agrees with some element of $\pi(\mathbb{Z})$.

Thus, the difference between the odometer and $\pi(\mathbb{Z}_p)$ only becomes clear when we look at their actions on the boundary of T. The odometer is countable, so each of its orbits on ∂T must be countable. Since ∂T is uncountable the odometer cannot be transitive on ∂T, whereas the tree representation theorem tells us that the uncountable group $\pi(\mathbb{Z}_p)$ does act transitively on ∂T. Indeed, we can find an explicit bijection between ∂T and $\mathbb{Z}_p$ so that the action of $\pi(\mathbb{Z}_p)$ on ∂T is just addition of p-adic integers. Let $\omega = (v_1, v_2, \ldots)$ be an end of T. Using the labelling described above, we may take $v_1 = (c_1, 1)$, $v_2 = (c_1 + c_2 p, 2)$, $v_3 = (c_1 + c_2 p + c_3 p^2, 3)$ and so on, where $c_i \in \{0, 1, \ldots, p-1\}$ for each i. Then we identify ω with the p-adic integer $\sum_{i=1}^{\infty} c_i p^i$. From the above description of the action of $\pi(\mathbb{Z}_p)$ on T, it follows that

$$\pi(v)(\omega) = v + \omega$$

for all $v \in \mathbb{Z}_p$. Transitivity is now obvious. This labelling is just a consequence of the fact that $\mathbb{Z}_p = \varprojlim \mathbb{Z}/p^n\mathbb{Z}$, where we are using the above identification of $\mathbb{Z}/p^n\mathbb{Z}$ with $\mathcal{L}_n$.

EXAMPLE 4.2 (Laurent series). Again let p be a prime and let $G = \mathbb{F}_p(\!(t)\!)$ be the additive group of formal Laurent series over the finite field $\mathbb{F}_p$. Elements of G have the form $f(t) = \sum_{n=N}^{\infty} a_n t^n$, where $a_i \in \mathbb{F}_p$ for each i, and $N \in \mathbb{Z}$. Addition is performed component-wise. Note the similarity to the p-adic expansion $\sum_{n=N}^{\infty} a_n p^n$; the difference is that for Laurent series there is no 'carrying' involved; addition of the separate coefficients is carried out in $\mathbb{F}_p$. This small difference results in a very different action of G on the tree, as we will see. Define the map $\alpha : G \to G$ by $\alpha(f(t)) = tf(t)$. Then α is a continuous automorphism of G, analogous to multiplication by p in the p-adic numbers.

The subgroup $V = \mathbb{F}_p[[t]]$ of G, defined by

$$\mathbb{F}_p[[t]] = \left\{ \sum_{n=0}^{\infty} a_n t^n \ : \ a_i \in \mathbb{F}_p \right\}$$

is compact and open in G. We have $\alpha(V) = t\mathbb{F}_p[[t]] \subset V$ and hence $V \subset \alpha^{-1}(V)$. Similarly to the p-adic case, we have $V_- = V$ and $V_+ = \{0\}$, which again yields $V_{--} = G$ and $V_{++} = \{0\}$, so V is tidy for α. We also have

$$s(\alpha^{-1}) = \left|\alpha^{-1}(V) : V \cap \alpha^{-1}(V)\right| = \left|t^{-1}\mathbb{F}_p[[t]] : \mathbb{F}_p[[t]]\right| = p$$

so once again the homogeneous tree X has valency $p+1$. Similarly to the p-adic case, the stabiliser of ξ_n is $\pi(t^n\mathbb{F}_p[[t]])$, so $\pi(\mathbb{F}_p[[t]])$ itself stabilises ξ_0 and acts on the rooted tree T. We also have a similar description of the vertices of T as ordered pairs of the form (x, n) where $x \in \mathbb{F}_p[[t]]$ and $n \geq 0$, where:

$$\mathcal{L}_n = \left\{ \left(\sum_{i=0}^{n-1} c_i t^i, n \right) : c_i \in \{0, 1, \dots, p-1\} \right\}.$$

The action of $\pi(\mathbb{F}_p[[t]])$ on $\mathcal{L}_n$ is similar as well. For all $v \in \mathbb{F}_p[[t]]$:

$$\pi(v)(x, n) = ((v + x) \bmod t^n, n). \tag{4.1}$$

The difference between $\mathbb{F}_p[[t]]$ and $\mathbb{Z}_p$ becomes apparent when we consider the action of $1 \in \mathbb{F}_p[[t]]$ on $\mathcal{L}_n$. In the case of $1 \in \mathbb{Z}_p$ this action was cyclic of order p^n. However, because addition in $\mathbb{F}_p[[t]]$ is term-by-term addition of power series (or polynomials) without the "carrying" that occurs in the p-adic numbers, the action of $1 \in \mathbb{F}_p[[t]]$ on $\mathcal{L}_n$ has order p rather than p^n.

In fact, a similar argument applies to the action of any element of $\mathbb{F}_p[[t]]$ — it is either trivial or has order p.

It is clear from (4.1) that the stabiliser of $\mathcal{L}_n$ is $\pi(t^n\mathbb{F}_p[[t]])$, so the action on $\mathcal{L}_n$ is isomorphic to the quotient $\mathbb{F}_p[[t]]/t^n\mathbb{F}_p[[t]]$ by Proposition 2.18. This group has order p^n, and in fact is isomorphic to the direct product $C_p \times \cdots \times C_p$ of n copies of the cyclic group of order p. To make this isomorphism explicit, the n-tuple $(a_0, \dots, a_{n-1})$ where each $a_i \in \mathbb{F}_p$ (note that the additive group of $\mathbb{F}_p$ is isomorphic to C_p) corresponds to the coset $q(t)\cdot t^n\mathbb{F}_p[[t]]$ where $q(t) = a_0 + a_1 t + \cdots + a_{n-1}t^{n-1}$. Compare this with the p-adic expansion of the integers modulo p^n from Example 4.1. It follows that the action of $\pi(\mathbb{F}_p[[t]])$ on $\mathcal{L}_n$ is generated by the set $\{\pi(1), \pi(t), \pi(t^2), \dots, \pi(t^{n-1})\}$ where each generator has order p, and $\pi(t^m)$ fixes $\mathcal{L}_m$ for each m.

Despite the initial similarity between this example and Example 4.1, they are at opposite extremes when it comes to the number of generators for their action on T — $\pi(\mathbb{Z}_p)$ is the closure of a cyclic group, and it follows from the above discussion that $\pi(\mathbb{F}_p[[t]])$ is not even topologically finitely generated.

Example 4.3 (Direct products). If G is a totally disconnected locally compact group with a continuous automorphism α, then the direct product $G^n = G \times \cdots \times G$ with the product topology is also a totally disconnected locally compact group. Define the map $\widetilde{\alpha} : G^n \to G^n$ by

$$\widetilde{\alpha}(g_1, g_2, \dots, g_n) = (g_2, \dots, g_n, \alpha(g_1)).$$

Then $\widetilde{\alpha}$ is a continuous automorphism of G^n. Furthermore, if V is tidy for α in G then $V^n = V \times \cdots \times V$ is tidy for $\widetilde{\alpha}$ in G^n, and it follows that $s(\widetilde{\alpha}) = s(\alpha)$ and $s(\widetilde{\alpha}^{-1}) = s(\alpha^{-1})$. Therefore the tree representation will act on the same homogeneous tree in both cases (up to isomorphism) although the actions will be different.

This kind of construction only works for direct products of copies of the same group — it is not so easy to generalise this to $G \times H$ where $G \neq H$, even in the case where the automorphisms have the same scale and G and H act on isomorphic trees.

CHAPTER 5

The family of groups with property $\mathcal{R}$

5.1. Introduction

In Chapter 4 we saw concrete examples of groups with property $\mathcal{R}$. Ultimately we would like to move towards a classification of these groups, but in the meantime we can look for systematic ways to generate new examples. A starting point for this is to follow Section 2.3 and look at the finite groups with property $\mathcal{R}_n$. The idea, which we develop in Sections 5.2 and 5.3, is to enlist the help of a computer to calculate all possible groups with property $\mathcal{R}_n$ for a fixed n (of which there are finitely many, since $\mathsf{Aut}(T_n)$ is finite) and then try to establish how each of those can be extended to produce groups with property $\mathcal{R}_{n+1}$. Specifically, for any subgroup G of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$, we wish to construct the set $\mathcal{E}_{n+1}(G)$ of subgroups of $\mathsf{Aut}(T_{n+1})$ with property $\mathcal{R}_{n+1}(G)$. The software package `MAGMA` has been used to calculate $\mathcal{E}_{n+1}(G)$ for every G with property $\mathcal{R}_n$ where $p = 2$ and $n \leq 4$. The results of these calculations are summarised in Appendix B. Ideally we would like to have a systematic description of $\mathcal{E}_{n+1}(G)$ for an arbitrary G but this is difficult to do in general. Section 5.3 covers some special cases where a clear pattern does emerge.

Even though we might not be able to construct the entire set $\mathcal{E}_{n+1}(G)$ in general, there are ways to construct at least one element of $\mathcal{E}_{n+1}(G)$ for an arbitrary G. Two such constructions, motivated by the results of calculations for small n, are described in detail in Chapters 6 and 7. Iterating such a construction produces a sequence of groups $(G_n)_{n=1}^{\infty}$ where each G_n has property $\mathcal{R}_n(G_{n-1})$. Then Theorem 2.32 tells us that the inverse limit $\varprojlim G_n$ has property $\mathcal{R}$.

The structure of groups with property $\mathcal{R}$ as inverse limits of groups with property $\mathcal{R}_n$ suggests a natural way to turn the family of groups with property $\mathcal{R}$ into a topological space — in fact, a metric space — by identifying it with the boundary of a rooted tree. This tree, which we shall call $\mathfrak{T}$, and the resulting metric space, which we shall call $\mathfrak{S}$, is described in Section 5.4. The points in our space $\mathfrak{S}$ are actually *conjugacy classes* of groups with property $\mathcal{R}$, for the reasons discussed in Section 3.2. We will show that $\mathfrak{S}$ is totally disconnected and compact, like $\mathsf{Aut}(T)$. However, unlike $\mathsf{Aut}(T)$, $\mathfrak{S}$ appears to have isolated points, so it is not homeomorphic to a Cantor set. Despite this additional complexity, we will prove that there exists a countable, dense subset of $\mathfrak{S}$. One such subset consists of the maximal groups that we will meet in Chapter 6.

For Sections 5.2 and 5.3, we fix $p = 2$ so that T is the rooted binary tree.

5.2. Levels 2 and 3 on the binary tree

Let us start with the simplest case $n = 1$. Here we are simply looking for subgroups G of $\mathsf{Aut}(T_1)$ which are transitive on $\mathcal{L}_1$. The reason for this is that $\mathsf{Aut}(T_0)$ is trivial so the self-similarity condition in property $\mathcal{R}_1$ is trivially satisfied. Since $\mathcal{L}_1$ contains only two vertices, $\mathsf{Aut}(T_1)$ is a group of order 2 so transitivity forces G to be equal to $\mathsf{Aut}(T_1)$. This greatly simplifies matters, since for $p > 2$ there would be more than one possibility on level 1, corresponding to all the transitive subgroups of $\mathsf{Sym}(p)$.

Now let $n = 2$. We are looking for subgroups $G \leq \mathsf{Aut}(T_2)$ with property $\mathcal{R}_2(\mathsf{Aut}(T_1))$. Since $\mathsf{Aut}(T_1)$ is transitive on $\mathcal{L}_1$ we need only check that $\psi_2(G) = \varphi_2(G) = \mathsf{Aut}(T_1)$. Obviously $G = \mathsf{Aut}(T_2)$ is one such group, so let us now assume that $G \neq \mathsf{Aut}(T_2)$. Note that $|\mathsf{Aut}(T_2)| = 2^{2^2-1} = 8$ and Proposition 3.19 tells us that $|G| \geq 2\,|\mathsf{Aut}(T_1)| = 4$ so we must have $|G| = 4$.

Let us enumerate the 8 elements of $\mathsf{Aut}(T_2)$ using the notation from Section 2.3. We will use σ to denote the transposition $(0\ 1)$ in $\mathsf{Sym}(2)$, which is the only nontrivial element of $\mathsf{Aut}(T_1)$. Thus:

$$\mathsf{Aut}(T_2) = \{e, (e,\sigma), (\sigma,e), (\sigma,\sigma), \sigma(e,e), \sigma(e,\sigma), \sigma(\sigma,e), \sigma(\sigma,\sigma)\}.$$

If G is cyclic then it must be generated by the automorphism $a = \sigma(e,\sigma)$, since a and a^{-1} are the only order 4 elements of $\mathsf{Aut}(T_2)$. This is precisely the group that arises from the action of $\mathbb{Z}_2$ restricted to T_2 which we have seen in Example 4.1; let us call this group $\mathfrak{Z}$ (this is a Fraktur 'Z' which reminds us of $\mathbb{Z}_2$). If G is not cyclic then it does not contain a or a^{-1}. The orbit-stabiliser theorem tells us that $|\mathsf{st}_G(0)| = 2$ so there are two elements of G that do not fix the vertex 0. They must therefore be $\sigma(e,e)$ and $\sigma(\sigma,\sigma)$, both of which have order 2, and thus these two elements generate G. This group is the one that arises from the action of $\mathbb{F}_2[[t]]$ on T_2 which was seen in Example 4.2; call this group $\mathfrak{L}$ (to remind us of 'Laurent series over $\mathbb{F}_2$'). In summary, there are only three potential groups with property $\mathcal{R}_2$:

- $\mathfrak{Z} := \langle \sigma(e,\sigma) \rangle = \{e, (\sigma,\sigma), \sigma(e,\sigma), \sigma(\sigma,e)\}$;
- $\mathfrak{L} := \langle \sigma(e,e), \sigma(\sigma,\sigma) \rangle = \{e, (\sigma,\sigma), \sigma(e,e), \sigma(\sigma,\sigma)\}$;
- $\mathsf{Aut}(T_2)$.

Observe that $\mathsf{st}_{\mathfrak{Z}}(0) = \mathsf{st}_{\mathfrak{L}}(0) = \{e, (\sigma,\sigma)\}$ so $\psi_2(\mathsf{st}_{\mathfrak{Z}}(0)) = \psi_2(\mathsf{st}_{\mathfrak{L}}(0)) = \mathsf{Aut}(T_1)$. Clearly $\varphi_2(\mathfrak{Z}) = \varphi_2(\mathfrak{L}) = \mathsf{Aut}(T_1)$ as well, so both groups have property $\mathcal{R}_2(\mathsf{Aut}(T_1))$. We therefore conclude:

PROPOSITION 5.1. *If $G \leq \mathsf{Aut}(T)$ has property $\mathcal{R}$ then $G_{[2]}$ is equal to either $\mathfrak{Z}$, $\mathfrak{L}$ or $\mathsf{Aut}(T_2)$.*

PROOF. Suppose that G has property $\mathcal{R}$. Then $G_{[n]}$ has property $\mathcal{R}_n$ for all n. Since G is transitive on $\mathcal{L}_1$ we must have $G_{[1]} = \mathsf{Aut}(T_1)$, so $G_{[2]}$ has property $\mathcal{R}_2(\mathsf{Aut}(T_1))$ and the result follows from the above discussion. ☠

Unfortunately this means that we are yet to discover any sign of new examples beyond those in Chapter 4. Let us move down to level 3 and consider the subgroups $G \leq \mathsf{Aut}(T_3)$ with property $\mathcal{R}_3$. We will split these groups into three cases according to the three possibilities for $G_{[2]}$. Since $|\mathsf{Aut}(T_3)| = 2^{2^3-1} = 128$, it would be tedious to enumerate all the possible groups as we did above. Instead, we have used `MAGMA`; see Appendix A for details of the algorithm.

The issue of conjugacy now arises (see Section 3.2), for we now have subgroups which are not normal in $\mathsf{Aut}(T_3)$. In total we find, up to conjugacy, 2 groups with property $\mathcal{R}_3(\mathfrak{Z})$, 4 groups with property $\mathcal{R}_3(\mathfrak{L})$, and 9 groups with property $\mathcal{R}_3(\mathsf{Aut}(T_2))$, for a total of 15 groups with property $\mathcal{R}_3$. These are enumerated in Appendix B.

In each of the three cases $G_{[2]} = \mathfrak{Z}$, $\mathfrak{L}$, $\mathsf{Aut}(T_2)$, we immediately notice that there is (for an appropriate choice of conjugacy class representatives) one group satisfying $\mathcal{R}_3(G_{[2]})$ which contains all the others. The subgroup lattices for the three families of groups are shown in Figure 5.1.

The three examples from level 2 extend to level 3: the action of $\mathbb{Z}_2$ on T_3 is the minimal subgroup with property $\mathcal{R}_3(\mathfrak{Z})$, and the action of $\mathbb{F}_2[[t]]$ on T_3 is one of the minimal subgroups with property $\mathcal{R}_3(\mathfrak{L})$. Obviously $\mathsf{Aut}(T_3)$ is the maximal subgroup with property $\mathcal{R}_3(\mathsf{Aut}(T_2))$.

Several other familiar examples can be identified among these groups. The direct product (as per Example 4.3) $\mathbb{Z}_2 \times \mathbb{Z}_2$ of two copies of the 2-adic integers acts on T_3 as one of the three minimal groups below $\mathfrak{L}$. On T_2 it is simply equal to $\mathfrak{L}$, which is isomorphic to $C_2 \times C_2$. The group generated by the actions of $\mathbb{Z}_2$ and $\mathbb{F}_2[[t]]$ on T_3 has property $\mathcal{R}_3(\mathsf{Aut}(T_2))$ and appears in

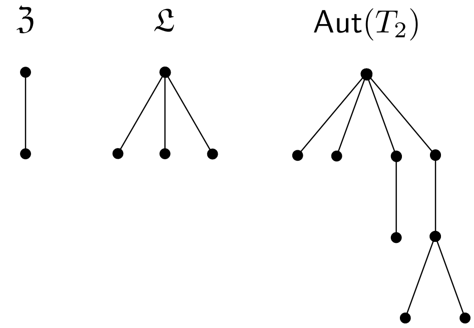

FIGURE 5.1. Subgroup lattices for groups with property $\mathcal{R}_3$

the lattice as the index 2 subgroup of $\mathsf{Aut}(T_3)$ with three descendants. Some known self-similar groups (at least, their actions on T_3) turn up in the lattice below $\mathsf{Aut}(T_2)$. The Basilica group (see [**GŻ02**] or [**DDMN10**]) appears as one of the minimal subgroups with index 2 in $\mathsf{Aut}(T_3)$, the infinite dihedral group D_∞ (see [**Šun07**, Example 1] or [**BGN03**]) appears as one of the minimal subgroups of index 8, and the lamplighter group (see [**GŻ01**] and Example 7.6) appears as the minimal subgroup of index 4. The Grigorchuk group (see Example 7.12) acts on T_3 as the full group $\mathsf{Aut}(T_3)$, but it appears in the lattice of groups with property $\mathcal{R}_4(\mathsf{Aut}(T_3))$ (see Figure 5.2) as one of the minimal subgroups of index 8 in $\mathsf{Aut}(T_4)$.

5.3. Going deeper: observed patterns, results and conjectures

The above patterns provide a taste of what happens in general as we increase the depth of the tree. In this section we will formulate some detailed conjectures about how these groups behave at arbitrary depth, and the consequences for groups acting on the infinite tree T with property $\mathcal{R}$.

Let us first set up some convenient notation.

DEFINITION 5.2. Let $G \leq \mathsf{Aut}(T_n)$ be a group with property $\mathcal{R}_n$. Define, for each $m > n$,

$$\mathcal{E}_m(G) = \{H \leq \mathsf{Aut}(T_m) \ : \ H \text{ has property } \mathcal{R}_m \text{ and } H_{[n]} = G\}.$$

Also define

$$\mathcal{E}(G) = \{\widetilde{G} \leq \mathsf{Aut}(T) \ : \ \widetilde{G} \text{ has property } \mathcal{R} \text{ and } \widetilde{G}_{[n]} = G\}.$$

We will usually consider the quotients of $\mathcal{E}_m(G)$ and $\mathcal{E}(G)$ by the equivalence relation of conjugacy. Diagrams such as those in Figure 5.1 will always depict conjugacy classes rather than individual groups; one conjugacy class will be depicted as being included in another if the inclusion holds for at least one representative of each class. Alternatively, the nodes in the diagram may be viewed as groups where one representative has been judiciously chosen from each conjugacy class.

The results shown in Figure 5.1, along with further calculations, suggest that for any G with property $\mathcal{R}_n$, the set $\mathcal{E}_{n+1}(G)$ contains a maximal group which is unique up to conjugacy. This turns out to be true (for arbitrary p, not just $p = 2$) and the construction of this maximal group is the subject of Chapter 6.

Another feature of Figure 5.1 is that in each subgroup lattice, each group is an index 2 subgroup of its immediate parent in the lattice. Since $\mathsf{Aut}(T_n)$ is a 2-group for all n, we know from Proposition 3.28 that this will always occur.

The patterns in Figure 5.1 can be generalised for $n > 3$. First, we generalise the three groups with property $\mathcal{R}_2$. Obviously the group $\mathsf{Aut}(T_2)$ generalises to $\mathsf{Aut}(T_n)$. The groups $\mathfrak{Z}$ and $\mathfrak{L}$ generalise as follows:

DEFINITION 5.3. For $n \geq 2$, define the following subgroups of $\mathsf{Aut}(T_n)$:

- $\mathfrak{Z}_n := G_{[n]}$ where G is the action of $\mathbb{Z}_2$ on T from Example 4.1. Note that $\mathfrak{Z}_2$ is just the group $\mathfrak{Z}$ we defined earlier.
- $\mathfrak{L}_n := G_{[n]}$ where G is the action of $\mathbb{F}_2[[t]]$ on T from Example 4.2. Again $\mathfrak{L}_2 = \mathfrak{L}$.

We saw in Examples 4.1 and 4.2 that $\mathfrak{Z}_n$ is cyclic of order 2^n and $\mathfrak{L}_n$ is isomorphic to a direct product of n copies of the group of order 2, so it also has order 2^n. Note that these groups are abelian. It follows from Proposition 3.19 that 2^n is the minimum possible size for a group with property $\mathcal{R}_n$. Consequently these groups, along with $\mathsf{Aut}(T_n)$, may be viewed as the extreme cases (in terms of group order) among the groups with property $\mathcal{R}_n$.

What can be said about the subgroup lattices of $\mathcal{E}_{n+1}(\mathfrak{Z}_n)$, $\mathcal{E}_{n+1}(\mathfrak{L}_n)$ and $\mathcal{E}_{n+1}(\mathsf{Aut}(T_n))$ for arbitrary n? It turns out that they resemble Figure 5.1 very closely. Figure 5.2 shows these lattices (up to conjugacy) for $n = 3$.

We immediately see that the lattice of $\mathcal{E}_4(\mathfrak{Z}_3)$ is the same as $\mathcal{E}_3(\mathfrak{Z})$, and $\mathcal{E}_4(\mathfrak{L}_3)$ differs from $\mathcal{E}_3(\mathfrak{L})$ only in that there is one more group. Further calculations reveal that this pattern continues for larger n. The lattice of $\mathcal{E}_{n+1}(\mathfrak{Z}_n)$ is always the same; it consists of a maximal group and an index 2 subgroup which is just $\mathfrak{Z}_{n+1}$. The lattice of $\mathcal{E}_{n+1}(\mathfrak{L}_n)$ always consists of a maximal group and $n + 1$ subgroups of index 2, one of which is $\mathfrak{L}_{n+1}$.

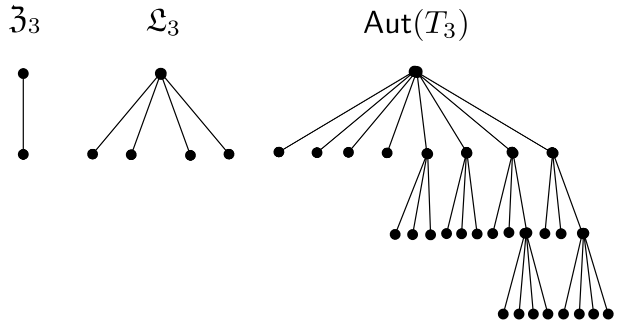

FIGURE 5.2. Subgroup lattices for groups with property $\mathcal{R}_4$ below $\mathfrak{Z}_3$, $\mathfrak{L}_3$ and $\mathsf{Aut}(T_3)$.

The picture below $\mathsf{Aut}(T_n)$ is much more complicated, but it is already possible to see similarities between the cases $n = 2$ and $n = 3$. Based on the patterns seen so far, we make the following conjecture. Further calculations confirm this result up to at least $n = 6$.

CONJECTURE 5.4. *For $n \geq 2$, the subgroup lattice (up to conjugacy) of $\mathcal{E}_{n+1}(\mathsf{Aut}(T_n))$ has the following structure:*

Layer 1: $\mathsf{Aut}(T_{n+1})$*, which has order* $2^{2^{n+1}-1}$*;*

Layer 2: 2^n *subgroups of index* 2 *in* $\mathsf{Aut}(T_{n+1})$*;*

Layer 3: $2^{n-1}(2^{n-1}-1)$ *subgroups of index* 4 *in* $\mathsf{Aut}(T_{n+1})$ *— half of the groups in Layer 2 each contain* $2^{n-1}-1$ *of these subgroups;*

Layer 4: 2^{2n-3} *subgroups of index* 8 *in* $\mathsf{Aut}(T_{n+1})$ *—* 2^{n-2} *of the groups in Layer 3, each of which is contained in a different Layer 2 group, each contain* 2^{n-1} *of these subgroups.*

This yields a total of $1 + 2^{n-1}(1 + 3 \cdot 2^{n-2})$ *conjugacy classes of groups with property* $\mathcal{R}_{n+1}(\mathsf{Aut}(T_n))$*.*

Another significant feature of these subgroup lattices involves the number of generators of each group. Let us use the notation $\sharp G$ to denote the minimum size of a generating set of a group G. Based on all the lattices that have been calculated, the following result appears to hold in general:

CONJECTURE 5.5. *Suppose that $G \leq \mathsf{Aut}(T_n)$ has property $\mathcal{R}_n(H)$ for some $H \leq \mathsf{Aut}(T_{n-1})$ and that $\sharp G = \sharp H$. Then G is a minimal element of $\mathcal{E}_n(H)$; that is, no proper subgroups of G have property $\mathcal{R}_n(H)$.*

One may naturally ask about the converse of Conjecture 5.5: if G is minimal in $\mathcal{E}_n(H)$ then does it follow that $\sharp G = \sharp H$? A counterexample shows the answer to be no in general: when $H = \mathfrak{L}_{n-1}$ we have $\sharp H =$

$n-1$. The group $G = \mathfrak{L}_n$ belongs to $\mathcal{E}_n(H)$ and is minimal (since it has the minimum possible order, 2^n) but $\sharp G = n$. However, this is the *only* counterexample among all the lattices that have been calculated. Thus we claim:

CONJECTURE 5.6. *Suppose that $G \leq \mathsf{Aut}(T_n)$ has property $\mathcal{R}_n(H)$ for some $H \leq \mathsf{Aut}(T_{n-1})$, where H is not conjugate to $\mathfrak{L}_{n-1}$, and suppose that G is minimal in $\mathcal{E}_n(H)$. Then $\sharp G = \sharp H$.*

Calculations have shown that if a counterexample to either Conjecture 5.5 or Conjecture 5.6 exists then n cannot be less than 6.

5.4. The topological space $\mathfrak{S}$ of groups with property $\mathcal{R}$

For this section, we relax our condition that $p = 2$ and let p be a fixed but arbitrary integer ≥ 2. For each $n \geq 1$ let $\mathfrak{R}_n$ be the set of all groups with property $\mathcal{R}_n$. Define

$$\mathfrak{R} := \bigcup_{n=1}^{\infty} \mathfrak{R}_n.$$

$\mathfrak{R}_n$ is finite for each n because $\mathsf{Aut}(T_n)$ is finite, hence $\mathfrak{R}$ is countable. Define a relation $\sim$ on $\mathfrak{R}$, where $G_1 \sim G_2$ iff G_1 and G_2 belong to the same $\mathfrak{R}_n$ and they are conjugate in $\mathsf{Aut}(T_n)$. It is clear that $\sim$ is an equivalence relation. Since $\sim$ can be restricted to each $\mathfrak{R}_n$, the set of equivalence classes $\mathfrak{R}/{\sim}$ is just the union $\bigcup_{n=1}^{\infty}(\mathfrak{R}_n/{\sim})$. We will use the notation $[G]$ for the equivalence class of G with respect to $\sim$. In other words, if G is a group in $\mathfrak{R}_n$ then $[G]$ is the set of all groups with property $\mathcal{R}_n$ that are conjugate to G in $\mathsf{Aut}(T_n)$. Note that $[G]$ is not the full conjugacy class of G in $\mathsf{Aut}(T_n)$ because it excludes groups that do not have property $\mathcal{R}_n$.

We will use the maps φ_n to endow $\mathfrak{R}/{\sim}$ with the structure of a rooted tree, which we will call $\mathfrak{T}$, defined as follows. For each $n \geq 1$, define the nth level of $\mathfrak{T}$ to be the set $\mathfrak{R}_n/{\sim}$. That is, the vertices on level n are the conjugacy classes of groups with property $\mathcal{R}_n$. Also define the root of $\mathfrak{T}$ to be the trivial group $\mathsf{Aut}(T_0)$ (recall that T_0 is simply the root of T). Define the edge relation on $\mathfrak{T}$ by defining the parent of a vertex $[G]$ on level n of $\mathfrak{T}$ to be the vertex $[\varphi_n(G)]$ on level $n-1$.

Let us check that the parent of each vertex is well-defined; in other words, that it does not depend on the choice of conjugacy class representative. Suppose that G and G' both have property $\mathcal{R}_n$ and that $[G] = [G']$. Then there exists $x \in \mathsf{Aut}(T_n)$ such that $x^{-1}Gx = G'$. Therefore $\varphi_n(G)$

and $\varphi_n(G')$ have property $\mathcal{R}_{n-1}$ and $\varphi_n(x)^{-1}\varphi_n(G)\varphi_n(x) = \varphi_n(G')$, hence $[\varphi_n(G)] = [\varphi_n(G')]$ as desired.

Note that each vertex on level 1 of $\mathfrak{T}$ is connected to the root, since $\varphi_1(G)$ is trivial for all subgroups G of $\mathsf{Aut}(T_1)$. Each vertex on level n for $n > 1$ is connected to its parent on level $n-1$, so it follows that every vertex of $\mathfrak{T}$ is connected to the root and thus $\mathfrak{T}$ is connected. It contains no cycles since each vertex on level n for $n \geq 1$ is adjacent only to vertices on level $n-1$ and level $n+1$. Thus, to form a cycle, there would have to exist a vertex on level n for some n that is adjacent to two distinct vertices on level $n-1$. This is impossible since the parent of each vertex is uniquely defined. Therefore $\mathfrak{T}$ is indeed a tree. It is locally finite since $\mathfrak{R}_n$ is finite for each n.

The tree $\mathfrak{T}$ for $p = 2$ has been calculated explicitly using `MAGMA` down to level 6 for most branches; the main obstacle to further calculations is the inefficiency of testing conjugacy in $\mathsf{Aut}(T_n)$. We have already seen in Section 5.2 the 3 vertices (i.e. conjugacy classes of groups) on level 2 of $\mathfrak{T}$ and the 15 vertices on level 3. There are 118 vertices on level 4 and 2207 vertices on level 5. See Appendix B for more details of these groups.

The idea now is to identify the ends of $\mathfrak{T}$ with the conjugacy classes of groups with property $\mathcal{R}$. Recall that, since $\mathfrak{T}$ is a rooted tree, the ends of $\mathfrak{T}$ are in one-to-one correspondence with the singly infinite paths in $\mathfrak{T}$ descending from the root. Thus we may consider $\partial\mathfrak{T}$ to be the set of these paths, and it will be convenient to do so without further comment from now on.

We need some more notation. Let $\mathfrak{R}_\infty$ be the set of subgroups of $\mathsf{Aut}(T)$ with property $\mathcal{R}$ and let $\sim$ be the conjugacy relation on $\mathfrak{R}_\infty$. As above, the equivalence class of $G \in \mathfrak{R}_\infty$ with respect to $\sim$ will be denoted $[G]$.

Definition 5.7. Define $\mathfrak{S}$ to be the quotient $\mathfrak{R}_\infty/{\sim}$. In other words, $\mathfrak{S}$ is the set of conjugacy classes of groups with property $\mathcal{R}$.

Proposition 5.8. *There is a bijection between $\mathfrak{S}$ and $\partial\mathfrak{T}$.*

Proof. Let $G \in \mathfrak{R}_\infty$. For each $n \geq 1$ the restriction $G_{[n]}$ has property $\mathcal{R}_n(G_{[n-1]})$, so the sequence $\left([G_{[n]}]\right)_{n=0}^{\infty}$ is a path descending from the root in $\mathfrak{T}$. Thus we may define the map

$$\begin{aligned} \theta : \mathfrak{S} &\longrightarrow \partial\mathfrak{T} \\ [G] &\longmapsto \left([G_{[n]}]\right)_{n=0}^{\infty}. \end{aligned}$$

We claim that θ is the required bijection. First we must check that θ does not depend on the choice of conjugacy class representatives. Suppose that G and G' have property $\mathcal{R}$ and $[G] = [G']$. Then there exists $x \in \mathsf{Aut}(T)$ such that $x^{-1}Gx = G'$, which implies that $x_{[n]}^{-1}G_{[n]}x_{[n]} = G'_{[n]}$ for all n. Thus $[G_{[n]}] = [G'_{[n]}]$ for all n and θ is well-defined.

To show that θ is injective, suppose that $\theta([G]) = \theta([G'])$ for some G and G' in $\mathfrak{R}_\infty$. Then $[G_{[n]}] = [G'_{[n]}]$ for all n, hence $G_{[n]}$ is conjugate to $G'_{[n]}$ for all n. It follows from Proposition 3.6 that G is conjugate to G', so $[G] = [G']$ and θ is injective.

To show that θ is surjective, let $\omega = ([G_n])_{n=0}^{\infty}$ be a path descending from the root in $\mathfrak{T}$. By definition of $\mathfrak{T}$, we may assume without loss of generality that $\varphi_n(G_n) = G_{n-1}$ for all $n \geq 1$. This means that G_n has property $\mathcal{R}_n(G_{n-1})$ for all $n \geq 1$. Now we invoke Theorem 2.32: let $G = \varprojlim G_n$. Then G has property $\mathcal{R}$ and $G_{[n]} = G_n$ for all n, so $\theta([G]) = ([G_n])_{n=0}^{\infty} = \omega$. Since ω was arbitrary, this proves that θ is surjective and we are done. ☠

In summary, the map θ simply identifies — up to conjugacy — an end of $\mathfrak{T}$ with the inverse limit of a sequence of groups representing the vertices along the corresponding path.

There is a natural way to define a metric on $\partial\mathfrak{T}$. The above bijection will then induce a corresponding metric on $\mathfrak{S}$, turning it into a compact, totally disconnected metric space. In order to define this metric we need to set up some notation. Let $\omega = (\omega_n)_{n=0}^{\infty}$ and $\varpi = (\varpi_n)_{n=0}^{\infty}$ be ends of $\mathfrak{T}$. These two paths both begin at the root of $\mathfrak{T}$ so $\omega_0 = \varpi_0$. If $\omega \neq \varpi$ then the two paths diverge at some point and do not reconnect because $\mathfrak{T}$ is a tree. Define

$$\ell(\omega, \varpi) = \sup\{n : \omega_n = \varpi_n\}$$

which is finite if $\omega \neq \varpi$.

Proposition 5.9. *Define the function $d : \partial\mathfrak{T} \times \partial\mathfrak{T} \longrightarrow \mathbb{R}$ as follows:*

$$d(\omega, \varpi) = \begin{cases} 2^{-\ell(\omega,\varpi)} & \text{if } \omega \neq \varpi \\ 0 & \text{if } \omega = \varpi. \end{cases}$$

Then d is a metric (in fact, an ultrametric) on $\partial\mathfrak{T}$.

Proof. Clearly $d(\omega, \varpi) \geq 0$ for all $\omega, \varpi \in \partial\mathfrak{T}$ and $d(\omega, \varpi) = 0$ if and only if $\omega = \varpi$. Symmetry of d follows from symmetry of ℓ. It remains to prove the ultrametric inequality, namely $d(\omega, \varpi) \leq \max\{d(\omega, \eta), d(\varpi, \eta)\}$

for all $\omega, \varpi, \eta \in \partial\mathfrak{T}$. This is trivial if any two of the three ends are equal, so suppose they are all distinct. Then the inequality is equivalent to

$$\ell(\omega, \varpi) \geq \min\{\ell(\omega, \eta), \ell(\varpi, \eta)\}. \tag{5.1}$$

If $\ell(\omega, \varpi) \geq \ell(\omega, \eta)$ then we are done. Suppose that $\ell(\omega, \varpi) < \ell(\omega, \eta)$. Let $(\omega_n)_{n=0}^{\infty}$, $(\varpi_n)_{n=0}^{\infty}$, and $(\eta_n)_{n=0}^{\infty}$ be the corresponding paths in $\mathfrak{T}$. For convenience, let $n = \ell(\omega, \varpi)$. Then $\omega_n = \varpi_n$. In fact, we have $\varpi_n = \omega_n = \eta_n$ since $n < \ell(\omega, \eta)$. This inequality is strict, so $n + 1 \leq \ell(\omega, \eta)$. Therefore $\omega_{n+1} = \eta_{n+1}$, but $\omega_{n+1} \neq \varpi_{n+1}$ since $\ell(\omega, \varpi) = n$, which means $\varpi_{n+1} \neq \eta_{n+1}$. Putting this together with the fact that $\varpi_n = \eta_n$ we conclude that $\ell(\varpi, \eta) = n$. This yields equality in (5.1) which suffices to complete the proof. ☠

PROPOSITION 5.10. *The metric space $(\partial\mathfrak{T}, d)$ is compact and totally disconnected.*

PROOF. First we will make the proof easier by giving an alternative characterisation of the topology on $(\partial\mathfrak{T}, d)$.

For each $n \geq 0$, denote level n of $\mathfrak{T}$ by $\mathfrak{T}_{[n]}$. Then for each $n \geq 1$ there is a surjective map $\chi_n : \mathfrak{T}_{[n]} \to \mathfrak{T}_{[n-1]}$ which maps each vertex to its parent. Putting the discrete topology on each level, we can think of $\partial\mathfrak{T}$ as the inverse limit of the topological spaces $\mathfrak{T}_{[n]}$ with respect to the maps χ_n. Denote the resulting topology by τ; it is the coarsest topology such that each of the projection maps $\widetilde{\chi}_n : \partial\mathfrak{T} \to \mathfrak{T}_{[n]}$ is continuous. To make this explicit, let $\omega = (\omega_n)_{n=0}^{\infty} \in \partial\mathfrak{T}$. Then $\widetilde{\chi}_n(\omega) = \omega_n$ for each n. A base for the topology τ consists of sets of the form $\widetilde{\chi}_n^{-1}(\mathfrak{U})$ where $\mathfrak{U}$ is any subset of $\mathfrak{T}_{[n]}$. Note that $\widetilde{\chi}_n^{-1}(\mathfrak{U})$ is just the set of all paths descending from the root which pass through a vertex in $\mathfrak{U}$. By [**RZ10**, Theorem 1.1.12], the topological space $(\partial\mathfrak{T}, \tau)$ is compact and totally disconnected.

We claim that τ is precisely the topology τ_d induced by the metric d. Consider an open ball $B(\omega, r)$ in $(\partial\mathfrak{T}, d)$, where $\omega \in \partial\mathfrak{T}$ and $r > 0$. Since d only takes values of the form 2^{-n} or 0, it suffices to consider $r = 2^{-n}$ for some nonnegative integer n. It is easy to verify that

$$B(\omega, 2^{-n}) = \{(\varpi_n)_{n=0}^{\infty} \in \partial\mathfrak{T} \ : \ \varpi_{n+1} = \omega_{n+1}\} \tag{5.2}$$

where the $n+1$ is to ensure that the inequality $d(\omega, \varpi) < 2^{-n}$ is strict. This set is precisely $\widetilde{\chi}_{n+1}^{-1}(\omega_{n+1}) \in \tau$. Since every open set in $(\partial\mathfrak{T}, d)$ is a union of open balls, we conclude that $\tau_d \subseteq \tau$.

Conversely, consider the open set $\widetilde{\chi}_n^{-1}(\mathfrak{U}) \in \tau$ where $\mathfrak{U}$ is a subset of $\mathfrak{T}_{[n]}$. Now $\widetilde{\chi}_n^{-1}(\mathfrak{U}) = \bigcup_{u \in \mathfrak{U}} \widetilde{\chi}_n^{-1}(u)$ and (5.2) implies that each $\widetilde{\chi}_n^{-1}(u)$ is equal to $B(\omega, 2^{-(n-1)})$ for any $\omega \in \chi_n^{-1}(u)$. Thus, as a union of open balls, $\widetilde{\chi}_n^{-1}(\mathfrak{U}) \in \tau_d$. Every open set in τ is a union of these $\widetilde{\chi}_n^{-1}(\mathfrak{U})$ so we conclude that $\tau \subseteq \tau_d$. Hence $\tau = \tau_d$ as claimed. ☠

Since the map θ defined in the proof of Proposition 5.8 is a bijection between $\mathfrak{S}$ and $\partial\mathfrak{T}$, we may use the existing metric on $\partial\mathfrak{T}$ to induce a metric on $\mathfrak{S}$. Define the function $\delta : \mathfrak{S} \times \mathfrak{S} \longrightarrow \mathbb{R}$ by:

$$\delta\left([G], [G']\right) = d\left(\theta\left([G]\right), \theta\left([G']\right)\right). \tag{5.3}$$

In other words, we define the metric on $\mathfrak{S}$ so as to make θ a bijective isometry from $\mathfrak{S}$ to $\partial\mathfrak{T}$. Then Propositions 5.9 and 5.10 immediately carry over, via the isometry θ^{-1}, to $\mathfrak{S}$:

Corollary 5.11. *$(\mathfrak{S}, \delta)$ is a compact, totally disconnected metric space.*

Following the proof of Proposition 5.10, the open balls in $(\mathfrak{S}, \delta)$ are:

$$B([G], 2^{-n}) = \left\{[G'] \in \mathfrak{S} \ : \ [G'_{[n+1]}] = [G_{[n+1]}]\right\}.$$

In other words, an open ball in $\mathfrak{S}$ is the set of all conjugacy classes of groups with property $\mathcal{R}$ which agree (up to conjugacy) on T_n for some n. Note that the set $\mathcal{E}(G)$ defined at the start of Section 5.3, where G is a group with property $\mathcal{R}_n$ for some n, is always an open ball in $\mathfrak{S}$.

Recall that a topological space is said to be *perfect* if it contains no isolated points; equivalently, every point is a limit point. That is, every neighbourhood of any point x in the space contains a point distinct from x. It is known that every compact, totally disconnected, perfect metric space is homeomorphic to the Cantor set [**HY61**, Corollary 2-98]. We should therefore check to see if $\mathfrak{S}$ is perfect or not. Surprisingly, perhaps, the answer appears to be no:

Conjecture 5.12. *$\mathfrak{S}$ contains isolated points.*

The justification for this conjecture will be given shortly, after we have examined the topologically finitely generated groups in $\mathfrak{S}$.

It would have been convenient to invoke [**Mun00**, Theorem 27.7] to prove that $\mathfrak{S}$ is uncountable. Unfortunately this does not work, due to the (probable) failure of $\mathfrak{S}$ to be perfect. However, we can prove that $\mathfrak{S}$ does have a countable, dense subset. Recall that such spaces are called *separable*.

PROPOSITION 5.13. *$\mathfrak{S}$ is separable.*

PROOF. Because of (5.3), it suffices to prove that $\partial\mathfrak{T}$ is separable. First we need to establish the following result:

LEMMA 5.14. *Suppose that there exists a function $\mathfrak{f} : \mathfrak{T} \longrightarrow \partial\mathfrak{T}$ such that for each vertex $\mathfrak{v} \in \mathfrak{T}$, the path descending from the root that corresponds to $\mathfrak{f}(\mathfrak{v})$ passes through $\mathfrak{v}$. Then the image of $\mathfrak{f}$ is countable and dense in $\partial\mathfrak{T}$.*

PROOF. Let $\mathfrak{f}$ be such a function. The tree $\mathfrak{T}$ has countably many vertices since $\mathfrak{R}$ is countable. Thus the image of $\mathfrak{f}$ is also countable.

Now fix $\omega \in \partial\mathfrak{T}$ and let $(\omega_n)_{n=0}^{\infty}$ be the corresponding path descending from the root. We claim that $\omega = \lim_{n\to\infty} \mathfrak{f}(\omega_n)$. Indeed, for each n, our assumption implies that the path corresponding to $\mathfrak{f}(\omega_n)$ passes through ω_n. Therefore $d(\omega, \mathfrak{f}(\omega_n)) \leq 2^{-n}$ for all n. These distances converge to 0 as $n \to \infty$ and the claim follows. Thus every $\omega \in \partial\mathfrak{T}$ is a limit point of the image of $\mathfrak{f}$, which means that the image of $\mathfrak{f}$ is dense in $\partial\mathfrak{T}$ as required. ☠

To prove Proposition 5.13, we will exhibit a function $\mathfrak{f}$ that satisfies the hypotheses of Lemma 5.14. Fix $n \geq 1$ and let $\mathfrak{v}$ be a vertex on level n of $\mathfrak{T}$. Then $\mathfrak{v}$ is a conjugacy class of groups with property $\mathcal{R}_n$. Define $\mathfrak{f}(\mathfrak{v})$ as follows. Let $\widehat{G}$ be any subgroup of $\mathsf{Aut}(T)$ with property $\mathcal{R}$ such that $\widehat{G}_{[n]} = G$ for some $G \in \mathfrak{v}$. Such a $\widehat{G}$ always exists; for example, following Section 6.3, we may define $\widehat{G} = \mathcal{M}^{\infty}(G)$ where G is any group in $\mathfrak{v}$ that has sufficient rigid automorphisms (and the existence of such a G follows from Proposition 3.10). Now $[\widehat{G}] \in \mathfrak{S}$, so let us define

$$\mathfrak{f}(\mathfrak{v}) = \theta([\widehat{G}])$$

which belongs to $\partial\mathfrak{T}$. It follows from the definition of θ that the path corresponding to $\mathfrak{f}(\mathfrak{v})$ passes through $\mathfrak{v}$. Doing this for all $\mathfrak{v} \in \mathfrak{T}$ yields a function $\mathfrak{f}$ with the required properties, and applying Lemma 5.14 to $\mathfrak{f}$ completes the proof. ☠

Let us return to the question of isolated points in $\mathfrak{S}$. In order for a point $[G] \in \mathfrak{S}$ to be isolated, there must exist $r > 0$ such that $B([G], r) = \{[G]\}$. Because of the definition of the metric on $\mathfrak{S}$, it suffices to consider $r = 2^{-n}$ for integers $n \geq 0$.

Fix an integer $n \geq 0$ and fix a group G with property $\mathcal{R}$. There is an easier way to visualise the open ball $B([G], 2^{-n})$ in $\mathfrak{S}$: it consists of the conjugacy classes whose members all agree with $G_{[n+1]}$ (up to conjugacy) on $\mathsf{Aut}(T_{n+1})$.

PROPOSITION 5.15. *Let $n \geq 0$ and suppose that G has property $\mathcal{R}$. Then*

$$B([G], 2^{-n}) = \{[H] \in \mathfrak{G} \ : \ H_{[n+1]} \sim G_{[n+1]}\}.$$

PROOF. Follows from (5.2) and the fact that θ is an isometry between $\mathfrak{G}$ and $\partial\mathfrak{T}$. ☠

We immediately derive the following characterisation of isolated points in $\mathfrak{G}$:

COROLLARY 5.16. *The following are equivalent:*

(a) $[G]$ *is isolated in* $\mathfrak{G}$*;*
(b) *There exists $n \geq 1$ such that if $[H] \in \mathfrak{G}$ and $H_{[n]} \sim G_{[n]}$ then* $[H] = [G]$.

In other words, $[G]$ is isolated in $\mathfrak{G}$ if and only if G is the *only* group (up to conjugacy) with property $\mathcal{R}$ that extends $G_{[n]}$ to T. Using the notation from earlier, this is equivalent to saying that every group in $\mathcal{E}(G_{[n]})$ is conjugate to G.

Recall Conjecture 5.5, which says that if G has property $\mathcal{R}_n(H)$ and $\sharp G = \sharp H$ then G is minimal (with respect to inclusion) among all groups with property $\mathcal{R}_n(H)$. This has consequences for topologically finitely generated subgroups of $\mathsf{Aut}(T)$ with property $\mathcal{R}$. We say *topologically* finitely generated (that is, the closure of a finitely generated group) because it is impossible for a finitely generated subgroup of $\mathsf{Aut}(T)$ to have property $\mathcal{R}$. This is a consequence of the fact that all finitely generated groups are countable, but all groups with property $\mathcal{R}$ are uncountable since they are transitive on the uncountable set ∂T.

PROPOSITION 5.17. *Suppose that Conjecture 5.5 holds, and suppose that G is a topologically finitely generated group with property $\mathcal{R}$. Then there exists $n \geq 1$ such that G is a minimal element of $\mathcal{E}(G_{[n]})$.*

PROOF. If G is topologically finitely generated then $G = \overline{H}$ for some finitely generated group H such that $H_{[n]} = G_{[n]}$ for all n. Since H is finitely generated, there exists N such that $\sharp H_{[n]} = \sharp H$ for all $n \geq N$, and hence $\sharp G_{[n+1]} = \sharp G_{[n]}$ for all $n \geq N$. Then Conjecture 5.5 implies that $G_{[n+1]}$ is minimal in $\mathcal{E}_{n+1}(G_{[n]})$ for all $n \geq N$. It follows that G is minimal in $\mathcal{E}(G_{[N]})$ as required. ☠

We are now in a position to justify Conjecture 5.12. Again we need to use a result from Chapter 6 regarding the maximal group $\mathcal{M}^{\infty}(G)$, which we also invoked in the proof of Proposition 5.13.

PROPOSITION 5.18. *Suppose that Conjecture 5.5 holds. Suppose that G is a topologically finitely generated group with property $\mathcal{R}$, and that G is a maximal element of $\mathcal{E}(G_{[n]})$ for some $n \geq 1$. Then $[G]$ is isolated in $\mathfrak{S}$.*

PROOF. Let G be as stated. By Proposition 5.17, there exists $n \geq 1$ such that G is minimal in $\mathcal{E}(G_{[n]})$. Since G is also maximal in $\mathcal{E}(G_{[n]})$, it follows from Theorem 6.3(e) that G is conjugate to the maximal group $\mathcal{M}^{\infty}(G_{[n]})$ and furthermore (since G is minimal) that every group in $\mathcal{E}(G_{[n]})$ is conjugate to G. Then $[G]$ is isolated in $\mathfrak{S}$ by Corollary 5.16. ☠

Thus, Conjecture 5.12 follows from Conjecture 5.5 provided that we can find a maximal group G that is topologically finitely generated. Indeed, one such example is the Grigorchuk group which is discussed later in Example 7.12. Such groups appear to be very rare, based on calculations so far. This makes intuitive sense because there is tension between the two conditions on such a group; being maximal makes G "large" and being topologically finitely generated tends to make G "small". Finding more groups with these properties — possibly along the lines of the "Grigorchuk-type" spinal groups in [**Šun07**] — would be an interesting direction for future work.

CHAPTER 6

Maximal groups

6.1. Introduction

This chapter introduces a construction — anticipated in Chapter 5 — which takes a group G with property $\mathcal{R}_n$ and produces a group, which we will call $\mathcal{M}(G)$, with property $\mathcal{R}_{n+1}(G)$. This group $\mathcal{M}(G)$ turns out to be maximal among groups with property $\mathcal{R}_{n+1}(G)$, and is in fact a maxi*mum* up to conjugacy. In Section 6.3 we iterate this construction and invoke Theorem 2.32 to produce a group $\mathcal{M}^\infty(G)$ with property $\mathcal{R}$ which agrees with G on T_n. The maximality properties of $\mathcal{M}(G)$ carry through to $\mathcal{M}^\infty(G)$. In Section 6.4 we prove that each group $\mathcal{M}^\infty(G)$ is a *finitely constrained group* defined by a set of forbidden patterns of size n, namely the complement of G in $\mathsf{Aut}(T_n)$.

Two purposes are served by the construction of $\mathcal{M}(G)$. Firstly, when searching for groups with property $\mathcal{R}_{n+1}(G)$, it means we can restrict our attention to subgroups of $\mathcal{M}(G)$ rather than the (often) much larger group $\mathsf{Aut}(T_{n+1})$. Secondly, and more fundamentally, it guarantees that there always exists at least one group (up to conjugacy) with property $\mathcal{R}_{n+1}(G)$. Then, by iterating the construction to produce the group $\mathcal{M}^\infty(G)$, we guarantee that every group with property $\mathcal{R}_n$ can be extended (up to conjugacy) to a group with property $\mathcal{R}$.

Recall that we already used this fact to prove Proposition 5.13, where we found a countable dense subset of the family $\mathfrak{S}$ of conjugacy classes of groups with property $\mathcal{R}$. This means that every group with property $\mathcal{R}$ can be approximated, in the sense of the metric we defined on $\mathfrak{S}$, by groups of the form $\mathcal{M}^\infty(G)$.

6.2. The finite maximal group $\mathcal{M}(G)$

Let G be a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$ for some $n \geq 2$. We will need to use the restriction of G to T_{n-1} in the construction, so let us be more specific and say that G has property $\mathcal{R}_n(H)$ where $H \leq \mathsf{Aut}(T_{n-1})$.

Assume further (without loss of generality, following Proposition 3.10) that G contains sufficient rigid automorphisms.

DEFINITION 6.1. Let G be as above. Define:

$$\mathcal{M}(G) := \{x \in \mathsf{Aut}(T_{n+1}) : \varphi_{n+1}(x) \in G \text{ and } x|_v \in G \text{ for all } v \in \mathcal{L}_1\}.$$

The following theorem proves that $\mathcal{M}(G)$ has the desired properties.

THEOREM 6.2. *Define $\mathcal{M}(G)$ as above. Then:*

- **(a)** *$\mathcal{M}(G)$ has property $\mathcal{R}_{n+1}(G)$;*
- **(b)** *$\mathcal{M}(G)$ has sufficient rigid automorphisms;*
- **(c)** *$|\mathcal{M}(G)| = |G|\,|\ker_G(\varphi_n)|^p = |G|^{p+1}/|H|^p$;*
- **(d)** *If $G' \leq G$ then $\mathcal{M}(G') \leq \mathcal{M}(G)$, with equality if and only if $G' = G$. In particular, $\mathcal{M}(G) = \mathsf{Aut}(T_{n+1})$ if and only if $G = \mathsf{Aut}(T_n)$;*
- **(e)** *Every subgroup of $\mathsf{Aut}(T_{n+1})$ with property $\mathcal{R}_{n+1}(G)$ is conjugate to a subgroup of $\mathcal{M}(G)$.*

PROOF. **(a)** First we must show that $\mathcal{M}(G)$ is a group. It is nonempty since it contains the identity. Suppose that $x, y \in \mathcal{M}(G)$. We must show that $xy^{-1} \in \mathcal{M}(G)$. We have $\varphi_{n+1}(xy^{-1}) = \varphi_{n+1}(x)\varphi_{n+1}(y)^{-1} \in G$ since $\varphi_{n+1}(x) \in G$ and $\varphi_{n+1}(y) \in G$ and G is a group. Now fix $v \in \mathcal{L}_1$. Then $(xy^{-1})|_v = x|_{y^{-1}(v)}(y|_{y^{-1}(v)})^{-1}$, which is in G since $x|_{y^{-1}(v)}$ and $y|_{y^{-1}(v)}$ are in G. Hence $xy^{-1} \in \mathcal{M}(G)$ and so $\mathcal{M}(G)$ is a group.

Now let us show that $\varphi_{n+1}(\mathcal{M}(G)) = \psi_{n+1}(\mathsf{st}_{\mathcal{M}(G)}(0)) = G$. The definition of $\mathcal{M}(G)$ immediately implies that $\varphi_{n+1}(\mathcal{M}(G)) \subseteq G$. It also follows that $\psi_{n+1}(\mathsf{st}_{\mathcal{M}(G)}(0)) \subseteq G$ since $\psi_{n+1}(g) = g|_0 \in G$ for all $g \in \mathsf{st}_{\mathcal{M}(G)}(0)$. It remains to show the reverse inclusions.

To show that $\varphi_{n+1}(\mathcal{M}(G)) \supseteq G$, fix $g \in G$. We must find $x \in \mathcal{M}(G)$ such that $\varphi_{n+1}(x) = g$. By the hypotheses on G and Proposition 3.11, $g|_v \in H$ for all $v \in \mathcal{L}_1$. For each $v \in \mathcal{L}_1$ let x_v be an element of G such that $\varphi_n(x_v) = g|_v$ (such an element exists because $\varphi_n(G) = H$). Now define $x \in \mathsf{Aut}(T_{n+1})$ by $x(v) = g(v)$ and $x|_v = x_v$ for all $v \in \mathcal{L}_1$. It follows that $\varphi_{n+1}(x) = g$ and $x \in \mathcal{M}(G)$ as required.

To show that $\psi_{n+1}(\mathsf{st}_{\mathcal{M}(G)}(0)) \supseteq G$, again fix $g \in G$. We must find $x \in \mathsf{st}_{\mathcal{M}(G)}(0)$ such that $\psi_{n+1}(x) = g$. Since G has property $\mathcal{R}_n(H)$ we have $\varphi_n(g) \in H$, so there is $y \in \mathsf{st}_G(0)$ such that $\psi_n(y) = \varphi_n(g)$. By the hypotheses on G and Proposition 3.11, $y|_v \in H$ for all $v \in \mathcal{L}_1$. For each $v \in \mathcal{L}_1$ with $v \neq 0$, let x_v be an element of G such that $\varphi_n(x_v) = y|_v$ (such an element exists because $\varphi_n(G) = H$). Let $x_0 = g$ and observe that

$\varphi_n(x_0) = \psi_n(y) = y|_0$, so we have $\varphi_n(x_v) = y|_v$ for all $v \in \mathcal{L}_1$. Now define $x \in \mathsf{Aut}(T_{n+1})$ by $x(v) = y(v)$ and $x|_v = x_v$ for all $v \in \mathcal{L}_1$. Then $x \in \mathsf{st}(0)$ since $y \in \mathsf{st}(0)$, and $\psi_{n+1}(x) = x|_0 = g$. In addition, $\varphi_{n+1}(x) = y \in G$ so $x \in \mathcal{M}(G)$ as required.

Finally, $\mathcal{M}(G)$ is transitive on $\mathcal{L}_1$ since G is assumed to be transitive on $\mathcal{L}_1$ and we have just shown that $\varphi_{n+1}(\mathcal{M}(G)) = G$. This completes the proof that $\mathcal{M}(G)$ has property $\mathcal{R}_{n+1}(G)$.

(b) We have already assumed that G has sufficient rigid automorphisms. By Lemma 3.8 it suffices to show that $\mathcal{M}(G)$ contains a $(v, v+1)$-rigid automorphism for each $v \in \{0, 1, \ldots, p-2\}$.

Fix $v \in \{0, 1, \ldots, p-2\}$. Then there exists $g \in G$ that is $(v, v+1)$-rigid, i.e. $g(v) = v + 1$ and $g|_v = e$. Define an automorphism $h \in \mathsf{Aut}(T_{n+1})$ as follows. Let $h_{[1]} = g_{[1]}$, so $h(w) = g(w)$ for all $w \in \mathcal{L}_1$. Then for each $w \in \mathcal{L}_1$ let $h|_w$ be any element of G such that $\varphi_n(h|_w) = g|_w$. Such an element exists because $g|_w \in H$ for all $w \in \mathcal{L}_1$ by our assumptions on G and Proposition 3.11, and $\varphi_n(G) = H$. In the case $w = v$, $g|_v = e$ so we choose $h|_v = e$ to ensure that h is $(v, v+1)$-rigid. Then by construction we have $\varphi_{n+1}(h) = g$, and $h|_w \in G$ for all $w \in \mathcal{L}_1$, so $h \in \mathcal{M}(G)$ as required. Since this works for each v, we are done.

(c) The second equality follows from the fact that $|G| = |H|\,|\ker_G(\varphi_n)|$ (see Proposition 3.19(a)). For the first equality, we use Proposisiton 3.19(a) again to obtain:

$$|\mathcal{M}(G)| = |G|\left|\ker_{\mathcal{M}(G)}(\varphi_{n+1})\right|. \tag{6.1}$$

Consider $x \in \ker_{\mathcal{M}(G)}(\varphi_{n+1})$. Since $\varphi_{n+1}(x) = e$ we have $\varphi_n(x|_v) = e$ for all $v \in \mathcal{L}_1$, and since $x \in \mathcal{M}(G)$ we also have $x|_v \in G$ for all $v \in \mathcal{L}_1$. Therefore $x|_v \in \ker_G(\varphi_n)$ for all $v \in \mathcal{L}_1$, which yields the decomposition $x = (x_0, \ldots, x_{p-1})$ where each $x_i \in \ker_G(\varphi_n)$. Furthermore, it is easy to verify from the definition that every element of $\mathsf{Aut}(T_{n+1})$ of this form belongs to $\mathcal{M}(G)$. It follows that $\left|\ker_{\mathcal{M}(G)}(\varphi_{n+1})\right| = |\ker_G(\varphi_n)|^p$ which, combined with (6.1), completes the proof.

(d) The inequality is immediate from the definition, since membership in G' implies membership in G. It follows from part (c) that equality is only possible if $G' = G$. It is also clear from the definition that $\mathcal{M}(\mathsf{Aut}(T_n)) = \mathsf{Aut}(T_{n+1})$. If $\mathcal{M}(G) = \mathsf{Aut}(T_{n+1})$ then by part (a), $G = \varphi_{n+1}(\mathcal{M}(G)) = \varphi_{n+1}(\mathsf{Aut}(T_{n+1})) = \mathsf{Aut}(T_n)$.

(e) Suppose that $M \leq \mathsf{Aut}(T_{n+1})$ has property $\mathcal{R}_{n+1}(G)$. By Propositions 3.9 and 3.11, there exists a group M' conjugate to M such that M'

has property $\mathcal{R}_{n+1}(G)$ and $x|_v \in G$ for all $x \in M'$ and all $v \in \mathcal{L}_1$. Since M' has property $\mathcal{R}_{n+1}(G)$, we also have $\varphi_{n+1}(x) \in G$ for all $x \in M'$. Therefore $M' \subseteq \mathcal{M}(G)$ as required. ☠

6.3. The infinite maximal group $\mathcal{M}^\infty(G)$

We will now see how to iterate the construction in Section 6.2 to produce a subgroup of $\mathsf{Aut}(T)$ with property $\mathcal{R}$. Let G be as above. Define the sequence of groups $\{\mathcal{M}^m(G)\}_{m=0}^\infty$ as follows:

$$\mathcal{M}^0(G) = G; \qquad \mathcal{M}^m(G) = \mathcal{M}(\mathcal{M}^{m-1}(G)) \text{ for } m \geq 1.$$

We must be careful here: our definition of $\mathcal{M}(G)$ assumes that G has sufficient rigid automorphisms. Theorem 6.2(b) ensures that this holds for $\mathcal{M}(G)$ as well, so our definition makes sense. By Theorem 6.2(a), $\mathcal{M}^m(G)$ has property $\mathcal{R}_{n+m}(\mathcal{M}^{m-1}(G))$ for all $m \geq 1$, so by Remark 2.33 we may define the inverse limit of this sequence of groups as in Theorem 2.32, and it acts on the infinite rooted tree T. Let

$$\mathcal{M}^\infty(G) = \varprojlim \mathcal{M}^m(G)$$

be this projective limit. We can extend the results of Theorem 6.2 (except, obviously, for the formula for the size of $\mathcal{M}(G)$) to analogous results about $\mathcal{M}^\infty(G)$.

THEOREM 6.3. *Let G be a group with property $\mathcal{R}_n$ that has sufficient rigid automorphisms. Then:*

- **(a)** *$\mathcal{M}^\infty(G)$ has property $\mathcal{R}$;*
- **(b)** *$\mathcal{M}^\infty(G)_{[n]} = G$;*
- **(c)** *$\mathcal{M}^\infty(G)$ is self-replicating; that is, $\mathsf{st}_{\mathcal{M}^\infty(G)}(v)|_v = \mathcal{M}^\infty(G)$ for all $v \in T$;*
- **(d)** *If $G' \leq G$ then $\mathcal{M}^\infty(G') \leq \mathcal{M}^\infty(G)$, with equality if and only if $G' = G$. In particular, $\mathcal{M}^\infty(G) = \mathsf{Aut}(T)$ if and only if $G = \mathsf{Aut}(T_n)$;*
- **(e)** *If $\widetilde{G} \leq \mathsf{Aut}(T)$ has property $\mathcal{R}$ and $\widetilde{G}_{[n]} = G$, then $\widetilde{G}$ is conjugate to a subgroup of $\mathcal{M}^\infty(G)$.*

PROOF. **(a)** Follows from Theorem 2.32.

(b) Also follows from Theorem 2.32.

(c) By Corollary 3.17 and part (a), it suffices to show that $\mathcal{M}^\infty(G)$ is self-similar; in other words, that $\mathsf{st}_{\mathcal{M}^\infty(G)}(v)|_v \subseteq \mathcal{M}^\infty(G)$ for all $v \in T$. Fix $v \in T$ and let $g \in \mathcal{M}^\infty(G)$ such that $g(v) = v$. We must show that

$g|_v \in \mathcal{M}^\infty(G)$. It follows from the definition of $\mathcal{M}^\infty(G)$ and Proposition 2.30 (along with Remark 2.33) that

$$\mathcal{M}^\infty(G) = \varprojlim \mathcal{M}^m(G) = \{g \in \mathsf{Aut}(T) : g_{[n+m]} \in \mathcal{M}^m(G) \text{ for all } m \geq 0\}.$$

Thus, we need to show that $(g|_v)_{[n+m]} \in \mathcal{M}^m(G)$ for all $m \geq 0$. Fix $m \geq 0$. Assume for the moment that $v \in \mathcal{L}_1$. Then

$$(g|_v)_{[n+m]} = g_{[n+m+1]}|_v. \tag{6.2}$$

Since $g \in \mathcal{M}^\infty(G)$, we know that $g_{[n+m+1]} \in \mathcal{M}^{m+1}(G)$. Now by induction and Theorem 6.2(b), $\mathcal{M}^{m+1}(G)$ has sufficient rigid automorphisms, so $g_{[n+m+1]}|_v \in \mathcal{M}^m(G)$ by Proposition 3.11. Then $(g|_v)_{[n+m]} \in \mathcal{M}^m(G)$ by (6.2), which is what we wanted to show. Since this holds for all $m \geq 0$, we have shown that $g|_v \in \mathcal{M}^\infty(G)$. So far, we have only shown this for all $v \in \mathcal{L}_1$, but the same induction argument from the proof of Proposition 3.13 extends this to all $v \in T$ and we are done.

(d) Suppose $G' \leq G$. By induction using Theorem 6.2(d) we have $\mathcal{M}^m(G') \leq \mathcal{M}^m(G)$ for all $m \geq 0$. The first assertion now follows from Proposition 2.31(d).

For the 'if' direction of the final assertion, induction using Theorem 6.2(d) yields $\mathcal{M}^m(\mathsf{Aut}(T_n)) = \mathsf{Aut}(T_{n+m})$ for all m and n. Therefore

$$\mathcal{M}^\infty(\mathsf{Aut}(T_n)) = \varprojlim \mathsf{Aut}(T_{n+m}) = \mathsf{Aut}(T),$$

and the other direction now follows from part (a) since $\mathsf{Aut}(T)_{[n]} = \mathsf{Aut}(T_n)$.

(e) Since $\widetilde{G}$ has property $\mathcal{R}$, Theorem 2.32 implies that $\widetilde{G} = \varprojlim \widetilde{G}_{[m]}$. Note that $\widetilde{G}_{[m]}$ has property $\mathcal{R}_m(\widetilde{G}_{[m-1]})$ for all $m \geq 1$; in particular, $\widetilde{G}_{[n+1]}$ has property $\mathcal{R}_{n+1}(G)$ (since $\widetilde{G}_{[n]} = G$) and so by Theorem 6.2(e), $\widetilde{G}_{[n+1]}$ is conjugate to a subgroup of $\mathcal{M}(G)$. We will prove by induction that for each $m > n$, $\widetilde{G}_{[m]}$ is conjugate to a subgroup of $\mathcal{M}^{m-n}(G)$. Suppose this claim holds for some $m > n$. Then, since $\widetilde{G}_{[m+1]}$ has property $\mathcal{R}_{m+1}(\widetilde{G}_{[m]})$ and $\mathcal{M}^{m+1-n}(G) = \mathcal{M}(\mathcal{M}^{m-n}(G))$, Theorem 6.2(e) implies that $\widetilde{G}_{[m+1]}$ is conjugate to a subgroup of $\mathcal{M}^{m+1-n}(G)$, completing the inductive step. It now follows from Proposition 3.6 that this conjugacy holds in the inverse limit as well, hence $\widetilde{G}$ is conjugate to a subgroup of $\mathcal{M}^\infty(G)$ as required. ☠

6.4. Finitely constrained groups

It turns out that all of the maximal groups $\mathcal{M}^\infty(G)$ belong to a class of groups known as finitely constrained groups. These groups were originally defined by Grigorchuk [**Gri05**, Definition 7.1] as *groups of finite type*, by analogy with subshifts of finite type in symbolic dynamics. The name *finitely constrained groups* has been preferred more recently, for example in [**Šun07**, Definition 5] and [**Šun11**, Definition 3]. In terms of our notation, the definition is as follows:

Definition 6.4. Let $\mathcal{F}$ be any subset of $\mathsf{Aut}(T_n)$ for some $n \geq 1$. We call $\mathcal{F}$ a set of *forbidden patterns* of size n. Define

$$\mathsf{FC}(\mathcal{F}) = \{g \in \mathsf{Aut}(T) : (g|_v)_{[n]} \notin \mathcal{F} \text{ for all } v \in T\}.$$

When $\mathsf{FC}(\mathcal{F})$ is a subgroup of $\mathsf{Aut}(T)$ we call it the *finitely constrained group* defined by the set of forbidden patterns $\mathcal{F}$. Note that $\mathcal{F}$ is finite since it is contained in $\mathsf{Aut}(T_n)$.

For our purposes, $\mathcal{F}$ will always be the complement of a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$. In this situation $\mathsf{FC}(\mathcal{F})$ will always be a group:

Proposition 6.5. *Let $\mathcal{F}$ be the complement of a subgroup of $\mathsf{Aut}(T_n)$. Then $\mathsf{FC}(\mathcal{F})$ is a subgroup of $\mathsf{Aut}(T)$.*

Proof. Let G be a subgroup of $\mathsf{Aut}(T_n)$ and let $\mathcal{F} = \mathsf{Aut}(T_n) \setminus G$. It is immediate from the definition that $e \in \mathsf{FC}(\mathcal{F})$ since $e \notin \mathcal{F}$. Suppose that $g, h \in \mathsf{FC}(\mathcal{F})$. Then:

$$\begin{aligned}\left((g^{-1}h)\,|_v\right)_{[n]} &= \left(g^{-1}|_{h(v)}h|_v\right)_{[n]} \\ &= \left(\left(g|_{g^{-1}h(v)}\right)_{[n]}\right)^{-1}(h|_v)_{[n]}\,. \end{aligned} \tag{6.3}$$

By our assumptions on g and h, both $\left(g|_{g^{-1}h(v)}\right)_{[n]}$ and $(h|_v)_{[n]}$ belong to G, and hence so does (6.3) since G is a group. Therefore by definition $g^{-1}h \in \mathsf{FC}(\mathcal{F})$ which completes the proof. ☠

It now follows from [**Šun07**, Theorem 3] that in this case $\mathsf{FC}(\mathcal{F})$ is always a closed, self-similar subgroup of $\mathsf{Aut}(T)$.

Fix $n \geq 2$, let G be a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$, and let $\mathcal{F}$ be the complement of G in $\mathsf{Aut}(T_n)$. We aim to show that $\mathcal{M}^\infty(G) = \mathsf{FC}(\mathcal{F})$. To do this, first observe that we can restate the definition of $\mathsf{FC}(\mathcal{F})$ in an obvious way:

$$\mathsf{FC}(\mathcal{F}) = \{g \in \mathsf{Aut}(T) : (g|_v)_{[n]} \in G \text{ for all } v \in T\}. \tag{6.4}$$

We will show that this group is equal to $\mathcal{M}^\infty(G)$ in two stages; first we will provide a similar characterisation for the finite groups $\mathcal{M}^m(G)$, and then use the inverse limit to pass to $\mathcal{M}^\infty(G)$.

PROPOSITION 6.6. *Let n, G and $\mathcal{F}$ be as above. Then:*

- **(a)** $\mathcal{M}^m(G) = \{g \in \mathsf{Aut}(T_{n+m}) : (g|_v)_{[n]} \in G$ *for all* $v \in T_m\}$ *for all* $m \geq 1$*;*
- **(b)** $\mathcal{M}^\infty(G) = \mathsf{FC}(\mathcal{F})$*.*

PROOF. **(a)** By induction. Let us start with the base step, $m = 1$. When v is the root of T_1 the restriction $g|_v$ is just g, so $(g|_v)_{[n]} = \varphi_{n+1}(g)$. Otherwise $v \in \mathcal{L}_1$ which means $g|_v \in \mathsf{Aut}(T_n)$, so $(g|_v)_{[n]} = g|_v$. Thus the claim for $m = 1$ simply reduces to the definition of $\mathcal{M}(G)$.

Suppose that the claim is true for some $m \geq 1$. Recall the definition:

$$\mathcal{M}^{m+1}(G) = \{g \in \mathsf{Aut}(T_{n+m+1}) : \varphi_{n+m+1}(g) \in \mathcal{M}^m(G) \text{ and } g|_v \in \mathcal{M}^m(G) \text{ for all } v \in \mathcal{L}_1\}.$$

We must show that this set is equal to

$$\{g \in \mathsf{Aut}(T_{n+m+1}) : (g|_v)_{[n]} \in G \text{ for all } v \in T_{m+1}\}.$$

To do this, we will show both inclusions separately.

$\subseteq$: Let $g \in \mathcal{M}^{m+1}(G)$. We must show that $(g|_v)_{[n]} \in G$ for all $v \in T_{m+1}$. First suppose that $v \in \mathcal{L}_k$ for some k with $0 \leq k \leq m$. Now

$$(g|_v)_{[n]} = (g_{[n+m]}|_v)_{[n]} = (\varphi_{n+m+1}(g)|_v)_{[n]}$$

since $g_{[n+m]}|_v \in \mathsf{Aut}(T_{n+m-k})$ and $n+m-k \geq n$. But $\varphi_{n+m+1}(g) \in \mathcal{M}^m(G)$, so by the inductive hypothesis $(g|_v)_{[n]} \in G$ as required. To handle the case where $v \in \mathcal{L}_{m+1}$, write $v = uw$ where $u \in \mathcal{L}_1$ and $w \in \mathcal{L}_m$. Now $g|_u \in \mathcal{M}^m(G)$ since $u \in \mathcal{L}_1$ so by the inductive hypothesis $((g|_u)|_w)_{[n]} \in G$. But $(g|_u)|_w = g|_v$ since $v = uw$, so we are done.

$\supseteq$: Let $g \in \mathsf{Aut}(T_{n+m+1})$ and suppose that $(g|_v)_{[n]} \in G$ for all $v \in T_{m+1}$. We must show that $g \in \mathcal{M}^{m+1}(G)$; that is, we need to show that $\varphi_{n+m+1}(g) \in \mathcal{M}^m(G)$ and $g|_v \in \mathcal{M}^m(G)$ for all $v \in \mathcal{L}_1$.

For the first part, fix $v \in T_m$, so that $v \in \mathcal{L}_k$ where $0 \leq k \leq m$. Observe that $\varphi_{n+m+1}(g)|_v \in \mathsf{Aut}(T_{n+m-k})$ and $n+m-k \geq n$, so $(\varphi_{n+m+1}(g)|_v)_{[n]} = (g|_v)_{[n]} \in G$ by our hypothesis on g. Since $v \in T_m$ was arbitrary, it follows from the inductive hypothesis that $\varphi_{n+m+1}(g) \in \mathcal{M}^m(G)$.

Now let $v \in \mathcal{L}_1$ and consider $g|_v$. Fix $w \in T_m$. Then $vw \in T_{m+1}$ and $((g|_v)|_w)_{[n]} = (g|_{vw})_{[n]} \in G$ by our hypothesis on g. This is true for all $w \in T_m$, so by the inductive hypothesis $g|_v \in \mathcal{M}^m(G)$. This completes the proof that $g \in \mathcal{M}^{m+1}(G)$.

Both inclusions have now been proven so the inductive step is complete and we are done.

(b) Recall (see the proof of Theorem 6.3(c) above) that

$$\mathcal{M}^\infty(G) = \{g \in \mathsf{Aut}(T) \,:\, g_{[n+m]} \in \mathcal{M}^m(G) \text{ for all } m \geq 0\}.$$

We will show that this is equal to $\mathsf{FC}(\mathcal{F})$ using (6.4). Again we will show each inclusion separately.

$\subseteq$: Suppose that $g \in \mathcal{M}^\infty(G)$. Fix $v \in T$; then $v \in \mathcal{L}_k$ for some $k \geq 0$. Observe that $(g|_v)_{[n]} = (g_{[n+k]}|_v)_{[n]} \in G$ by part (a), since $g_{[n+k]} \in \mathcal{M}^k(G)$. Since this is true for all $v \in T$, it follows from (6.4) that $g \in \mathsf{FC}(\mathcal{F})$.

$\supseteq$: Suppose that $g \in \mathsf{FC}(\mathcal{F})$. Then by (6.4), $(g|_v)_{[n]} \in G$ for all $v \in T$. We must show that $g_{[n+m]} \in \mathcal{M}^m(G)$ for all $m \geq 0$, which by part (a) is equivalent to $(g_{[n+m]}|_v)_{[n]} \in G$ for all $m \geq 0$ and all $v \in T_m$. Fix $m \geq 0$ and let $v \in T_m$, so $v \in \mathcal{L}_k$ where $0 \leq k \leq m$. Since $g_{[n+m]}|_v \in \mathsf{Aut}(T_{n+m-k})$ and $n+m-k \geq n$, we have $(g_{[n+m]}|_v)_{[n]} = ((g|_v)_{[n+m-k]})_{[n]} = (g|_v)_{[n]} \in G$. This holds for all $m \geq 0$ and $v \in T_m$ so the proof is complete. ☠

Thus each maximal group $\mathcal{M}^\infty(G)$ is a finitely constrained group defined by a set of forbidden patterns of size n, namely $\mathsf{Aut}(T_n) \setminus G$. Since $\mathcal{M}^\infty(G)$ always has property $\mathcal{R}$ by Theorem 6.3(a) and is self-replicating by Theorem 6.3(c), we obtain the following result about finitely constrained groups, which appears to be new:

COROLLARY 6.7. *Let G be a group with property $\mathcal{R}_n$ that has sufficient rigid automorphisms, and let $\mathcal{F}$ be the complement of G in $\mathsf{Aut}(T_n)$. Then the finitely constrained group $\mathsf{FC}(\mathcal{F})$ is a self-replicating group with property $\mathcal{R}$.*

The new part of this result is that for this particular $\mathcal{F}$, the group $\mathsf{FC}(\mathcal{F})$ is self-replicating; as stated above, it is already known that all finitely constrained groups are self-similar.

6.5. Hausdorff dimension

Since $\mathcal{M}^\infty(G)$ is a closed subgroup of $\mathsf{Aut}(T)$, it has a well-defined Hausdorff dimension, as shown in [**BS97**]. For subgroups of the pro-p group $\mathsf{Aut}_p(T)$, this dimension can be taken relative to $\mathsf{Aut}_p(T)$ rather than $\mathsf{Aut}(T)$ as in [**Šun07**]. We can give a simple formula in both cases:

PROPOSITION 6.8. *Let G and H be as above, so that G has property $\mathcal{R}_n(H)$. Then the Hausdorff dimension of $\mathcal{M}^\infty(G)$ in $\mathsf{Aut}(T)$ is*

$$\begin{aligned}\dim_{\mathcal{H}}(\mathcal{M}^\infty(G)) &= p^{1-n}\log_{p!}|\ker_G(\varphi_n)| \\ &= p^{1-n}\log_{p!}\frac{|G|}{|H|}.\end{aligned}$$

In addition, if $G \leq \mathsf{Aut}_p(T_n)$, then $\mathcal{M}^\infty(G) \leq \mathsf{Aut}_p(T)$ and its Hausdorff dimension in $\mathsf{Aut}_p(T)$ is

$$\begin{aligned}\dim_{\mathcal{H}}(\mathcal{M}^\infty(G)) &= p^{1-n}\log_{p}|\ker_G(\varphi_n)| \\ &= p^{1-n}\log_{p}\frac{|G|}{|H|}\end{aligned}$$

which is always rational.

PROOF. The Hausdorff dimension is given by the formula from [**BS97**]:

$$\begin{aligned}\dim_{\mathcal{H}}(\mathcal{M}^\infty(G)) &= \liminf_{m\to\infty}\frac{\log[\mathcal{M}^\infty(G) : \mathsf{st}_{\mathcal{M}^\infty(G)}(\mathcal{L}_m)]}{\log[\mathsf{Aut}(T) : \mathsf{st}(\mathcal{L}_m)]} \\ &= \liminf_{m\to\infty}\frac{\log\left|\mathcal{M}^\infty(G)_{[m]}\right|}{\log\left|\mathsf{Aut}(T)_{[m]}\right|} \\ &= \liminf_{m\to\infty}\frac{\log|\mathcal{M}^m(G)|}{\log|\mathsf{Aut}(T_{n+m})|} \qquad \text{(replacing } m \text{ with } n+m) \\ &= \liminf_{m\to\infty}\frac{\log|\mathcal{M}^m(G)|}{\frac{p^{n+m}-1}{p-1}\log p!} \qquad \text{by (2.17)} \\ &= \frac{p-1}{\log p!}\liminf_{m\to\infty}\frac{\log|\mathcal{M}^m(G)|}{p^{n+m}-1}, \end{aligned} \tag{6.5}$$

so it remains to find a formula for $|\mathcal{M}^m(G)|$. Let $a_m = |\mathcal{M}^m(G)|$ for $m \geq 0$. We have $a_0 = |G|$ and $a_1 = |\mathcal{M}(G)| = |G|\,|\ker_G(\varphi_n)|^p$ by Theorem 6.2(c). For $m \geq 2$ we have the recurrence relation

$$a_m = |\mathcal{M}^m(G)| = \frac{\left|\mathcal{M}^{m-1}(G)\right|^{p+1}}{|\mathcal{M}^{m-2}(G)|^p} = \frac{{a_{m-1}}^{p+1}}{{a_{m-2}}^p}.$$

If we let $b_m = \log a_m$ then we get the linear recurrence

$$b_m = (p+1)b_{m-1} - pb_{m-2}, \qquad m \geq 2$$

with $b_0 = \log|G|$ and $b_1 = \log|G| + p\log|\ker_G(\varphi_n)|$. The characteristic equation for this recurrence relation is

$$0 = \lambda^2 - (p+1)\lambda + p = (\lambda - p)(\lambda - 1)$$

which has roots 1 and p, so the solution is $b_m = c_0 + c_1 p^m$ for some constants c_0 and c_1. Using the values of b_0 and b_1 we get $c_0 = \log|G| - \frac{p}{p-1}\log|\ker_G(\varphi_n)|$ and $c_1 = \frac{p}{p-1}\log|\ker_G(\varphi_n)|$. This yields

$$b_m = \log|G| + \frac{p(p^m-1)}{p-1}\log|\ker_G(\varphi_n)|$$

and hence

$$|\mathcal{M}^m(G)| = a_m = \exp(b_m) = |G|\,|\ker_G(\varphi_n)|^{\frac{p(p^m-1)}{p-1}}.$$

Substituting this back into (6.5) yields:

$$\begin{aligned}
\dim_{\mathcal{H}}(\mathcal{M}^\infty(G)) &= \frac{p-1}{\log p!}\liminf_{m\to\infty}\frac{\log|G| + \frac{p(p^m-1)}{p-1}\log|\ker_G(\varphi_n)|}{p^{n+m}-1} \\
&= \frac{p-1}{p^n\log p!}\liminf_{m\to\infty}\frac{\frac{p(p^m-1)}{p-1}\log|\ker_G(\varphi_n)|}{p^m-p^{-n}} \\
&= \frac{p\log|\ker_G(\varphi_n)|}{p^n\log p!}\liminf_{m\to\infty}\frac{p^m-1}{p^m-p^{-n}} \\
&= \frac{\log|\ker_G(\varphi_n)|}{p^{n-1}\log p!}\cdot 1 \\
&= p^{1-n}\log_{p!}|\ker_G(\varphi_n)|
\end{aligned}$$

as claimed. The second formula is obtained from Proposition 3.19(a).

The calculation is identical for the Hausdorff dimension in $\mathsf{Aut}_p(T)$, except that $|\mathsf{Aut}_p(T_n)| = p^{\frac{p^n-1}{p-1}}$ by (3.5), which has the sole effect of replacing $\log p!$ with $\log p$. The dimension is always rational in this case because G is a p-group so $|\ker_G(\varphi_n)|$ is a power of p. This completes the proof. ☠

Remark. This formula for the Hausdorff dimension of a finitely constrained group has been independently discovered for the special case $p = 2$ in [**PŠ**, Lemma 8].

Note that the Hausdorff dimension of $\mathcal{M}^\infty(G)$ depends only on two quantities: the level n on which G acts and the size of $\ker_G(\varphi_n)$, or equivalently, the ratio of $|G|$ to $|H|$, which can be viewed as the "expansion factor" of G when moving from T_{n-1} to T_n.

We can deduce from this formula that every maximal group $\mathcal{M}^\infty(G)$ has strictly positive Hausdorff dimension, and only $\mathsf{Aut}(T)$ achieves the maximum possible dimension of 1. These are known results about finitely constrained groups; Theorem 4(a) from [**Šun07**] shows that $\dim_{\mathcal{H}}(\mathcal{M}^\infty(G))$ is always positive — and rational in $\mathsf{Aut}_p(T)$ — when G is a p-group.

PROPOSITION 6.9. *Let G be a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$. Then*

(a) $\dim_{\mathcal{H}}(\mathcal{M}^\infty(G)) > 0$*, and*

(b) $\dim_{\mathcal{H}}(\mathcal{M}^\infty(G)) = 1$ *if and only if* $G = \mathsf{Aut}(T_n)$.

PROOF. Part (a) is immediate, since $\ker_G(\varphi_n)$ is always nontrivial by Proposition 3.19(b).

By Proposition 6.3(d), $\mathcal{M}^\infty(\mathsf{Aut}(T_n)) = \mathsf{Aut}(T)$, so the 'if' direction of part (b) says that $\dim_{\mathcal{H}}(\mathsf{Aut}(T)) = 1$. This can be readily deduced from the Hausdorff dimension formula. To prove the 'only if' direction, suppose that $\dim_{\mathcal{H}}(\mathcal{M}^\infty(G)) = 1$. Then, by Proposition 6.8, $\log_{p!} \frac{|G|}{|H|} = p^{n-1}$ so $|G| = p!^{p^{n-1}} |H|$, and hence $|\ker_G(\varphi_n)| = p!^{p^{n-1}}$ by Proposition 3.19(a). We will show that $G = \mathsf{Aut}(T_n)$.

Lemma 2.23 implies that $G_{[k]}$ has property $\mathcal{R}_k(G_{[k-1]})$ for $2 \le k \le n$ (note that $H = G_{[n-1]}$). Let $K_k = \ker_{G_{[k]}}(\varphi_k)$ for $2 \le k \le n$, and for convenience let $K_1 = G_{[1]}$ (this is consistent, based on the understanding that level 0 is just the root so $G_{[0]}$ is trivial and φ_1 is the trivial map). Proposition 3.19(d) implies that $|K_k| \le |K_{k-1}|^p$ for each k, so:

$$|K_n| \le |K_{n-1}|^p \le |K_{n-2}|^{p^2} \le \cdots \le |K_1|^{p^{n-1}}. \tag{6.6}$$

But $|K_n| = |\ker_G(\varphi_n)| = p!^{p^{n-1}}$, and $|K_1| \le |\mathsf{Aut}(T_1)| = p!$ so there must be equality throughout (6.6) and $|K_k| = p!^{p^{k-1}}$ for each k. Now $\left|G_{[k]}\right| = |K_k| \left|G_{[k-1]}\right|$ for each k by Proposition 3.19(a), hence:

$$\begin{aligned} |G| &= \left|G_{[n]}\right| \\ &= |K_n| \cdots |K_2| \left|G_{[1]}\right| \\ &= |K_n| \cdots |K_2| \, |K_1| \\ &= p!^{(p^{n-1}+\cdots+p+1)} \\ &= p!^{\frac{p^n-1}{p-1}} \\ &= |\mathsf{Aut}(T_n)|. \end{aligned}$$

It follows that $G = \mathsf{Aut}(T_n)$ as required. ☠

CHAPTER 7

Automaton groups

7.1. Introduction

In Chapter 6 we described one way to extend any subgroup G of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$ to a maximal subgroup $\mathcal{M}^\infty(G)$ of $\mathsf{Aut}(T)$ with property $\mathcal{R}$. We also saw in Section 5.4 that these maximal groups form a countable dense subset of all (conjugacy classes of) groups with property $\mathcal{R}$.

We may ask if there are any other general constructions which will take a subgroup G of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$ and extend it to a subgroup of $\mathsf{Aut}(T)$ with property $\mathcal{R}$, such that the resulting group is non-maximal, or even minimal. The purpose of this chapter is to describe one such construction (strictly speaking, a family of constructions) using finite automata.

Groups generated by automata, or *automaton groups*, have been studied extensively (see [**GNS00**] for example) and provide a rich family of interesting examples of self-similar groups. Unfortunately (for our purposes) these groups are not always self-replicating. The class of spherically transitive, self-replicating automaton groups does not appear to have been studied specifically. Indeed, in general it can be extremely difficult to check whether or not an automaton group is self-replicating. This means that although it is easy to find examples and constructions of automaton groups in the literature, many of them are not directly useful to us. Still, plenty of interesting and relevant examples have been studied, some of which we have already seen such as the odometer (in Example 4.1), and other well-known groups such as the lamplighter group [**GŻ01**], the first Grigorchuk group [**Gri05**], and *spinal groups* similar to the Grigorchuk group [**Šun07**].

The important property of all the groups constructed in this chapter is that they are *topologically finitely generated*; that is, they are the topological closure in $\mathsf{Aut}(T)$ of a finitely generated group. Although the original aim of this construction was to produce non-maximal groups, it can actually produce $\mathcal{M}^\infty(G)$ if $\mathcal{M}^\infty(G)$ happens to be topologically finitely generated. As we saw in Section 5.4, such cases appear to be very rare. The most well-known example of this type, the Grigorchuk group, will be discussed in Example 7.16.

7.2. Preliminaries

We begin with some terminology and notation.

DEFINITION 7.1. An *automaton* (strictly speaking, a *synchronous transducer*) is a quadruple $\mathcal{A} = (\mathsf{X}, Q, \tau, \lambda)$ where:

- X is a finite set of symbols called the *alphabet*;
- Q is a set, called the set of *states*;
- $\tau : Q \times \mathsf{X} \to Q$ is the *transition function*;
- $\lambda : Q \times \mathsf{X} \to \mathsf{X}$ is the *output function.*

Following Section 2.3, we will always use the alphabet $\mathsf{X} = \{0, 1, \dots, p-1\}$. An automaton is said to be *finite* if Q is finite. An automaton is *invertible* if for each state $q \in Q$ the corresponding output function $\lambda(q, -) : \mathsf{X} \to \mathsf{X}$ is invertible (i.e. a permutation of X). The function $\lambda(q, -)$ will be written more concisely as λ_q. For our purposes we will only consider finite, invertible automata.

We can represent an automaton by a labelled directed graph as follows. The vertices are the states Q, with each $q \in Q$ labelled by the corresponding output permutation λ_q. There is a directed edge labelled a from q to q' if and only if $\tau(q, a) = q'$. We refer to these as *incoming* edges (or incoming transitions) to the state q', or *outgoing* edges (transitions) from the state q.

Our examples here will focus on the case where $p = 2$, i.e. $\mathsf{X} = \{0, 1\}$, although the results will apply to arbitrary $p \geq 2$. For convenience we will denote the only nontrivial permutation of $\{0, 1\}$ by σ and the trivial permutation by e.

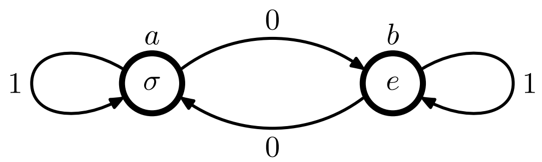

FIGURE 7.1. A simple automaton with two states a and b over the alphabet $\mathsf{X} = \{0, 1\}$.

EXAMPLE 7.2. Let $p = 2$ and $Q = \{a, b\}$ with output permutations $\lambda_a = \sigma$ and $\lambda_b = e$, and transitions $\tau(a, 0) = b$, $\tau(a, 1) = a$, $\tau(b, 0) = a$ and $\tau(b, 1) = b$. The graph of this automaton is shown in Figure 7.1.

The idea is that each state q of an invertible automaton represents an automorphism of the tree T, which we may identify with X^* as described in Section 2.3 ($\mathsf{X}^* = \bigcup_{n\geq 0} \mathsf{X}^n$ where X^n is the set of words of length n over X, previously referred to as $\mathcal{L}_n$). We think of the automaton "reading" a word $v \in \mathsf{X}^*$ (representing a vertex $v \in T$) one symbol at a time. If the automaton is in the state q and reads the symbol $w \in \mathsf{X}$ then it "writes" the symbol $\lambda(q, w)$ and transitions into the state $\tau(q, w)$, ready to read the next symbol. When the whole word v has been read, the word formed by the written symbols represents $q(v)$.

More formally, the transition and output functions can be extended by induction to functions $\widetilde{\tau} : Q \times \mathsf{X}^* \to Q$ and $\widetilde{\lambda} : Q \times \mathsf{X}^* \to \mathsf{X}^*$ as follows. For all $w \in \mathsf{X}$, define $\widetilde{\tau}(q, w) = \tau(q, w)$ and $\widetilde{\lambda}(q, w) = \lambda(q, w)$. For $n \geq 2$ and $v = w_1 \cdots w_n \in \mathsf{X}^n$, define:

$$\begin{aligned} &\widetilde{\tau}(q, w_1 \cdots w_n) = \widetilde{\tau}(\tau(q, w_1), w_2 \cdots w_n) \\ \text{and} \qquad &\widetilde{\lambda}(q, w_1 \cdots w_n) = \lambda(q, w_1)\widetilde{\lambda}(\tau(q, w_1), w_2 \cdots w_n). \end{aligned}$$

For each $q \in Q$ we will use the notation $\widetilde{\lambda}_q$ when convenient, to denote the function $\widetilde{\lambda}(q, -) : \mathsf{X}^* \to \mathsf{X}^*$. The inductive formula for $\widetilde{\lambda}$ above may then be rewritten more concisely as:

$$\widetilde{\lambda}_q(w_1 \cdots w_n) = \lambda_q(w_1)\widetilde{\lambda}_{\tau(q,w_1)}(w_2 \cdots w_n).$$

EXAMPLE 7.3. Referring again to the automaton shown in Figure 7.1, let us calculate $\widetilde{\lambda}_a(001)$. Observing that $\lambda_a = \sigma$ swaps the symbols 0 and 1, and $\lambda_b = e$ fixes both 0 and 1, we have:

$$\begin{aligned} \widetilde{\lambda}_a(001) &= \lambda(a, 0)\widetilde{\lambda}_{\tau(a,0)}(01) \\ &= 1\widetilde{\lambda}_b(01) \\ &= 1\lambda_b(0)\widetilde{\lambda}_{\tau(b,0)}(1) \\ &= 10\widetilde{\lambda}_a(1) \\ &= 100. \end{aligned}$$

It is clear from this calculation that each $\widetilde{\lambda}_q$ preserves the length of words; that is, $\left|\widetilde{\lambda}_q(v)\right| = |v|$ for all $v \in \mathsf{X}^*$ and all $q \in Q$.

PROPOSITION 7.4. *If $\mathcal{A}$ is an invertible automaton, then for each state $q \in Q$ the function $\widetilde{\lambda}_q : \mathsf{X}^* \to \mathsf{X}^*$ is an automorphism of the rooted tree X^*.*

PROOF. We must show that each $\widetilde{\lambda}_q$ preserves the edge relation in the tree; that is, for all $q \in Q$, all $v \in \mathsf{X}^*$ and all $w \in \mathsf{X}$, $\widetilde{\lambda}_q(vw) = \widetilde{\lambda}_q(v)w'$ for some $w' \in \mathsf{X}$. We prove this by induction on $|v|$. For the base step, suppose that $|v| = 1$, so that $v \in \mathsf{X}$. Now for all $q \in Q$ and all $w \in \mathsf{X}$:

$$\widetilde{\lambda}_q(vw) = \lambda_q(v)\widetilde{\lambda}_{\tau(q,v)}(w) = \widetilde{\lambda}_q(v)w'$$

where $w' = \lambda_{\tau(q,v)}(w)$, as required. Note that $\widetilde{\lambda}_q(v) = \lambda_q(v)$ since $v \in \mathsf{X}$. Now fix $n \geq 1$ and suppose for all $q \in Q$, all $v \in \mathsf{X}^n$ and all $w \in \mathsf{X}$, that there exists $w' \in \mathsf{X}$ such that $\widetilde{\lambda}_q(vw) = \widetilde{\lambda}_q(v)w'$. Fix $v \in \mathsf{X}^{n+1}$, $q \in Q$ and $w \in \mathsf{X}$. Write $v = w_1 \cdots w_{n+1}$. Then:

$$\begin{aligned} \widetilde{\lambda}_q(v) &= \lambda_q(w_1)\widetilde{\lambda}_{\tau(q,w_1)}(w_2 \cdots w_{n+1}) \\ \text{and} \qquad \widetilde{\lambda}_q(vw) &= \lambda_q(w_1)\widetilde{\lambda}_{\tau(q,w_1)}(w_2 \cdots w_{n+1}w). \end{aligned}$$

Now $w_2 \cdots w_{n+1} \in \mathsf{X}^n$ so, by the inductive hypothesis, there exists $w' \in \mathsf{X}$ such that $\widetilde{\lambda}_{\tau(q,w_1)}(w_2 \cdots w_{n+1}w) = \widetilde{\lambda}_{\tau(q,w_1)}(w_2 \cdots w_{n+1})w'$, whence $\widetilde{\lambda}_q(vw) = \widetilde{\lambda}_q(v)w'$ as required.

It remains to show that $\widetilde{\lambda}_q$ is a bijection on X^*. This follows immediately from the fact that $\mathcal{A}$ is invertible and so λ_q is a bijection for each $q \in Q$. ☠

We will abuse notation and use the same symbol to denote both a state and its corresponding automorphism; so for $q \in Q$ and $v \in \mathsf{X}^*$ we simply write $q(v)$ instead of $\widetilde{\lambda}_q(v)$. For instance, in Example 7.3 we can write $a(001) = 100$.

A careful reading of the proof of Proposition 7.4 and the definitions in Section 2.3, leads immediately to the following results:

PROPOSITION 7.5. *Let $\mathcal{A}$ be an invertible automaton, and let q be the automorphism of X^* corresponding to a state $q \in Q$. Then:*

- **(a)** *The output permutation $\lambda_q \in \mathsf{Sym}(\mathsf{X})$ is the root permutation of q (i.e. the action of q on level 1 of the tree);*
- **(b)** *For each word $v \in \mathsf{X}^*$, the automorphism $\tau(q, v)$ is equal to the vertex restriction $q|_v$;*
- **(c)** *The wreath product decomposition (see (2.13)) of q is*

$$q = \lambda_q\left(\tau(q, 0), \ldots, \tau(q, p-1)\right).$$

The *group* of the invertible automaton $\mathcal{A}$, denoted $\mathcal{G}(\mathcal{A})$, is defined to be the subgroup of $\mathsf{Aut}(\mathsf{X}^*)$ generated by the automorphisms corresponding to the states of $\mathcal{A}$. That is, $\mathcal{G}(\mathcal{A}) = \langle q : q \in Q\rangle$. Obviously $\mathcal{G}(\mathcal{A})$ is finitely generated if $\mathcal{A}$ is finite. By Proposition 7.5(b) all vertex restrictions of the generators of $\mathcal{G}(\mathcal{A})$ belong to $\mathcal{G}(\mathcal{A})$ and so it follows that $\mathcal{G}(\mathcal{A})$ is self-similar.

Proposition 7.5(c) shows us how to read off a recursive decomposition of each generator directly from the automaton.

Example 7.6. For the automaton in Figure 7.1, the recursive decompositions are $a = \sigma(b, a)$ and $b = (a, b)$. This group $\mathcal{G}(\mathcal{A}) = \langle a, b\rangle$ is actually the lamplighter group (see [**GŻ01**]), isomorphic to $(\mathbb{Z}/2\mathbb{Z}) \wr \mathbb{Z}$.

Example 7.7. Let $\mathcal{A}$ be the automaton in Figure 7.2. Notice that the output permutation of b is trivial so b acts trivially on level 1. In fact, the recursive decomposition $b = (b, b)$ shows that b itself is trivial so that $\mathcal{G}(\mathcal{A}) = \langle a, b\rangle = \langle a\rangle$. Therefore $a = \sigma(e, a)$ which, as we saw in Example 4.1, generates the odometer group, isomorphic to $\mathbb{Z}$; the closure of this group in $\mathsf{Aut}(T)$ is the action of $\mathbb{Z}_2$ on T.

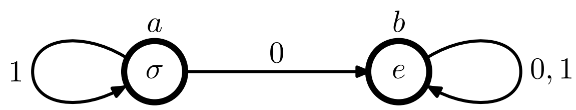

$a = \sigma(b, a)$ $\qquad$ $b = (b, b)$

Figure 7.2. The automaton generating the odometer group for $p = 2$.

Although it is easy to show that any automaton group $\mathcal{G}(\mathcal{A})$ is self-similar, it is much more difficult to determine whether or not it is self-replicating. However we can give a sufficient condition that is not too onerous to check:

Lemma 7.8. *Let $\mathcal{A}$ be an invertible automaton. Suppose that for every $q \in Q$ there is $q' \in Q$ and $w \in \mathsf{X}$ such that $\tau(q', w) = q$ and $\lambda(q', w) = w$. Then $\mathcal{G}(\mathcal{A})$ is self-replicating.*

Proof. Immediate from the definitions and from Proposition 7.5.

This condition can be weakened if we assume the existence of rigid automorphisms in $\mathcal{G}(\mathcal{A})$ (see Proposition 3.13):

Proposition 7.9. *Let $\mathcal{A}$ be an invertible automaton. Suppose that for each pair $w, w' \in \mathsf{X}$ there exists $g \in \mathcal{G}(\mathcal{A})$ such that $g(w) = w'$ and $g|_w = e$. Suppose further that for every $q \in Q$ there is $q' \in Q$ and $w \in \mathsf{X}$ such that $\tau(q', w) = q$. Then $\mathcal{G}(\mathcal{A})$ is self-replicating.*

Proof. Follows from Lemma 7.8 and Corollary 3.17.

In other words, a sufficient condition for an automaton group to be self-replicating is that it must contain sufficient rigid automorphisms, and every state of the automaton must have an incoming transition from another (not necessarily different) state.

EXAMPLE 7.10. The odometer group generated by the automaton shown in Figure 7.2 is self-replicating. Indeed, $a(0) = 1$ and $a|_0 = b = e$ and hence $a^{-1}(1) = 0$ and $a^{-1}|_1 = a|_0 = e$, so the group contains the required rigid automorphisms. Finally $\tau(a, 0) = b$ and $\tau(a, 1) = a$ so Proposition 7.9 applies.

We can now proceed to our construction.

7.3. The automata $\mathcal{A}_G$

Let G be a subgroup of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n$. We would like to define a finite automaton $\mathcal{A}_G$ such that the closure of the corresponding automaton group $\overline{\mathcal{G}(\mathcal{A}_G)}$ has property $\mathcal{R}$ and agrees with G on T_n, i.e. $\mathcal{G}(\mathcal{A}_G)_{[n]} = G$. We will say that such an automaton *extends G to T*.

It is necessary to take the closure of this group in $\mathsf{Aut}(T)$ because $\mathcal{G}(\mathcal{A}_G)$ is finitely generated when $\mathcal{A}_G$ is finite and, as we saw at the end of Chapter 5, no finitely generated group can have property $\mathcal{R}$. We must therefore make do with *topologically* finitely generated groups.

The hard part will be to ensure that $\mathcal{G}(\mathcal{A}_G)$ is self-replicating. To this end, following Proposition 3.10 we may assume without loss of generality that G has sufficient rigid automorphisms; that is, for every pair of vertices $v, v' \in \mathcal{L}_1$ there exists $g \in G$ such that $g(v) = v'$ and $g|_v = e$. For convenience, let $H = \varphi_n(G)$ be the restriction of G to T_{n-1}, so that G satisfies $\mathcal{R}_n(H)$. As above, we take the alphabet to be $\mathsf{X} = \{0, 1, \ldots, p-1\}$ where p is the degree of T.

We must ask first of all what necessary conditions $\mathcal{A}_G$ must satisfy. If $Q = \{q_1, \ldots, q_r\}$ then in order to ensure that $\mathcal{G}(\mathcal{A}_G)_{[n]} = G$ we must have $\langle (q_1)_{[n]}, \ldots, (q_r)_{[n]} \rangle = G$ (where we are using q_i to represent both the state and its corresponding automorphism). The only other condition is that $\mathcal{G}(\mathcal{A}_G)$ be self-replicating, which can be met by satisfying the hypotheses of Proposition 7.9. This requires our assumption about the existence of rigid automorphisms in G, and also an appropriate choice of generating set Q for G. Let us see how this works in a couple of examples.

EXAMPLE 7.11. Let G be the subgroup of $\mathsf{Aut}(T_n)$ corresponding to the action of $\mathbb{Z}_2$ on the binary tree. Then G is cyclic of order 2^n, and as we have seen it is equal to the odometer group restricted to T_n. This means we already have a self-replicating automaton which extends G to T, namely the one in Figure 7.2.

EXAMPLE 7.12. The (closure of the) first Grigorchuk group Γ, whose many interesting properties are described in [**Gri05**], acts on the binary tree and is equal to $\mathcal{M}^\infty(G)$ for a certain subgroup G of $\mathsf{Aut}(T_4)$ of order 2^{12} that has property $\mathcal{R}_4(\mathsf{Aut}(T_3))$. There is a well-known automaton which generates Γ and hence extends G to T. It is shown in Figure 7.3.

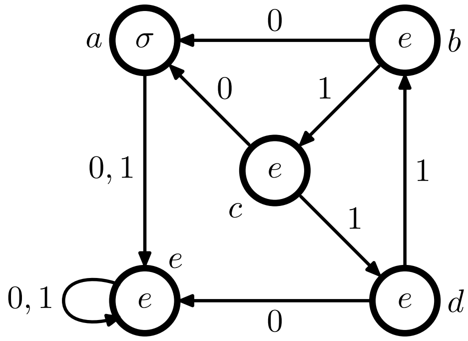

$$a = \sigma(e, e),\ e = (e, e)$$
$$b = (a, c),\ c = (a, d),\ d = (e, b)$$

FIGURE 7.3. The automaton generating the first Grigorchuk group Γ.

This example demonstrates that finite automata are not guaranteed to produce non-maximal groups. In fact, in this case, it is conjectured (see the end of Section 5.4) that $\overline{\Gamma}$ is the *only* subgroup of $\mathsf{Aut}(T)$ with property $\mathcal{R}$ that agrees with G on T_4. If this conjecture is true then *all* possible automata $\mathcal{A}_G$ from Theorem 7.13 would have $\overline{\mathcal{G}(\mathcal{A}_G)} = \overline{\Gamma}$.

It has been verified with a computer search that any such group must agree with Γ on T_8, so if there exists an automaton $\mathcal{A}_G$ such that $\overline{\mathcal{G}(\mathcal{A}_G)} \neq \overline{\Gamma}$ then the difference is not visible until at least level 9. This seems highly unlikely.

Of course, it is not enough to describe a few examples. We must provide a general construction which takes an arbitrary subgroup $G \leq \mathsf{Aut}(T_n)$ satisfying $\mathcal{R}_n$ and produces an automaton $\mathcal{A}_G$ which extends G to T. As usual, because of Proposition 3.10, we will assume that G has sufficient rigid automorphisms.

THEOREM 7.13. *Suppose that $G \leq \mathsf{Aut}(T_n)$ has property $\mathcal{R}_n$ and has sufficient rigid automorphisms. Then there exists a finite automaton $\mathcal{A}_G$ such that the group $\overline{\mathcal{G}(\mathcal{A}_G)}$ has property $\mathcal{R}$ and agrees with G on T_n.*

PROOF. We can define $\mathcal{A}_G$ as follows. We define the set of states Q to be the elements of G. Note that G is finite so $\mathcal{A}_G$ will be finite as well. The output function $\lambda : G \times \mathsf{X} \to \mathsf{X}$ is defined by

$$\lambda(g, w) = g(w) \text{ for all } g \in G \text{ and } w \in \mathsf{X}.$$

The notation $g(w)$ makes sense because we identify X with level 1 of T. The transition function $\tau : G \times \mathsf{X} \to G$ presents a difficulty. In light of Proposition 7.5 we would like to define something like $\tau(g, w) = g|_w$. However, this will not work because $g|_w$ belongs to $\mathsf{Aut}(T_{n-1})$ (in fact it belongs to H, by Proposition 3.11) and not to G itself. Instead, we need to choose an automorphism $g' \in G$ such that $\varphi_n(g') = g|_w$ and then set $\tau(g, w) = g'$. Such a g' is guaranteed to exist because G has property $\mathcal{R}_n(H)$ and $g|_w \in H$.

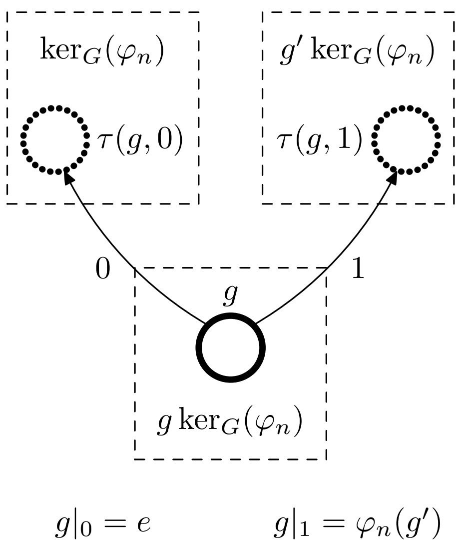

FIGURE 7.4. Choosing transitions: $\tau(g, w)$ can be any $g' \in G$ such that $\varphi_n(g') = g|_w$. The set of such g' is a coset of $\ker_G(\varphi_n)$ in G.

The problem with this definition of τ is that there could be more than one choice for g'. Indeed, the set of possible g' is a coset of $\ker(\varphi_n)$ in G which by Proposition 3.19(a) has cardinality $\frac{|G|}{|H|} \geq p$ (see Figure 7.4).

Unfortunately, there really is no way in general to systematically choose g'. This means that $\mathcal{A}_G$ will be one of a number of possible automata rather than a single well-defined automaton. This freedom actually turns out to be extemely useful, as we can use it to place further restrictions on the definition of τ and guarantee that the automaton generates a self-replicating group.

Firstly, to take advantage of Proposition 7.9, we need to ensure that sufficient rigid automorphisms exist in the automaton group. This is why we assumed above that these automorphisms exist in G. Let $v, w \in \mathsf{X}$ and suppose that $g \in G$ is (v, w)-rigid, i.e. $g(v) = w$ and $g|_v = e$. The above argument implies that $\tau(g, v)$ can be any $g' \in G$ such that $\varphi_n(g') = e$, which means $g' \in \ker_G(\varphi_n)$. In this case however, we need to set $\tau(g, v) = e$ so that the corresponding automorphism $g \in \mathcal{G}(\mathcal{A}_G)$ is (v, w)-rigid as well. Doing this for all rigid automorphisms in G guarantees that $\mathcal{G}(\mathcal{A}_G)$ contains the required rigid automorphisms, thereby satisfying the first hypothesis of Proposition 7.9.

Now, since we need $\tau(g, v) = e$ for some g and v, there needs to be a state representing e in $\mathcal{A}_G$. The state q corresponding to $e \in G$ need not be trivial since we have only specified that $\tau(q, w) \in \ker_G(\varphi_n)$ for each $w \in \mathsf{X}$; this guarantees that q is trivial on level n but not necessarily on all of T. The simplest solution is to set $\tau(q, w) = q$ for all $w \in \mathsf{X}$, for then we have the decomposition $q = (q, \ldots, q)$ which means that q is trivial. In terms of the graph, this state will look like the one in Figure 7.5.

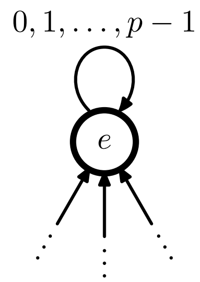

FIGURE 7.5. The trivial state in $\mathcal{A}_G$.

In order to satisfy the second hypothesis of Proposition 7.9 we need to ensure that for every $g \in G$ there exists $h \in G$ and $w \in \mathsf{X}$ such that $\tau(h, w) = g$, i.e. every state should have at least one incoming transition.

For convenience let $k = \frac{|G|}{|H|} = |\ker_G(\varphi_n)|$ and recall from Proposition 3.19(a) that $k \geq p$ (in fact p divides k). Now there are, in total, $p\,|G|$ edges in the graph of $\mathcal{A}_G$ — one outgoing transition $\tau(g, w)$ from each state $g \in G$ for each $w \in \mathsf{X}$ — and there are k choices for each transition, within the appropriate coset of $\ker_G(\varphi_n)$. There are $|H|$ such cosets (by the first isomorphism theorem, since $\varphi_n(G) = H$) so we expect, on average, $\frac{p|G|}{|H|} = pk$ incoming transitions into each $\ker_G(\varphi_n)$-coset. It turns out that this is not just the average, it is the *exact* number:

LEMMA 7.14. *Let G, H and k be as above. Then for each $h \in H$ there exist exactly pk pairs (g, w) where $g \in G$ and $w \in \mathsf{X}$, such that $g|_w = h$.*

PROOF. Fix $h \in H$. For each $w, w' \in \mathsf{X}$ define

$$S_{w,w'} = \{g \in G : g(w) = w' \text{ and } g|_w = h\}.$$

Observe that, for a fixed w, $S_{w,w'} \cap S_{w,w''} = \varnothing$ when $w' \neq w''$. We claim that $\left|S_{w,w'}\right| = |S_{w,w}|$ for all $w, w' \in \mathsf{X}$. To see this, fix w and w' and let x be a (w, w')-rigid automorphism in G; that is, $x(w) = w'$ and $x|_w = e$. Such an x exists by our assumptions on G. Then for each $g \in S_{w,w}$ we have $(xg)(w) = x(w) = w'$ and

$$(xg)|_w = x|_{g(w)} g|_w = x|_w h = h$$

so $xg \in S_{w,w'}$. Conversely, by a similar caluclation, if $g \in S_{w,w'}$ then $x^{-1}g \in S_{w,w}$. Thus, left multiplication by x is a bijection from $S_{w,w}$ to $S_{w,w'}$ and the claim is established.

The next claim is that $|S_{w,w}| = |S_{0,0}|$ for all $w \in \mathsf{X}$. To prove this, fix $w \in \mathsf{X}$ and let y be a $(0, w)$-rigid automorphism in G; that is, $y(0) = w$ and $y|_0 = e$. Such a y exists by our assumptions on G. Note that y^{-1} is $(w, 0)$-rigid; that is, $y^{-1}(w) = 0$ and $y^{-1}|_w = e$. Let $g \in S_{w,w}$. Then $(y^{-1}gy)(0) = (y^{-1}g)(w) = y^{-1}(w) = 0$ and

$$(y^{-1}gy)|_0 = (y^{-1}g)|_{y(0)} y|_0 = (y^{-1}g)|_w = y^{-1}|_{g(w)} g|_w = y^{-1}|_w h = h$$

so $y^{-1}gy \in S_{0,0}$. Similarly if $g \in S_{0,0}$ then $ygy^{-1} \in S_{w,w}$ so conjugation by y is the required bijection between $S_{w,w}$ and $S_{0,0}$.

We now calculate $|S_{0,0}|$. Since $S_{0,0} \subseteq G_0$ we may write

$$S_{0,0} = \{g \in G_0 : g|_0 = h\} = \{g \in G_0 : \psi_n(g) = h\}.$$

Therefore $S_{0,0}$ is a coset of $\ker_{G_0}(\psi_n)$. Then by Proposition 3.19(a),

$$|S_{0,0}| = |\ker_{G_0}(\psi_n)| = \frac{1}{p} |\ker_G(\varphi_n)| = \frac{k}{p}.$$

Now for each $w \in \mathsf{X}$ define

$$T_w = \{g \in G : g|_w = h\}.$$

Note that $T_w = \bigcup_{w' \in \mathsf{X}} S_{w,w'}$ and this is a disjoint union. Since each $\left|S_{w,w'}\right| = |S_{0,0}|$, we obtain

$$|T_w| = |\mathsf{X}|\,|S_{0,0}| = p\frac{k}{p} = k.$$

It is important to note that the sets T_w are not necessarily disjoint, since a single $g \in G$ could have $g|_w = h = g|_{w'}$ for some $w \neq w'$.

Finally, the set of pairs (g, w) in the statement of the Lemma is equal to

$$U = \{(g, w) \in G \times \mathsf{X} : g \in T_w\} = \bigcup_{w \in \mathsf{X}} T_w \times \{w\}$$

and this is now a disjoint union due to the pairing of T_w with w. Since $|T_w \times \{w\}| = |T_w| = k$ for all $w \in \mathsf{X}$, we have $|U| = |\mathsf{X}|\,k = pk$ as required. ☠

Returning to the proof of the Theorem, observe that in order to make the construction work, all that matters is that for each state q at least one of the pk incoming transitions into the coset $q\ker_G(\varphi_n)$ is assigned to q. Since there are only k states in each coset, this yields a large number of possibilities (see Example 7.15). The important thing is that it can be done.

Let us clarify one remaining point. The existence of rigid automorphisms makes the situation more complicated in the case of the coset $\ker_G(\varphi_n)$ itself. There are two potential problems: firstly, some of the pk incoming transitions are $\tau(g, w)$ for an automorphism g which is rigid at w, so these must be equal to the identity. Secondly, for the identity itself, we must have $\tau(e, w) = e$ for all $w \in \mathsf{X}$. We need to check that there are enough incoming transitions left over so that each state $q \in \ker_G(\varphi_n)$ can be equal to at least one of them. It follows from Lemma 3.8 that we really only need automorphisms that are $(v, v+1)$-rigid for $v = 0, 1, \ldots, p-2$. This means we only need to guarantee that $p-1$ of the incoming transitions from rigid automorphisms are equal to e. Taking into account the p transitions from e to itself, there are $pk - 2p + 1$ incoming transitions left over to assign to the remaining $k-1$ nontrivial states in $\ker_G(\varphi_n)$. This suffices, since $k \geq p \geq 2$ and

$$pk - 2p + 1 = k - 1 + (p-1)(k-2) \geq k - 1. \tag{7.1}$$

Thus, Proposition 7.9 guarantees that the group $\mathcal{G}(\mathcal{A}_G)$ will be self-replicating. It is clear from the construction that the restriction of $\mathcal{G}(\mathcal{A}_G)$ to T_n is equal to G which means it is transitive on $\mathcal{L}_n$. By Proposition 2.14 the closure $\overline{\mathcal{G}(\mathcal{A}_G)}$ is transitive on ∂T and therefore has property $\mathcal{R}$ as required. ☠

EXAMPLE 7.15. Let G be the subgroup of $\mathsf{Aut}(T_2)$ corresponding to the action of $\mathbb{Z}_2$ on the binary tree, as in Example 7.11 with $n = 2$. Let us follow Theorem 7.13 to construct all possible $\mathcal{A}_G$.

First, some notation: as before, let $p = 2$, $\mathsf{X} = \{0, 1\}$ and let σ denote the transposition $(0\ 1)$. Define $a \in \mathsf{Aut}(T)$ by the recursive decomposition $a = \sigma(e, a)$ and let a_n be the restriction of a to T_n for each n. Note that $a_1 = \sigma$ and $a_n = \sigma(e, a_{n-1})$ for $n \geq 2$. We saw in Example 4.1 that a generates the odometer group and therefore a_n has order 2^n and generates the restriction of this group to T_n. In particular, $G = \langle a_2 \rangle = \{e, a_2, a_2^2, a_2^3\}$. For convenience let $b = a_2$. We have the decompositions $b = \sigma(e, \sigma)$, $b^2 = (\sigma, \sigma)$ and $b^3 = b^{-1} = \sigma(\sigma, e)$.

Following the construction specified above, let $Q = G$. We have $\lambda_q = \sigma$ for $q = b$ and $q = b^3$, and $\lambda_q = e$ for $q = e$ and $q = b^2$. We also have $\tau(e, 0) = \tau(e, 1) = e$. Since $p = 2$ we only need $p - 1 = 1$ rigid automorphism in our group. Now b is $(0, 1)$-rigid in G, as can be seen from the decomposition $b = \sigma(e, \sigma)$. We insist, therefore, that $\tau(b, 0) = e$. We are now free to choose the other transitions, subject to the condition that each $q \in Q$ is equal to at least one of them.

In how many ways can this be done? Note that there are 5 transitions still to choose, namely $\tau(b, 1)$, $\tau(b^2, 0)$, $\tau(b^2, 1)$, $\tau(b^3, 0)$ and $\tau(b^3, 1)$. The subgroup $K = \ker_G(\varphi_2)$ has order 2; it is equal to $\{e, b^2\}$, since $\varphi_2(e) = \varphi_2(b^2) = e$, and the only other coset is $bK = \{b, b^3\}$, with $\varphi_2(b) = \varphi_2(b^3) = \sigma$. From the decompositions we can see that the transition $\tau(b^3, 1)$ must belong to K (since $b^3|_1 = e$) and the other four must all belong to bK (since $b|_1 = b^2|_0 = b^2|_1 = b^3|_0 = \sigma$). Since the transitions $\tau(b, 0)$, $\tau(e, 0)$ and $\tau(e, 1)$ have already been assigned to the identity, we are forced to choose $\tau(b^3, 1) = b^2$; note that equality holds in (7.1) in this case because $p = k = 2$. The other four transitions can be chosen freely from bK, provided each of b and b^3 is chosen at least once. This yields a total of $2^4 - 2 = 14$ possible $\mathcal{A}_G$.

Some of these automata are shown in Figure 7.6. Moving in a clockwise direction from the top left in each diagram are the states corresponding to e, b, b^3 and b^2 respectively. Note that these labels are purely formal; there is no guarantee that the automorphism of T corresponding to b^2 will be the square of the automorphism of T corresponding to b. The construction of $\mathcal{A}_G$ only ensures that this relationship holds on T_2, so that we get $\mathcal{G}(\mathcal{A}_G)_{[2]} = G$.

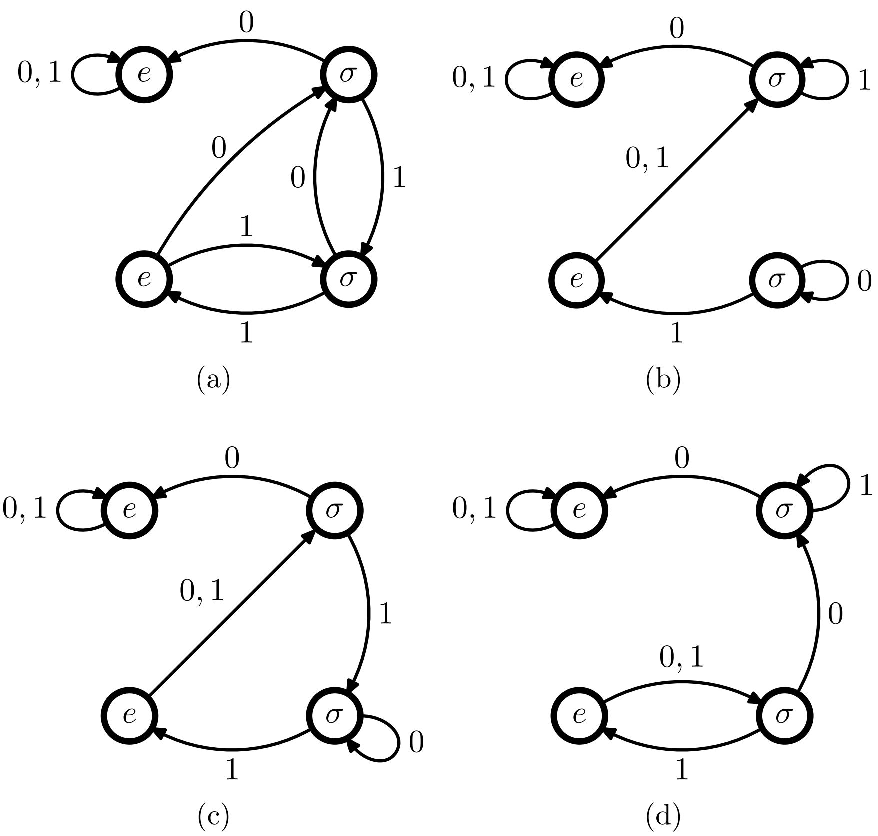

FIGURE 7.6. Some possible $\mathcal{A}_G$.

Recall from Section 5.2 that there are two subgroups of $\mathsf{Aut}(T_3)$ satisfying $\mathcal{R}_3(G)$, up to conjugacy, namely $\mathcal{M}(G)$ of order 2^4 and the level 3 odometer group which is cyclic of order 2^3. Both of these groups appears as $\mathcal{G}(\mathcal{A}_G)_{[3]}$ for some $\mathcal{A}_G$ in Figure 7.6. Automata (a) and (b) both generate $\mathcal{M}(G)$ on T_3, although they generate different subgroups of $\mathsf{Aut}(T_4)$; (a) generates $\mathcal{M}^2(G)$ of order 2^8 and (b) generates an index 2 subgroup of $\mathcal{M}^2(G)$. Automaton (d) generates the odometer on T_3 (but not on T_4 where it is maximal of order 2^5) and (c) generates a group which is conjugate to the odometer on T_3 (and is also conjugate to the odometer on T_4 and T_5, but not on T_6 where it has order 2^7).

Finally, it is known from [**Šun11**, Proposition 2] that the maximal group $\mathcal{M}^\infty(G)$ in this example is not topologically finitely generated, so each of the 14 possible $\mathcal{A}_G$ must generate a group with property $\mathcal{R}$ distinct from $\mathcal{M}^\infty(G)$.

EXAMPLE 7.16. Let Γ be the Grigorchuk group and let $G = \Gamma_{[4]}$, as in Example 7.12. The automaton in that example had only 5 states, whereas each of the $\mathcal{A}_G$ constructed in the proof of Theorem 7.13 would have $|G| = 2^{12} = 4096$ states. Note that $K = \ker_G(\varphi_4)$ has index 2^7 in G since $G/K \cong \varphi_4(G) = \Gamma_{[3]} = \mathsf{Aut}(T_3)$, so the states of $\mathcal{A}_G$ are partitioned into the 2^7 cosets of K, each containing $|K| = 2^5$ states. This means that the number of possible $\mathcal{A}_G$ is vast, in the vicinity of $|K|^{2|G|-3} = 2^{40945}$. Obviously it would not be practical to describe them all explicitly!

As mentioned in Example 7.12, it is conjectured that *all* of these automata generate $\mathcal{M}^\infty(G) = \overline{\Gamma}$. This situation stands in stark contrast to Example 7.15 where *none* of the $\mathcal{A}_G$ generate $\mathcal{M}^\infty(G)$.

In general, except for very small groups, the construction of $\mathcal{A}_G$ using $Q = G$ is rather unwieldy. Its purpose is simply to show that an automaton extending G to T always exists. In practice one can usually find a much smaller automaton to do the job, as we saw in Examples 7.11 and 7.12. It would be worthwhile, as a future project, to modify the construction in the proof of Theorem 7.13 and turn it into a more practical algorithm for generating $\mathcal{A}_G$ with a relatively small number of states.

APPENDIX A

Algorithm to find groups with property $\mathcal{R}_n$

This appendix describes the algorithm that was used to calculate all the possible groups with property $\mathcal{R}_n(G)$ for a given group G with property $\mathcal{R}_{n-1}$. The diagrams in Chapter 5 are based on the data obtained using an implementation of this algorithm in `MAGMA`.

The algorithm begins with the maximal group $\mathcal{M}(G)$, which can be generated in a variety of ways using Definition 6.1. For example, one could keep choosing random elements of $\mathcal{M}(G)$ until they generate a group of the size given by Theorem 6.2(c). Having calculated $\mathcal{M}(G)$, the algorithm proceeds by generating all the maximal subgroups of $\mathcal{M}(G)$, and discarding those that do not have property $\mathcal{R}_n(G)$. Then, for each remaining subgroup H, the algorithm calculates the maximal subgroups of H, again discarding those that do not have property $\mathcal{R}_n(G)$. Conjugates are removed at each step. This process continues until no more subgroups with property $\mathcal{R}_n(G)$ are found. Proposition 3.24 guarantees that the algorithm will find, up to conjugacy, all subgroups of $\mathcal{M}(G)$ with property $\mathcal{R}_n(G)$.

The set of maximal subgroups of a group G can be calculated in `MAGMA` with the command `Subgroups(G : Al:="Maximal")`. The built-in algorithms for calculating maximal subgroups tend to be much more efficient than those for generic subgroups, which is why the algorithm was designed in this way.

A technicality should be noted here — the `Subgroups` command in `MAGMA` actually returns a set of conjugacy class representatives for the subgroups it finds. One may therefore ask: what happens if there is a subgroup H with property $\mathcal{R}_n(G)$ but the command picks a conjugate of H that does not have property $\mathcal{R}_n(G)$?

This situation cannot occur when G is a p-group because all of its maximal subgroups are normal (by Lemma 3.26), and thus there will only be one representative in each conjugacy class. Note that the conjugate removal step would still still necessary in this case because subgroups that are not conjugate in G can still be conjugate in $\mathsf{Aut}(T_n)$. Most importantly, Proposition 3.23 ensures that no generality is lost by restricting to p-groups for $p = 2$, which was the focus of the calculations in Chapter 5.

The algorithm still works for non-p-groups, provided *all* maximal subgroups are checked and not just a set of conjugacy class representatives. The efficiency of the algorithm would suffer as a result, particularly because conjugacy testing can be computationally expensive in large groups.

Algorithm A.1 Finding all conjugacy classes of groups with property $\mathcal{R}_n(G)$ for a given group G.

Input : A subgroup G of $\mathsf{Aut}(T_{n-1})$ with property $\mathcal{R}_{n-1}$
Output: A full set of conjugacy class representatives for the subgroups of $\mathsf{Aut}(T_n)$ with property $\mathcal{R}_n(G)$

```
allGroups ← {M(G)};
currentLayer ← {M(G)};
while currentLayer is nonempty do
    newGroups ← {};
    for currentGroup in currentLayer do
        subgroups ← MaximalSubgroups (currentGroup);
        for subgroup in subgroups do
            if subgroup does not have property R_n(G) then
                remove subgroup from subgroups;
            end
        end
        newGroups ← newGroups ∪ subgroups;
    end
    RemoveConjugates (newGroups);
    currentLayer ← newGroups;
    allGroups ← allGroups ∪ newGroups;
end
return allGroups;
```

APPENDIX B

Groups with property $\mathcal{R}_n$ for $p = 2$ and $n \leq 4$

In this appendix we fix $p = 2$ so T is the binary tree. From Chapter 5, in particular Proposition 5.1, we know that the only groups with property $\mathcal{R}_2$ are the following:

- $\mathfrak{Z} = \langle \sigma(e, \sigma) \rangle = \{e, (\sigma, \sigma), \sigma(e, \sigma), \sigma(\sigma, e)\}$;
- $\mathfrak{L} = \langle \sigma(e, e), \sigma(\sigma, \sigma) \rangle = \{e, (\sigma, \sigma), \sigma(e, e), \sigma(\sigma, \sigma)\}$;
- $\mathsf{Aut}(T_2)$.

Tables B.1, B.2 and B.3 list one representative from each conjugacy class of groups with property $\mathcal{R}_3$, following Proposition 3.10 to choose representatives that contain sufficient rigid automorphisms. For clarity we split the groups into the three cases enumerated above, according to $G_{[2]}$. For each group G with property $\mathcal{R}_3$ we also state the number of conjugacy classes of groups with property $\mathcal{R}_4(G)$ (in the $|\mathcal{E}_4(G)|$ column).

Since $p = 2$, by Proposition 3.23 all these groups are 2-groups and hence $|G|$ is a power of 2. For simplicity we give $\log_2 |G|$ rather than $|G|$ in the tables below. Recall that $\sharp G$ means the (minimal) size of a generating set for G.

We use σ to denote the element of $\mathsf{Aut}(T_n)$ (where n is clear from context) that transposes the vertices 0 and 1 and $\sigma|_0 = \sigma|_1 = e$.

$\log_2 \lvert G\rvert$	$\sharp G$	generators	$\lvert\mathcal{E}_4(G)\rvert$	notes
4	2	$\sigma(e, \sigma(e, \sigma))$, $(e, (\sigma, \sigma))$	4	$\mathcal{M}(\mathfrak{Z})$
3	1	$\sigma(e, \sigma(e, \sigma))$	2	$\mathfrak{Z}_3$

TABLE B.1. Groups below $\mathfrak{Z}$

$\log_2 \lvert G\rvert$	$\sharp G$	generators	$\lvert\mathcal{E}_4(G)\rvert$	notes
4	3	$\sigma,\ (\sigma,\sigma),\ (e,(\sigma,\sigma))$	18	$\mathcal{M}(\mathfrak{L})$
3	2	$\sigma,\ (\sigma,\sigma(\sigma,\sigma))$	1	
3	2	$\sigma(e,(\sigma,\sigma)),\ (\sigma,\sigma)$	3	abelian
3	3	$\sigma,\ (\sigma,\sigma),\ ((\sigma,\sigma),(\sigma,\sigma))$	5	$\mathfrak{L}_3$

TABLE B.2. Groups below $\mathfrak{L}$

$\log_2 \lvert G\rvert$	$\sharp G$	generators	$\lvert\mathcal{E}_4(G)\rvert$	notes
7	3	$\sigma,\ (e,\sigma),\ (e,(e,\sigma))$	29	$\mathsf{Aut}(T_3)$
6	3	$\sigma,\ (\sigma,\sigma),\ (e,\sigma(e,\sigma))$	9	$\langle\mathfrak{Z}_3,\mathfrak{L}_3\rangle$
6	3	$\sigma,\ (e,\sigma),\ ((e,\sigma),(e,\sigma))$	24	
6	2	$\sigma(e,(e,\sigma)),\ (e,\sigma(e,\sigma))$	1	
6	2	$\sigma(e,(e,\sigma)),\ (e,\sigma)$	1	Basilica
5	3	$\sigma(e,\sigma(e,\sigma)),\ ((e,\sigma),\sigma),\ (e,(\sigma,\sigma))$	6	
5	2	$\sigma,\ ((e,\sigma),\sigma(e,\sigma))$	11	Lamplighter
4	2	$\sigma(e,\sigma(e,\sigma)),\ ((e,\sigma),\sigma)$	1	
4	2	$\sigma(e,\sigma(e,\sigma)),\ ((\sigma,e),\sigma)$	3	D_∞

TABLE B.3. Groups below $\mathsf{Aut}(T_2)$

Bibliography

[BGN03] L. Bartholdi, R. Grigorchuk, and V. Nekrashevych, *From fractal groups to fractal sets*, Springer, 2003.

[BS97] Y. Barnea and A. Shalev, *Hausdorff dimension, pro-p groups, and Kac-Moody algebras*, T. Am. Math. Soc. **349** (1997), no. 12, 5073–5091.

[BW04] U. Baumgartner and G. Willis, *Contraction groups and scales of automorphisms of totally disconnected locally compact groups*, Israel J. Math. **142** (2004), 221–248.

[DD89] W. Dicks and M. Dunwoody, *Groups acting on graphs*, Cambridge University Press, 1989.

[DDMN10] D. D'Angeli, A. Donno, M. Matter, and T. Nagnibeda, *Schreier graphs of the basilica group*, J. Modern Dynamics **4** (2010), no. 1, 167–205.

[Gle51] A. Gleason, *The structure of locally compact groups*, Duke Math. J. **18** (1951), no. 1, 85–104.

[Glö98] H. Glöckner, *Scale functions on p-adic lie groups*, Manuscripta Mathematica **97** (1998), no. 2, 205–215.

[Glö06] ———, *Locally compact groups built up from p-adic lie groups, for p in a given set of primes*, J. Group Theory **9** (2006), no. 4, 427–454.

[GNS00] R. Grigorchuk, V. Nekrashevich, and V. Sushchanskii, *Automata, dynamical systems and groups*, Proc. Steklov Inst. Math. **231** (2000), no. 4, 128–203.

[Gri05] R. Grigorchuk, *Solved and unsolved problems around one group*, Infinite groups: geometric, combinatorial and dynamical aspects, Springer, 2005, pp. 117–218.

[GŻ01] R. Grigorchuk and A. Żuk, *The lamplighter group as a group generated by a 2-state automaton, and its spectrum*, Geom. Dedicata **87** (2001), no. 1-3, 209–244.

[GŻ02] ———, *On a torsion-free weakly branch group defined by a three state automaton*, Int. J. Algebr. Comput. **12** (2002), no. 01n02, 223–246.

[HNN49] G. Higman, B. H. Neumann, and H. Neumann, *Embedding theorems for groups*, J. London Math. Soc. **s1-24** (1949), no. 4, 247–254.

[HY61] J. Hocking and G. Young, *Topology*, Addison-Wesley, 1961.

[Mun00] J. Munkres, *Topology*, Prentice Hall, 2000.

[OOR04] R. Orellana, M. Orrison, and D. Rockmore, *Rooted trees and iterated wreath products of cyclic groups*, Adv. in Appl. Math. **33** (2004), 531–547.

[PŠ] A. Penland and Z. Šunić, *Finitely constrained groups of maximal Hausdorff dimension*, (preprint) arXiv:1406.0522.

[RZ10] L. Ribes and P. Zalesskii, *Profinite groups*, Springer, 2010.

[Ser80] J-P Serre, *Trees*, Springer, 1980.

[Šun07] Z. Šunić, *Hausdorff dimension in a family of self-similar groups*, Geom. Dedicata **124** (2007), no. 1, 213–236.

[Šun11] ———, *Pattern closure of groups of tree automorphisms*, Bull. Math. Sci. **1** (2011), no. 1, 115–127.

[Tit70] J. Tits, *Sur le groupe des automorphismes d'un arbre*, Essays on topology and related topics (Mémoires dédiés à Georges de Rham), Springer, 1970, pp. 188–211.

[vD36] D. van Dantzig, *Zur topologischen Algebra III. Brouwersche und Cantorsche Gruppen*, Compositio Math. **3** (1936), 408–426.

[Wil94] G. Willis, *The structure of totally disconnected locally compact groups*, Math. Ann. **300** (1994), 341–363.

[Wil01] ———, *Further properties of the scale function on a totally disconnected locally compact group*, J. Algebra **237** (2001), 142–164.

[Wil04a] ———, *A canonical form for automorphisms of totally disconnected locally compact groups*, Random walks and geometry, 2004, pp. 295–316.

[Wil04b] ———, *Tidy subgroups for commuting automorphisms of totally disconnected locally compact groups: An analogue of simultaneous triangularisation of matrices*, New York J. Math. **10** (2004), 1–35.

[Yam53] H. Yamabe, *A generalization of a theorem of Gleason*, Annals of Math. **58** (1953), no. 2, 351–365.