The Automorphism Tower Problem

Simon Thomas

Rutgers University, New Brunswick, New Jersey 08903

E-mail address: `sthomas@math.rutgers.edu`

Contents

Preface

This is a preliminary version of an unpublished book on the automorphism tower problem, which was intended to be intelligible to beginning graduate students in both logic and algebra.

Simon Thomas

Introduction

If G is a centreless group, then there is a natural embedding of G into its automorphism group $\mathrm{Aut}(G)$, obtained by sending each $g \in G$ to the corresponding inner automorphism $i_g \in \mathrm{Aut}(G)$. It is easily shown that the group $\mathrm{Inn}(G)$ of inner automorphisms is a normal subgroup of $\mathrm{Aut}(G)$ and that $C_{\mathrm{Aut}(G)}(\mathrm{Inn}(G)) = 1$. In particular, $\mathrm{Aut}(G)$ is also a centreless group. This enables us to define the automorphism tower of G to be the ascending chain of centreless groups

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \dots G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \dots$$

such that for each ordinal α

(a) $G_{\alpha+1} = \mathrm{Aut}(G_\alpha)$; and

(b) if α is a limit ordinal, then $G_\alpha = \bigcup_{\beta<\alpha} G_\beta$.

(At each successor step, we identify G_α with $\mathrm{Inn}(G_\alpha)$ via the natural embedding.)

The automorphism tower is said to terminate if there exists an ordinal α such that $G_{\alpha+1} = G_\alpha$. Of course, this occurs if and only if there exists an ordinal α such that $\mathrm{Aut}(G_\alpha) = \mathrm{Inn}(G_\alpha)$. In this case, the height $\tau(G)$ of the automorphism tower is defined to be the least ordinal α such that $G_{\alpha+1} = G_\alpha$. In 1939, Wielandt proved that if G is a finite centreless group, then its automorphism tower terminates after finitely many steps. However, there exist natural examples of infinite centreless groups whose automorphism towers do not terminate in finitely many steps. For example, it can be shown that the automorphism tower of the infinite dihedral group D_∞ terminates after exactly $\omega+1$ steps. The classical version of the automorphism tower problem asks whether the automorphism tower of an arbitrary centreless group G eventually terminates, perhaps after a transfinite number of steps.

In the 1970s, a number of special cases of the automorphism tower problem were solved. For example, Rae and Roseblade proved that the automorphism tower of a centreless Černikov group terminates after finitely many steps and Hulse proved that the automorphism tower of a centreless polycyclic group terminates after countably many steps. In each of these special cases, the proof depended upon a detailed understanding of the groups G_α occurring in the automorphism towers of the relevant groups G. But the problem was not solved in full generality until 1984, when I proved that the automorphism tower of an arbitrary centreless group G terminates after at most $\left(2^{|G|}\right)^+$ steps. The proof is extremely simple and uses only the most basic results on automorphism towers, together with some elementary properties of the infinite cardinal numbers. Of course, this still leaves open the problem of determining the best possible upper bound for the height $\tau(G)$ of the automorphism tower of an infinite centreless group G. For example, it is natural to ask whether there exists a fixed cardinal κ such that $\tau(G) \leq \kappa$ for every infinite centreless group G. This question turns out to be relatively straightforward. For each ordinal α, it is possible to construct an infinite centreless group G such that $\tau(G) = \alpha$. However, if we ask whether it is true that $\tau(G) \leq 2^{|G|}$ for every infinite centreless group G, then matters become much more interesting. For example, it is independent of the classical ZFC axioms of set theory whether $\tau(G) \leq 2^{\aleph_1}$ for every centreless group G of cardinality $\aleph_1$; i.e. this statement can neither be proved nor disproved using ZFC.

This book presents a self-contained account of the automorphism tower problem, which is intended to be intelligible to beginning graduate students in both logic and algebra. There are essentially no set-theoretic prerequisites. The only requirement is a basic familiarity with some of the fundamental notions of algebra. The first half of the book presents those results which can be proved using ZFC; and also includes an account of the necessary set-theoretic background, such as the notions of a regular cardinal and a stationary set. This is followed by a short introduction to set-theoretic forcing, which is aimed primarily at algebraists. The final three chapters explain why a number of natural problems concerning automorphism towers are independent of ZFC.

In more detail, the book is organised as follows. In Chapter 1, we introduce the notion of the automorphism tower of a centreless group and illustrate this notion by computing the automorphism towers of a number of groups. In Chapter 2, we present a proof of Wielandt's theorem that the automorphism tower of a finite centreless group terminates after finitely many steps. In Chapter 3, we prove that the automorphism tower of an infinite centreless group G terminates after strictly less than $\left(2^{|G|}\right)^+$ steps. This chapter also contains an account of the basic theory of regular cardinals and stationary sets. Much of the remainder of this book is concerned with the problem of constructing centreless groups with extremely long automorphism towers. Unfortunately it is usually very difficult to compute the successive groups in the automorphism tower of a centreless group. In Chapter 4, we introduce the normaliser tower technique, which enables us to almost entirely bypass this problem. Instead, throughout most of this book, we only have to deal with the much easier problem of computing the successive normalisers of a subgroup H of a group G. As a first application of this technique, we prove that if κ is an infinite cardinal, then for each ordinal $\alpha < \kappa^+$, there exists a centreless group G of cardinality κ such that the automorphism tower of G terminates after exactly α steps. In Chapter 5, we present an account of Hamkins' work on the automorphism towers of arbitrary (not necessarily centreless) groups. Chapter 6 contains an introduction to set-theoretic forcing, which is intended to be intelligible to beginning graduate students. In Chapter 7, we show that it is impossible to find a better bound in ZFC than $\left(2^{|G|}\right)^+$ for the height of the automorphism tower of an infinite centreless group G. For example, it is consistent that there exists a centreless group G of cardinality $\aleph_1$ such that the automorphism tower of G has height strictly greater than $2^{\aleph_1}$. On the other hand, in Chapter 9, we show that it is consistent that the height of the automorphism tower of every centreless group G of cardinality $\aleph_1$ is strictly less than $2^{\aleph_1}$. In Chapter 8, we consider the relationship between the heights of the automorphism towers of a single centreless group G computed in two different models $M \subset N$ of ZFC.

I am very grateful to David Nacin and Steve Warner for carefully reading the manuscript of this book and for supplying me with a substantial number of corrections.

Thanks are due to the many friends and colleagues with whom I have discussed the material in this book, especially Greg Cherlin, Warren Dicks, Ulrich Felgner, Ed Formanek, Ken Hickin, Wilfrid Hodges, Otto Kegel, Felix Leinen, Peter Neumann, Dick Phillips and John Wilson. Particular thanks are due to my coauthors Joel

Hamkins and Winfried Just. Of course, like so many logicians, my greatest debt of gratitude is owed to Saharon Shelah, who I would like to thank for all that he has taught me over the last twenty years.

Finally it is a pleasure to acknowledge the generous financial support of the Alexander von Humboldt Foundation, Rutgers University and the National Science Foundation.

CHAPTER 1

The Automorphism Tower Problem

In this chapter, we shall introduce the notion of the automorphism tower of a centreless group, and we shall illustrate this notion by computing the automorphism towers of a number of groups. In particular, we shall prove that for any integer $n \in \omega$, there exists a finite centreless group F such that the automorphism tower of F terminates after exactly n steps; and we shall show that the automorphism tower of the infinite dihedral group D_∞ terminates after exactly $\omega + 1$ steps. Section 1.2 contains some fundamental algebraic results which will be used repeatedly throughout this book. As we shall see later, Theorem 1.1.10 is the algebraic heart of the proof of the automorphism tower theorem; and Theorem 1.2.8 is the key to the construction of centreless groups with extremely long automorphism towers.

1.1. Automorphism towers

If G is a group, then the automorphism group of G is denoted by $\mathrm{Aut}(G)$. In this book, we shall always work with the left action of $\mathrm{Aut}(G)$ on G. Thus if φ, $\psi \in \mathrm{Aut}(G)$, then the product $\varphi\psi$ is the automorphism defined by

$$\varphi\psi(x) = \varphi(\psi(x))$$

for each $x \in G$. If $g \in G$, then the corresponding *inner automorphism* $i_g \in \mathrm{Aut}(G)$ is defined by

$$i_g(x) = gxg^{-1}$$

for each $x \in G$. Notice that if g, $h \in G$, then

$$i_g i_h(x) = ghxh^{-1}g^{-1} = (gh)x(gh)^{-1} = i_{gh}(x)$$

for all $x \in G$. Thus there is a natural homomorphism from G to $\mathrm{Aut}(G)$, obtained by sending each $g \in G$ to the corresponding inner automorphism $i_g \in \mathrm{Aut}(G)$. The image of this homomorphism is the group $\mathrm{Inn}(G)$ of inner automorphisms of G and the kernel is the centre $Z(G)$ of G.

LEMMA 1.1.1. *Let G be an arbitrary group.*

(a) *If $g \in G$ and $\pi \in \operatorname{Aut}(G)$, then $\pi i_g \pi^{-1} = i_{\pi(g)}$.*
(b) *$\operatorname{Inn} G$ is a normal subgroup of $\operatorname{Aut}(G)$.*

PROOF. (a) For each $x \in G$, we have that

$$\begin{aligned}(\pi i_g \pi^{-1})(x) &= \pi(g\pi^{-1}(x)g^{-1}) \\ &= \pi(g)x\pi(g)^{-1} \\ &= i_{\pi(g)}(x).\end{aligned}$$

(b) This is an immediate consequence of (a). □

$\operatorname{Inn}(G)$ may be regarded as the group of "obvious" automorphisms of the group G. An automorphism $\pi \in \operatorname{Aut}(G) \smallsetminus \operatorname{Inn}(G)$ is called an *outer automorphism* and the quotient group

$$\operatorname{Out}(G) = \operatorname{Aut}(G)/\operatorname{Inn}(G)$$

is called the *outer automorphism group* of G. In most of this book, we shall be interested in the case when G is a centreless group. In this case, the natural homomorphism $g \mapsto i_g$ is an embedding and $G \simeq \operatorname{Inn}(G)$.

LEMMA 1.1.2. *If G is a centreless group, then $C_{\operatorname{Aut}(G)}(\operatorname{Inn}(G)) = 1$. In particular, $\operatorname{Aut}(G)$ is also a centreless group.*

PROOF. If $\pi \in C_{\operatorname{Aut}(G)}(\operatorname{Inn}(G))$, then for each $g \in G$, we have that

$$i_{\pi(g)} = \pi i_g \pi^{-1} = i_g$$

and hence $\pi(g) = g$. Thus $\pi = 1$. □

EXAMPLE 1.1.3. Let n be an integer such that $n > 2$ and $n \neq 6$. Let $\operatorname{Sym}(n)$ be the group of all permutations of the set $\{0, 1, 2, \dots, n-1\}$ and let $\operatorname{Alt}(n)$ be the subgroup of even permutations. Then $\operatorname{Alt}(n)$ is a normal subgroup of $\operatorname{Sym}(n)$. Hence each $\pi \in \operatorname{Sym}(n)$ yields a corresponding automorphism c_π of $\operatorname{Alt}(n)$, defined by $c_\pi(\varphi) = \pi\varphi\pi^{-1}$ for each $\varphi \in \operatorname{Alt}(n)$. In fact, every automorphism of $\operatorname{Alt}(n)$ arises in this manner. (See Suzuki [**48**, Section 3.2].) Thus $\operatorname{Aut}(\operatorname{Alt}(n)) \simeq \operatorname{Sym}(n)$ and $\operatorname{Out}(\operatorname{Alt}(n))$ is the cyclic group of order 2.

It is well-known that if $n > 2$ and $n \neq 6$, then every automorphism of $\mathrm{Sym}(n)$ is inner. (Once again, this result can be found in Suzuki [**48**, Section 3.2].) Thus we have found a natural embedding of $\mathrm{Alt}(n)$ into a complete group, i.e. a centreless group with no "nonobvious" automorphisms.

DEFINITION 1.1.4. A centreless group H is *complete* if $\mathrm{Aut}(H) = \mathrm{Inn}(H)$.

We obtain a similar situation if G is an arbitrary centreless group. Identify G with $\mathrm{Inn}(G)$ via the natural embedding, $g \mapsto i_g$, so that $G \trianglelefteq G_1 = \mathrm{Aut}(G)$. By Lemma 1.1.2, G_1 is a centreless group; and by Lemma 1.1.1(a), every automorphism of G is induced by an inner automorphism of G_1. Of course, it is possible that G_1 might possess outer automorphisms. In this case, we can continue on to the larger centreless groups $G_2 = \mathrm{Aut}(G_1)$, $G_3 = \mathrm{Aut}(G_2)$, $\dots$, etc. And by making suitable identifications, we obtain an ascending chain of groups

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_n \trianglelefteq G_{n+1} \trianglelefteq \cdots$$

such that $G_{n+1} = \mathrm{Aut}(G_n)$ for each $n \in \omega$. This chain is the beginning of the automorphism tower of G. A classical result of Wielandt [**52**] says that if G is finite, then there exists an integer n such that G_n is a complete group. In other words, the automorphism tower of a finite centreless group G terminates after finitely many steps. (We shall present a proof of Wielandt's theorem in Section 2.1.) However, there are natural examples of infinite centreless groups G such that the automorphism tower of G does not terminate after finitely many steps.

EXAMPLE 1.1.5. A group D is said to be a *dihedral group* if D is generated by two involutions a and b. (An *involution* is an element of order 2.) A dihedral group D is determined up to isomorphism by the order of the element $c = ab$. If c has finite order $n > 1$, then D is the dihedral group D_{2n} of order $2n$. If c has infinite order, then D is the infinite dihedral group D_∞.

Let $D_\infty = \langle a, b \rangle$ be the infinite dihedral group. Then $D_\infty = \langle a \rangle * \langle b \rangle$ is the free product of its cyclic subgroups $\langle a \rangle$ and $\langle b \rangle$. It follows easily that D_∞ is a centreless group and that D_∞ has an outer automorphism π of order 2 which interchanges the elements a and b. In Section 1.4, we will prove that $\mathrm{Aut}(D_\infty) = \langle \pi, i_a \rangle$. Thus

$\mathrm{Aut}(D_\infty)$ is also an infinite dihedral group and so

$$\mathrm{Inn}(D_\infty) < \mathrm{Aut}(D_\infty) \simeq D_\infty.$$

It follows that for each $n \in \omega$, the nth group in the automorphism tower of D_∞ is isomorphic to D_∞; and hence the automorphism tower of D_∞ does not terminate after finitely many steps.

The above example suggests that we should extend the notion of the automorphism tower past the finite stages and on into the transfinite.

DEFINITION 1.1.6. Let G be a centreless group. Then the *automorphism tower* of G is the ascending chain of groups

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \cdots$$

such that for each ordinal α

(a) $G_{\alpha+1} = \mathrm{Aut}(G_\alpha)$; and
(b) if α is a limit ordinal, then $G_\alpha = \bigcup_{\beta<\alpha} G_\beta$.

(At each successor step, we identify G_α with $\mathrm{Inn}\, G_\alpha$ via the natural embedding.)

The automorphism tower of G is said to terminate if there exists an ordinal α such that $G_\alpha = G_{\alpha+1}$. Of course, this occurs iff there exists an ordinal α such that G_α is a complete group.

EXAMPLE 1.1.7. Let n be an integer such that $n > 3$ and $n \neq 6$. Then our earlier remarks show that the automorphism tower of $\mathrm{Alt}(n)$ terminates after exactly 1 step.

EXAMPLE 1.1.8. In Section 1.4, we shall prove that the automorphism tower of the infinite dihedral group D_∞ terminates after exactly $\omega + 1$ steps.

The original version of the automorphism tower problem simply asked whether the automorphism tower of an arbitrary centreless group G eventually terminates. This question will be answered positively by the automorphism tower theorem, which we shall prove in Section 3.1. The proof is extremely simple and uses only the most basic results on automorphism towers, together with some elementary properties of the infinite cardinal numbers. To some extent, this is to be expected.

It is inconceivable that the proof could make use of a detailed understanding of the structure of the automorphism tower of an arbitrary centreless group, since no such understanding exists, nor is it ever likely to exist. However, the automorphism tower theorem is not merely a set-theoretic triviality. To see this, consider the following superficially similar problem. (I am grateful to Warren Dicks for pointing this out to me.)

Let R be a ring and let $\operatorname{End}_{\mathbb{Z}}(R)$ be the ring of endomorphisms $\varphi : R \to R$ of the additive group of R. Then there is a natural ring homomorphism of R into $\operatorname{End}_{\mathbb{Z}}(R)$, obtained by sending each $r \in R$ to the corresponding left multiplication map $\lambda_r \in \operatorname{End}_{\mathbb{Z}}(R)$, defined by $\lambda_r(x) = rx$ for each $x \in R$. Since $\lambda_r(1) = r$, this homomorphism is an embedding; and so we can identify R with the subring $\{\lambda_r \mid r \in R\}$ of $\operatorname{End}_{\mathbb{Z}}(R)$. Continuing in this fashion, we obtain an ascending chain of rings

$$R = R_0 \leqslant R_1 \leqslant R_2 \leqslant \cdots \leqslant R_\alpha \leqslant R_{\alpha+1} \leqslant \cdots$$

such that for each ordinal α,

(a) $R_{\alpha+1} = \operatorname{End}_{\mathbb{Z}}(R_\alpha)$; and
(b) if α is a limit ordinal, then $R_\alpha = \bigcup_{\beta<\alpha} R_\beta$.

This chain is called the *endomorphism tower* of R. The endomorphism tower problem asks whether the endomorphism tower of an arbitrary ring eventually terminates. The following result shows that most endomorphism rings never terminate.

THEOREM 1.1.9. *If R is a noncommutative ring, then R_α is a proper subring of $R_{\alpha+1}$ for all ordinals α.*

PROOF. If R is a noncommutative ring, then R_α is also noncommutative for each ordinal α. Hence it is enough to prove that if S is a noncommutative ring, then $\{\lambda_s \mid s \in S\}$ is a proper subring of $\operatorname{End}_{\mathbb{Z}}(S)$. To see this, let $s \in S$ be a noncentral element and let $\rho_s \in \operatorname{End}_{\mathbb{Z}}(S)$ be the corresponding right multiplication map, defined by $\rho_s(x) = xs$ for all $x \in S$. Suppose that there exists $t \in S$ such that $\rho_s = \lambda_t$. Then $s = \rho_s(1) = \lambda_t(1) = t$. This means that $xs = \rho_s(x) = \lambda_s(x) = sx$ for all $x \in S$, which contradicts the assumption that s is a noncentral element. $\square$

So what is the essential algebraic difference between automorphism towers of centreless groups and endomorphism towers of rings? Recall that Lemma 1.1.2

says that if G is a centreless group, then $C_{\mathrm{Aut}(G)}(\mathrm{Inn}(G)) = 1$; i.e. $C_{G_1}(G_0) = 1$. In Section 3.1, we shall see that the automorphism tower theorem is an easy set-theoretic consequence of the following generalisation of Lemma 1.1.2.

THEOREM 1.1.10. *Let G be a centreless group, and let*

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \cdots$$

be the automorphism tower of G. Then $C_{G_\alpha}(G_0) = 1$ for all ordinals α.

We shall prove Theorem 1.1.10 in the next section. Notice that, by taking $\alpha = 2$, it follows that if G is a centreless group, then each element $\pi \in G_2 = \mathrm{Aut}(G_1)$ is uniquely determined by its restriction $\pi \restriction G$. The analogous statement is false for every noncommutative ring.

PROPOSITION 1.1.11. *If R is a noncommutative ring, then there exist elements $\varphi \neq \psi$ in $R_2 = End_{\mathbb{Z}}(R_1)$ such that $\varphi \restriction R = \psi \restriction R$.*

PROOF. Once again, if S is a ring and $s \in S$, then λ_s, $\rho_s \in \mathrm{End}_{\mathbb{Z}}(S)$ will denote the corresponding left multiplication and right multiplication maps. Let R be a noncommutative ring and let $r \in R$ be a noncentral element. Consider the endomorphisms $\varphi = \lambda_{\rho_r}$ and $\psi = \rho_{\rho_r}$ of $R_1 = \mathrm{End}_{\mathbb{Z}}(R)$. Since ρ_r is a noncentral element of R_1, it follows that $\varphi \neq \psi$. On the other hand, in passing from R to R_1, each element $s \in R$ becomes identified with the endomorphism λ_s; and so

$$\varphi(s) = \lambda_{\rho_r}(\lambda_s) = \rho_r \circ \lambda_s = \lambda_s \circ \rho_r = \rho_{\rho_r}(\lambda_s) = \psi(s).$$

Consequently $\varphi \restriction R = \psi \restriction R$. □

1.2. Some fundamental results

In this section, we shall prove a number of fundamental group-theoretic results, which will be used repeatedly throughout this book. In particular, we shall prove Theorem 1.1.10, which forms the algebraic heart of the proof of the automorphism tower theorem.

During the proof of Theorem 1.1.10, we shall make use of some basic properties of commutators and commutator subgroups.

DEFINITION 1.2.1. Let G be a group.

(a) If a, $b \in G$, then the *commutator* of a and b is defined to be

$$[a, b] = a^{-1}b^{-1}ab.$$

(b) If A, B are nonempty subsets of G, then the *commutator subgroup* of A and B is defined to be

$$[A, B] = \langle [a, b] \mid a \in A, b \in B \rangle.$$

LEMMA 1.2.2. *If G is a group and a, b, $c \in G$, then*

(a) $[a, b]^{-1} = [b, a]$,
(b) $[ab, c] = [a, c]^b[b, c]$,
(c) $[c, ab] = [c, b][c, a]^b$.

PROOF. Each of these identities is easily checked using the definition of the commutator. For example,

$$\begin{aligned} [ab, c] &= (ab)^{-1}c^{-1}abc \\ &= b^{-1}a^{-1}c^{-1}abc \\ &= b^{-1}(a^{-1}c^{-1}ac)b(b^{-1}c^{-1}bc) \\ &= b^{-1}[a, c]b[b, c] \\ &= [a, c]^b[b, c]. \end{aligned}$$

□

LEMMA 1.2.3. *Let A and B be subgroups of the group G.*

(a) $[A, B] = [B, A]$.
(b) $[A, B] \leqslant A$ *iff* $B \leqslant N_G(A)$.

PROOF. Clearly (a) is an immediate consequence of Lemma 1.2.2(a). Thus it is enough to prove (b). First suppose that $[A, B] \leqslant A$. Then for any $a \in A$ and $b \in B$, we have that $a^{-1}b^{-1}ab = [a, b] \in A$ and so $b^{-1}ab \in A$. It follows that $b^{-1}Ab = A$ for all $b \in B$ and so $B \leqslant N_G(A)$. Conversely, suppose that $B \leqslant N_G(A)$. Then for each $a \in A$ and $b \in B$, $[a, b] = a^{-1}(b^{-1}ab) \in A$ and so $[A, B] \leqslant A$. □

Suppose that G is a centreless group and that

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \cdots$$

is the automorphism tower of G. Then for each ordinal α,

$$C_{G_{\alpha+1}}(G_\alpha) = C_{\mathrm{Aut}(G_\alpha)}(\mathrm{Inn}(G_\alpha)) = 1.$$

Thus Theorem 1.1.10 is an immediate consequence of the following slightly more general result.

Lemma 1.2.4. *Let G be a centreless group and let*

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \cdots$$

be an ascending chain of groups such that

(a) $C_{G_{\alpha+1}}(G_\alpha) = 1$ *for all α; and*
(b) *if α is a limit ordinal, then $G_\alpha = \bigcup_{\beta<\alpha} G_\beta$.*

Then $C_{G_\alpha}(G_0) = 1$ for all ordinals α.

Proof. Suppose that the result fails and let α be the least ordinal such that $C = C_{G_\alpha}(G_0) \neq 1$. Since G is centreless, it follows that $\alpha > 0$. Also it is clear that α is not a limit ordinal. Hence there exists an ordinal β such that $\alpha = \beta + 1$.

Claim 1.2.5. $[G_\gamma, C] = 1$ *for all $\gamma \leq \beta$.*

Proof of Claim 1.2.5. Once again, suppose that the result fails and let γ be the least ordinal such that $[G_\gamma, C] \neq 1$. Since $[G_0, C_{G_\alpha}(G_0)] = 1$, it follows that $\gamma > 0$; and it is clear that γ is not a limit ordinal. Hence there exists an ordinal δ such that $\gamma = \delta + 1$. By the minimality of γ, we have that $[G_\delta, C] = 1$ and so

$$C_{G_\alpha}(G_0) = C \leqslant C_{G_\alpha}(G_\delta).$$

Since $G_0 \leqslant G_\delta$, we also have that

$$C_{G_\alpha}(G_\delta) \leqslant C_{G_\alpha}(G_0) = C.$$

Hence $C = C_{G_\alpha}(G_\delta)$.

Next we shall show that G_γ normalises $C = C_{G_\alpha}(G_\delta)$. Suppose that $g \in G_\gamma$. Since $G_\delta \trianglelefteq G_{\delta+1} = G_\gamma$, it follows that $gG_\delta g^{-1} = G_\delta$. Because $\gamma < \alpha$, we have that $g \in G_\alpha$ and so $gG_\alpha g^{-1} = G_\alpha$. Consequently, $gC_{G_\alpha}(G_\delta)g^{-1} = C_{G_\alpha}(G_\delta)$, as required.

Since G_γ normalises C, Lemma 1.2.3 yields that $[G_\gamma, C] = [C, G_\gamma] \leqslant C$. Now notice that $G_\gamma \leqslant G_\beta$ and that $C = C_{G_\alpha}(G_\delta) \leqslant G_\alpha = G_{\beta+1}$. Since $G_{\beta+1}$ normalises G_β, it follows that $[G_\gamma, C] \leqslant [G_\beta, G_{\beta+1}] \leqslant G_\beta$. Thus we have established that

$$[G_\gamma, C] \leqslant C \cap G_\beta = C_{G_\alpha}(G_0) \cap G_\beta = C_{G_\beta}(G_0).$$

By the minimality of α, we have that $C_{G_\beta}(G_0) = 1$ and so $[G_\gamma, C] = 1$, which is a contradiction. $\square$

By Claim 1.2.5, $C \leqslant C_{G_\alpha}(G_\beta) = C_{G_{\beta+1}}(G_\beta) = 1$, which is the final contradiction. This completes the proof of Lemma 1.2.4. $\square$

In the next two sections, we shall compute the automorphism towers of some well-known groups. The following result provides a useful criterion for recognising when an automorphism tower has reached a complete group.

DEFINITION 1.2.6. Let H be a subgroup of the group G. Then H is said to be a *characteristic subgroup* of G, written H char G, if $\pi[H] = H$ for every automorphism $\pi \in \mathrm{Aut}(G)$.

If H is a characteristic subgroup of the group G, then $gHg^{-1} = i_g[H] = H$ for every $g \in G$ and so H is a normal subgroup of G. However, the converse need not be true. For example, let $G = \langle u \rangle \times \langle v \rangle$ be an elementary abelian p-group of order p^2. Then $H = \langle u \rangle$ is a normal subgroup which is not characteristic in G. Of course, if G is a complete group and $H \leqslant G$, then $H \trianglelefteq G$ iff H char G.

THEOREM 1.2.7 (Burnside). *If G is a centreless group, then the following statements are equivalent.*

(a) $\mathrm{Aut}(G)$ *is a complete group.*
(b) $\mathrm{Inn}(G)$ *is a characteristic subgroup of* $\mathrm{Aut}(G)$.

PROOF. If $\mathrm{Aut}(G)$ is a complete group, then the normal subgroup $\mathrm{Inn}(G)$ is clearly a characteristic subgroup of $\mathrm{Aut}(G)$. Conversely, suppose that $\mathrm{Inn}(G)$ is a characteristic subgroup of $\mathrm{Aut}(G)$. Let

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \cdots$$

be the automorphism tower of G. Notice that $\mathrm{Aut}(G)$ is a complete group iff $G_1 = G_2$. Let $g \in G_2$. Then g induces an automorphism of G_1 via conjugation; and since G_0 char G_1, it follows that $gG_0g^{-1} = G_0$. Hence there exists an element $h \in G_1 = \mathrm{Aut}(G_0)$ such that $hxh^{-1} = gxg^{-1}$ for all $x \in G_0$. This implies that $g^{-1}h \in C_{G_2}(G_0) = 1$ and so $g = h \in G_1$. Hence $G_1 = G_2$. □

Using Theorem 1.2.7, we can now easily prove that the automorphism tower of any simple nonabelian group terminates after at most 1 step.

THEOREM 1.2.8 (Burnside). *Let S be a simple nonabelian group, and let G be a group such that* $\mathrm{Inn}(S) \leqslant G \leqslant \mathrm{Aut}(S)$.

(a) $\mathrm{Inn}(S)$ *is the unique minimal nontrivial normal subgroup of* G.

(b) $\mathrm{Aut}(S)$ *is a complete group.*

PROOF. (a) It is clear that $\mathrm{Inn}(S)$ is a minimal nontrivial normal subgroup of G. Suppose that there exists a second minimal nontrivial normal subgroup N. Applying Lemma 1.2.3, we find that $[\mathrm{Inn}(S), N] \leqslant \mathrm{Inn}(S) \cap N$. As $\mathrm{Inn}(S)$ and N are distinct minimal nontrivial normal subgroups, $\mathrm{Inn}(S) \cap N = 1$ and so $N \leqslant C_{\mathrm{Aut}(S)}(\mathrm{Inn}(S)) = 1$, which is a contradiction.

(b) Since $\mathrm{Inn}(S)$ is the unique minimal nontrivial normal subgroup of $\mathrm{Aut}(S)$, it follows that $\mathrm{Inn}(S)$ is a characteristic subgroup of $\mathrm{Aut}(S)$. By Theorem 1.2.7, $\mathrm{Aut}(S)$ is a complete group. □

1.3. Some examples of automorphism towers

In this section, we shall present some examples of automorphism towers. First we shall show that for each integer $n \in \omega$, there exists a finite centreless group F such that the automorphism tower of F terminates after exactly n steps. Then we shall compute the automorphism towers of some well-known infinite permutation groups: the alternating group $\mathrm{Alt}(\kappa)$ for $\kappa \geq \omega$ and the automorphism group $\mathrm{Aut}(\mathbb{Q})$ of the linearly ordered set $\mathbb{Q}$ of rational numbers. We shall also introduce some important notions and notation from the theory of permutation groups.

EXAMPLE 1.3.1. For each $n \in \omega$, we shall construct an example of a finite centreless group $F(n)$ such that the automorphism tower of $F(n)$ terminates after exactly n steps. First we define an auxiliary sequence of groups T_n inductively as follows.

(a) $T_0 = \mathrm{Sym}(5)$.

(b) $T_{n+1} = [T_n \times T_n] \rtimes \langle \sigma_n \rangle$, where σ_n is an involution which interchanges the factors $T_n \times 1$ and $1 \times T_n$ of the direct product $T_n \times T_n$; i.e.

$$\sigma_n(t,1)\sigma_n^{-1} = (1,t)$$

for all $t \in T_n$.

We shall show that $F(n) = T_0 \times \prod_{\ell<n} T_\ell$ satisfies our requirements.

Before becoming involved in the details of the proof, we shall explain the idea behind this construction. Consider $F(3) = T_0 \times T_0 \times T_1 \times T_2$. Then $F(3)$ has an obvious outer automorphism σ of order 2, which interchanges the first two factors of the product $T_0 \times T_0 \times T_1 \times T_2$. We shall see that $\mathrm{Aut}(F(3)) = \langle \mathrm{Inn}(F(3)), \sigma \rangle$, so that $F(3)_1 = \mathrm{Aut}(F(3)) \simeq T_1 \times T_1 \times T_2$. Similarly we shall find that $F(3)_2 \simeq T_2 \times T_2$ and that the automorphism tower of $F(3)$ terminates with the group $F(3)_3 \simeq T_3$.

Our argument is based upon the uniqueness of the Remak decomposition of $F(n)$, together with a result of Peter Neumann on the structure of wreath products of groups. First we shall recall some of the basic theory of Remak decompositions of finite groups.

Definition 1.3.2. A nontrivial group G is said to be *decomposable* if there exist nontrivial normal subgroups H_1, H_2 of G such that $G = H_1 \times H_2$. Otherwise, G is said to be *indecomposable*.

Clearly if G is a nontrivial finite group, then G can be expressed as the direct product of finitely many indecomposable subgroups. Such a direct product decomposition, $G = H_1 \times \cdots \times H_r$, is called a *Remak decomposition* of G. The set of factors $\{H_1, \dots, H_r\}$ is generally not uniquely determined. For example, if $G = \langle u, v \rangle$ is an elementary abelian p-group of order p^2, then $\langle u \rangle \times \langle v \rangle$ and $\langle u \rangle \times \langle uv \rangle$ are Remak decompositions of G with distinct sets of factors. However, if G is a nontrivial finite centreless group, then there is essentially a unique Remak decomposition of G.

Theorem 1.3.3. *Suppose that G is a nontrivial finite centreless group. If*

$$G = H_1 \times \cdots \times H_r = K_1 \times \cdots \times K_s$$

are Remak decompositions of G*, then* $r = s$ *and there exists a permutation* π *of* $\{1, \dots, r\}$ *such that* $H_i = K_{\pi(i)}$ *for all* $1 \le i \le r$*.*

Proof. See Suzuki [**48**, Section 2.4]. □

Next we shall introduce the notion of the (standard) wreath product $A \mathrm{~Wr~} G$ of the groups A and G. First we define the *base group* B to be the group of all functions $b : G \to A$. If $b_1, b_2 \in B$, then their product $b_1 b_2 \in B$ is the function such that $b_1 b_2(x) = b_1(x) b_2(x)$ for all $x \in G$. Thus $B = \prod_{x \in G} A_x$, where each $A_x = \{b \in B \mid b(y) = 1 \text{ for all } y \neq x\} \simeq A$. The *wreath product* $A \mathrm{~Wr~} G$ is defined to be the semidirect product $B \rtimes G$, where $(gbg^{-1})(x) = b(g^{-1}x)$ for each $b \in B$ and $g, x \in G$. Thus $gA_x g^{-1} = A_{gx}$ for each $g, x \in G$.

Notice that $T_0 = \mathrm{Sym}(5)$ is a complete indecomposable group and that T_{n+1} is isomorphic to $T_n \mathrm{~Wr~} C_2$. Hence the following result of Peter Neumann implies that T_n is a complete group for each $n \in \omega$.

Lemma 1.3.4. *Let* A *be a complete indecomposable group such that* $A \neq D_6$*, the dihedral group of order 6. Then* A *Wr* C_2 *is also a complete indecomposable group.*

Proof. See the end of Section 10 of Neumann [**34**]. (Neumann does not mention explicitly that $A \mathrm{~Wr~} C_2$ is indecomposable. However, this is an immediate consequence of Theorem 6.1 [**34**], which says that if a wreath product $A \mathrm{~Wr~} G$ decomposes nontrivially into a direct product $P \times Q$, then one of the factors P, Q must be central in $A \mathrm{~Wr~} G$.) □

It is now easy to prove the following result.

Theorem 1.3.5. *For each* $0 \le m \le n$*, the* m*th group in the automorphism tower of* $F(n)$ *is given by*

$$F(n)_m = T_m \times \prod_{m \le \ell < n} T_\ell.$$

Hence the automorphism tower of $F(n)$ *terminates after exactly* n *steps.*

PROOF. Fix some integer n. We shall argue by induction on $m \leq n$. Suppose that $m < n$ and that

$$F(n)_m = T_m \times \prod_{m \leq \ell < n} T_\ell = T_m \times T_m \times \prod_{m+1 \leq \ell < n} T_\ell.$$

By Lemma 1.3.4, each of the groups T_ℓ is indecomposable. Thus we have just displayed a Remak decomposition of $F(n)_m$. Furthermore, since $F(n)_m$ is centreless, this is the unique Remak decomposition of $F(n)_m$, up to the order of the factors. Let $\sigma \in \operatorname{Aut}(F(n)_m)$ be the obvious outer automorphism of order 2 such that

(a) σ interchanges the first two factors of the above Remak decomposition, and
(b) σ acts as the identity automorphism on the remaining factors.

Then

$$\langle \operatorname{Inn}(F(n)_m), \sigma \rangle \simeq T_{m+1} \times \prod_{m+1 \leq \ell < n} T_\ell$$

and so we must show that $\operatorname{Aut}(F(n)_m) = \langle \operatorname{Inn} F(n)_m, \sigma \rangle$. Let $\varphi \in \operatorname{Aut}(F(n)_m)$ be an arbitrary automorphism. Then φ must permute the factors of the Remak decomposition of $F(n)_m$. After replacing φ by $\sigma\varphi$ if necessary, we can suppose that $\varphi[T] = T$ for each factor T. Since each factor T is a complete group, there exist elements $g_T \in T$ such that $\varphi \restriction T = i_{g_T} \restriction T$. Let $g = \prod_T g_T \in F(n)_m$. Then $\varphi = i_g \in \operatorname{Inn} F(n)_m$.

Finally note that $F(n)_n = T_n$ is a complete group. Thus the automorphism tower of $F(n)$ terminates after exactly n steps. □

Next we need to introduce some notions and notation from the theory of permutation groups. Let Ω be a nonempty set. Then $\operatorname{Sym}(\Omega)$ denotes the group of all permutations of Ω. A group G is said to be a permutation group on Ω if G is a subgroup of $\operatorname{Sym}(\Omega)$. Suppose that G is a permutation group on Ω and that $\Delta \subseteq \Omega$. Then

$$G_{\{\Delta\}} = \{g \in G \mid g[\Delta] = \Delta\}$$

and

$$G_{(\Delta)} = \{g \in G \mid g(x) = x \text{ for all } x \in \Delta\}$$

are the setwise and pointwise stabilisers of Δ in G. When $\Delta = \{x\}$ is a singleton, then we usually write G_x for the stabiliser of Δ in G. If $x \in \Omega$, then

$$\mathrm{Orb}_G(x) = \{g(x) \mid g \in G\}$$

is the G-orbit containing x. In this book, we shall make repeated use of the fact that

$$[G : G_x] = |\mathrm{Orb}_G(x)|$$

for all $x \in \Omega$. (For example, see Suzuki [**48**, Section 1.7].)

EXAMPLE 1.3.6. Let κ be an infinite cardinal. Remember that κ is the set of ordinals α such that $\alpha < \kappa$ and so κ is itself a canonical example of a set of cardinality κ. If $\varphi \in \mathrm{Sym}(\kappa)$, then the fixed point set of φ is

$$\mathrm{fix}(\varphi) = \{\alpha \in \kappa \mid \varphi(\alpha) = \alpha\}$$

and the support of φ is

$$\mathrm{supp}(\varphi) = \{\alpha \in \kappa \mid \varphi(\alpha) \neq \alpha\}.$$

If λ is an infinite cardinal such that $\lambda \leq \kappa$, then we define

$$\mathrm{Sym}_\lambda(\kappa) = \{\pi \in \mathrm{Sym}(\kappa) \mid |\,\mathrm{supp}(\pi)| < \lambda\}.$$

Clearly $\mathrm{Sym}_\lambda(\kappa)$ is a normal subgroup of $\mathrm{Sym}(\kappa)$. If $\lambda = \omega$, then we also write $\mathrm{Fin}(\kappa)$ for the group $\mathrm{Sym}_\omega(\kappa)$ of finite permutations of κ. The alternating group $\mathrm{Alt}(\kappa)$ on κ is the subgroup of finite even permutations of κ.

We shall compute the automorphism tower of $A = \mathrm{Alt}(\kappa)$. Since A is a simple group, Theorem 1.2.8 says that $\mathrm{Aut}(A)$ is a complete group and so we need only determine $\mathrm{Aut}(A)$. Since $\mathrm{Alt}(\kappa) \lhd \mathrm{Sym}(\kappa)$, each $\varphi \in \mathrm{Sym}(\kappa)$ yields a corresponding automorphism c_φ of $\mathrm{Alt}(\kappa)$, defined by $c_\varphi(x) = \varphi x \varphi^{-1}$ for each $x \in \mathrm{Alt}(\kappa)$. We shall eventually show that every automorphism of $\mathrm{Alt}(\kappa)$ arises in this manner. Thus $\mathrm{Aut}(A) = \mathrm{Sym}(\kappa)$.

To begin our analysis, we shall define a natural action of $\mathrm{Aut}(A)$ on an associated graph $\Gamma = \langle \Omega, E \rangle$. Let

$$\Omega = \{\langle \sigma \rangle \mid \sigma \in \mathrm{Alt}(\kappa) \text{ is a 3-cycle }\}.$$

Notice that the elements of Ω are exactly those subgroups T of $\mathrm{Alt}(\kappa)$ such that

(a) $T \simeq C_3$, the cyclic group of order 3, and

(b) $C_{\mathrm{Alt}(\kappa)}(T) \simeq C_3 \times \mathrm{Alt}(\kappa)$.

Hence if $\varphi \in \mathrm{Aut}(A)$, then φ induces a permutation $\overline{\varphi}$ of Ω, defined by $\overline{\varphi}(T) = \varphi[T]$. For each $T = \langle \sigma \rangle \in \Omega$, let $\Delta_T = \mathrm{supp}(\sigma)$. Note that if T_1, T_2 are distinct elements of Ω, then the subgroup $\langle T_1, T_2 \rangle$ is isomorphic to

(i) $\mathrm{Alt}(4)$, if $|\Delta_{T_1} \cap \Delta_{T_2}| = 2$;

(ii) $\mathrm{Alt}(5)$, if $|\Delta_{T_1} \cap \Delta_{T_2}| = 1$;

(iii) $C_3 \times C_3$, if $|\Delta_{T_1} \cap \Delta_{T_2}| = 0$.

We define a graph structure on Ω by specifying that

$$\{T_1, T_2\} \in E \quad \text{iff} \quad \langle T_1, T_2 \rangle \simeq \mathrm{Alt}(4).$$

Clearly if $\varphi \in \mathrm{Aut}(A)$, then the corresponding permutation $\overline{\varphi}$ of Ω is an automorphism of $\Gamma = \langle \Omega; E \rangle$. Let $\pi : \mathrm{Aut}(A) \to \mathrm{Aut}(\Gamma)$ be the homomorphism defined by $\pi(\varphi) = \overline{\varphi}$.

LEMMA 1.3.7. *π is an embedding.*

PROOF. Let $\psi \in A \smallsetminus 1$. Let $\alpha \in \mathrm{supp}(\psi)$ and let $\sigma = (\,\alpha\,\beta\,\gamma\,)$ be a 3-cycle such that $\mathrm{supp}(\psi) \cap \mathrm{supp}(\sigma) = \{\alpha\}$. Then

$$i_\psi(\sigma) = \psi\sigma\psi^{-1} = (\,\psi(\alpha)\,\psi(\beta)\,\psi(\gamma)\,) \notin \langle \sigma \rangle.$$

Thus $\overline{i_\psi}(\langle \sigma \rangle) \neq \langle \sigma \rangle$ and so $i_\psi \notin \ker \pi$. Hence $\ker \pi$ is a normal subgroup of $\mathrm{Aut}(A)$ such that $\ker \pi \cap \mathrm{Inn}(A) = 1$. So Lemma 1.3.7 is an immediate consequence of Lemma 1.3.8. □

LEMMA 1.3.8. *Let G be a centreless group and let N be a normal subgroup of $\mathrm{Aut}(G)$. If $N \cap \mathrm{Inn}(G) = 1$, then $N = 1$.*

PROOF. By Lemma 1.2.3, we have that $[N, \mathrm{Inn}(G)] \leqslant N \cap \mathrm{Inn}(G) = 1$ and hence $N \leqslant C_{\mathrm{Aut}(G)}(\mathrm{Inn}(G)) = 1$. □

In order to make the argument as transparent as possible, we will identify Γ with the isomorphic graph $\langle V, \sim \rangle$, where

(1) V is the set of 3-subsets of κ; and

(2) $\Delta_1 \sim \Delta_2$ iff $|\Delta_1 \cap \Delta_2| = 2$.

Of course, this corresponds to identifying each subgroup $T = \langle \sigma \rangle \in \Omega$ with its support $\Delta_T \in V$. Notice that if $\varphi \in \mathrm{Sym}(\kappa)$ and $T = \langle (\, \alpha \, \beta \, \gamma \,) \rangle$, then

$$c_\varphi[T] = \varphi T \varphi^{-1} = \langle (\, \varphi(\alpha) \, \varphi(\beta) \, \varphi(\gamma) \,) \rangle.$$

Hence the associated action of $\mathrm{Sym}(\kappa)$ on $\langle V, \sim \rangle$ is the natural one, defined by

$$\Delta \stackrel{\varphi}{\mapsto} \varphi[\Delta]$$

for each $\varphi \in \mathrm{Sym}(\kappa)$ and $\Delta \in V$. We shall show that every automorphism of $\Gamma = \langle V, \sim \rangle$ is induced by an element of $\mathrm{Sym}(\kappa)$ in this manner; and thus complete the proof that $\mathrm{Aut}(A) = \mathrm{Sym}(\kappa)$.

DEFINITION 1.3.9. (a) If $\alpha < \beta < \kappa$, then $C_{\{\alpha,\beta\}} = \{\Delta \in V \mid \alpha, \beta \in \Delta\}$.
(b) If $\alpha < \kappa$, then $D_{\{\alpha\}} = \{\Delta \in V \mid \alpha \in \Delta\}$.

We shall show that each $\psi \in \mathrm{Aut}(\Gamma)$ induces permutations of the collections $\mathcal{C} = \{C_{\{\alpha,\beta\}} \mid \alpha < \beta < \kappa\}$ and $\mathcal{D} = \{D_{\{\alpha\}} \mid \alpha < \kappa\}$. The next lemma shows that $\mathcal{C}$ is invariant under the action of $\mathrm{Aut}(\Gamma)$.

LEMMA 1.3.10. *If C is an infinite maximal complete subgraph of Γ, then there exist $\alpha < \beta < \kappa$ such that $C = C_{\{\alpha,\beta\}}$.*

PROOF. Let C be an infinite maximal complete subgraph of Γ. Fix some $\Delta_0 \in C$. If $\Delta \in C \smallsetminus \{\Delta_0\}$, then $|\Delta_0 \cap \Delta| = 2$. Hence there is an infinite subset $\{\Delta_n \mid 1 \leq n < \omega\}$ of C and a pair of ordinals $\alpha < \beta < \kappa$ such that $\Delta_0 \cap \Delta_n = \{\alpha, \beta\}$ for all $1 \leq n < \omega$. It is easily checked that if $\Delta \sim \Delta_n$ for all $n < \omega$, then $\{\alpha, \beta\} \subset \Delta$. Thus $C \subseteq C_{\{\alpha,\beta\}}$. By the maximality of C, we must have that $C = C_{\{\alpha,\beta\}}$. $\square$

Notice that $C_{\{\alpha,\beta\}} \cap C_{\{\gamma,\delta\}} \neq \emptyset$ iff $\{\alpha, \beta\} \cap \{\gamma, \delta\} \neq \emptyset$. Using this observation, we can easily give an $\mathrm{Aut}(\Gamma)$-invariant characterisation of the elements of $\mathcal{D}$.

LEMMA 1.3.11. *Suppose that D is a subgraph of Γ which is maximal subject to the following conditions.*

(a) *There exists an infinite subset $\{C_i \mid i \in I\}$ of $\mathcal{C}$ such that $D = \bigcup_{i \in I} C_i$; and*
(b) *if i, $j \in I$, then $C_i \cap C_j \neq \emptyset$.*

Then $D = D_{\{\alpha\}}$ for some $\alpha < \kappa$.

PROOF. Suppose that D satisfies the hypotheses of Lemma 1.3.11. For each $i \in I$, let $C_i = C_{\{\alpha_i,\beta_i\}}$. Then $\{\{\alpha_i, \beta_i\} \mid i \in I\}$ is an infinite collection of 2-subsets of κ, which is maximal subject to the condition that $\{\alpha_i, \beta_i\} \cap \{\alpha_j, \beta_j\} \neq \emptyset$ for all $i, j \in I$. Arguing as in the proof of Lemma 1.3.10, we see that there exists $\alpha < \kappa$ such that $\{\{\alpha_i, \beta_i\} \mid i \in I\} = \{\{\alpha, \beta\} \mid \alpha \neq \beta < \kappa\}$. Hence $D = D_{\{\alpha\}}$. □

If $\Delta = \{\alpha, \beta, \gamma\} \in V$, then $\{\Delta\} = D_{\{\alpha\}} \cap D_{\{\beta\}} \cap D_{\{\gamma\}}$. Thus each $\psi \in \mathrm{Aut}(\Gamma)$ is uniquely determined by its action on $\mathcal{D}$. Moreover, for each $\psi \in \mathrm{Aut}(\Gamma)$, there is a corresponding permutation $\varphi \in \mathrm{Sym}(\kappa)$ such that $\psi\left[D_{\{\alpha\}}\right] = D_{\{\varphi(\alpha)\}}$ for each $\alpha < \kappa$. But we clearly have that $\varphi\left[D_{\{\alpha\}}\right] = D_{\{\varphi(\alpha)\}}$ for each $\alpha < \kappa$. Consequently, $\psi = \varphi \in \mathrm{Sym}(\kappa)$. This completes the proof of the following result.

THEOREM 1.3.12. *If* $\kappa \geq \omega$, *then* $\mathrm{Aut}(A) = \mathrm{Sym}(\kappa)$.

□

Theorem 1.3.12 is an important and useful result, which we shall make use of in Section 3.2. In contrast, the final example of this section is neither important nor useful. But it does present a pleasant application of the strong small index property.

EXAMPLE 1.3.13. Let $\mathcal{M}$ be a countable structure. Then the automorphism group $G = \mathrm{Aut}(\mathcal{M})$ is said to have the *strong small index property* if whenever H is a subgroup of G with $[G : H] < 2^\omega$, then there exists a finite subset $X \subset \mathcal{M}$ such that $G_{(X)} \leqslant H \leqslant G_{\{X\}}$. Let $\mathbb{Q}$ be the rational numbers regarded as a linearly ordered set. Then Truss [**51**] has shown that $G = \mathrm{Aut}(\mathbb{Q})$ has the strong small index property. Of course, since every finite linear ordering is rigid, we have that $G_{(X)} = G_{\{X\}}$ for each finite subset X of $\mathbb{Q}$. (Recall that a structure $\mathcal{M}$ is said to be *rigid* iff the identity map $id_{\mathcal{M}}$ is the only automorphism of $\mathcal{M}$.) Thus if $[G : H] < 2^\omega$, then there exists a finite subset X of $\mathbb{Q}$ such that $H = G_{(X)}$. In particular, the maximal proper subgroups H with $[G : H] < 2^\omega$ are precisely those of the form $H = G_q$ for some $q \in \mathbb{Q}$. We shall use this result to compute the automorphism tower of $G = \mathrm{Aut}(\mathbb{Q})$. (It is an easy exercise to show that G is a centreless group.)

First we shall compute $\mathrm{Aut}(G)$. Let $\varphi \in \mathrm{Aut}(G)$. Then φ must permute the set of those maximal proper subgroups H of G such that $[G : H] < 2^\omega$. Thus there is an associated permutation $\overline{\varphi} \in \mathrm{Sym}(\mathbb{Q})$ such that $\varphi\,[G_q] = G_{\overline{\varphi}(q)}$ for each $q \in \mathbb{Q}$. It is easily checked that the mapping $\pi : \mathrm{Aut}(G) \to \mathrm{Sym}(\mathbb{Q})$, defined by $\pi(\varphi) = \overline{\varphi}$, is a homomorphism.

LEMMA 1.3.14. (a) *$\pi(i_g) = g$ for each $g \in G$.*
(b) *π is an embedding.*
(c) *$\pi[\mathrm{Aut}(G)] = N_{\mathrm{Sym}(\mathbb{Q})}(G)$.*

PROOF. (a) If $g \in G$ and $q \in \mathbb{Q}$, then $i_g\,[G_q] = gG_qg^{-1} = G_{g(q)}$.

(b) Since $\ker\pi$ is a normal subgroup of $\mathrm{Aut}(G)$ such that $\ker\pi \cap \mathrm{Inn}(G) = 1$, Lemma 1.3.8 yields that $\ker\pi = 1$.

(c) Using the facts that $\mathrm{Inn}(G) \trianglelefteq \mathrm{Aut}(G)$ and $\pi[\mathrm{Inn}(G)] = G$, it follows that $\pi[\mathrm{Aut}(G)] \leqslant N_{\mathrm{Sym}(\mathbb{Q})}(G)$. But clearly every element of $N_{\mathrm{Sym}(\mathbb{Q})}(G)$ induces an automorphism of G via conjugation. Hence $\pi[\mathrm{Aut}(G)] = N_{\mathrm{Sym}(\mathbb{Q})}(G)$. □

We have just seen that $\mathrm{Aut}(G)$ can be naturally identified with $N_{\mathrm{Sym}(\mathbb{Q})}(G)$. To compute $N_{\mathrm{Sym}(\mathbb{Q})}(G)$, we shall make use of the following basic result on permutation groups.

LEMMA 1.3.15. *Let (H, Ω) be a permutation group. Suppose that $N \trianglelefteq H$ and that Δ is an orbit of N on Ω. Then $h[\Delta]$ is also an N-orbit for each $h \in H$.*

PROOF. By assumption, there exists $x \in \Omega$ such that

$$\Delta = \mathrm{Orb}_N(x) = \{g(x) \mid g \in N\}.$$

Let $h \in H$ and $y = h(x)$. Since $N \trianglelefteq H$, we have that $hN = Nh$ and hence

$$h[\Delta] = \{hg(x) \mid g \in N\} = \{g(y) \mid g \in N\}$$

is the N-orbit containing y. □

We shall apply the above lemma to the action of $N_{\mathrm{Sym}(\mathbb{Q})}(G)$ on $\mathbb{Q} \times \mathbb{Q}$. Note that $G = \mathrm{Aut}(\mathbb{Q})$ has three orbits on $\mathbb{Q} \times \mathbb{Q}$; namely,

(a) $\Delta_0 = \{(q, r) \in \mathbb{Q} \times \mathbb{Q} \mid q = r\}$,
(b) $\Delta_1 = \{(q, r) \in \mathbb{Q} \times \mathbb{Q} \mid q < r\}$,
(c) $\Delta_2 = \{(q, r) \in \mathbb{Q} \times \mathbb{Q} \mid q > r\}$.

Let $g \in N_{\mathrm{Sym}(\mathbb{Q})}(G)$. Then g must permute the three G-orbits Δ_0, Δ_1 and Δ_2. Clearly $g[\Delta_0] = \Delta_0$. Thus there are only two cases to consider. If $g[\Delta_1] = \Delta_1$, then g is an order-preserving permutation of $\mathbb{Q}$ and so $g \in G = \mathrm{Aut}(\mathbb{Q})$. Otherwise, $g[\Delta_1] = \Delta_2$ and g is an order-reversing permutation of $\mathbb{Q}$. Thus $N_{\mathrm{Sym}(\mathbb{Q})}(G)$ is the group B of permutations of $\mathbb{Q}$ which either preserve or reverse the order. (B is the group of those permutations of $\mathbb{Q}$ which preserve the ternary "betweenness" relation R, defined by $R(a,b,c)$ iff either $a < b < c$ or $c < b < a$.) Notice that the product of two order-reversing permutations is an order-preserving permutation. It follows that $[B : G] = 2$.

Finally we shall prove that B is a complete group. By Theorem 1.2.7, it is enough to show that G is a characteristic subgroup of B. Let $\pi \in \mathrm{Aut}(B)$ be an arbitrary automorphism. Then $[B : \pi[G]] = 2$ and so $[G : \pi[G] \cap G] \leq 2$. If $[G : \pi[G] \cap G] = 1$, then $G \leqslant \pi[G]$ and so $G = \pi[G]$. Otherwise, $[G : \pi[G] \cap G] = 2$. Let $H = \pi[G] \cap G$. Since $1 < [G : H] < 2^\omega$, there exists a finite nonempty subset X of $\mathbb{Q}$ such that $H = G_{(X)}$. But then $[G : H] = \omega$, which is a contradiction. We have now completed the proof of the following result.

THEOREM 1.3.16. *The automorphism tower of* $\mathrm{Aut}(\mathbb{Q})$ *terminates after exactly 1 step.*

□

1.4. The infinite dihedral group

In this section, we shall show that the automorphism tower of the infinite dihedral group $G = D_\infty$ terminates after exactly $\omega + 1$ steps. Before we can compute $\mathrm{Aut}(G)$, we must first present a more detailed account of the structure of G. Remember that $G = \langle a \rangle * \langle b \rangle$ is the free product of the cyclic subgroups generated by the involutions a and b. Let $c = ab$. Then c is an element of infinite order and

$$aca^{-1} = aaba = ba = c^{-1}.$$

Thus $C = \langle c \rangle$ is a normal subgroup of G and G is the semidirect product $C \rtimes \langle a \rangle$. In fact, C is a characteristic subgroup of G. To see this, note that each element $g \in G \smallsetminus C$ has the form $g = c^n a$ for some $n \in \mathbb{Z}$ and that

$$(c^n a)^2 = c^n a c^n a^{-1} = c^n c^{-n} = 1.$$

Thus $G \smallsetminus C$ is the set of involutions of G.

LEMMA 1.4.1. (a) $[G,G] = \langle c^2 \rangle$.

(b) *G has two conjugacy classes of involutions; namely,* $a^G = \{c^{2n}a \mid n \in \mathbb{Z}\}$ *and* $b^G = \{c^{2n+1}a \mid n \in \mathbb{Z}\}$.

PROOF. (a) Note that $c^2 = abab = [a,b]$ and so $\langle c^2 \rangle \leqslant [G,G]$. Since the quotient group $G/\langle c^2 \rangle$ is abelian, it follows that $[G,G] = \langle c^2 \rangle$.

(b) An easy calculation shows that $c^n a c^{-n} = c^{2n}a$. Since every element of G has the form c^n or $c^n a$ for some $n \in \mathbb{Z}$, it follows that $a^G = \{c^{2n}a \mid n \in \mathbb{Z}\}$. Using the fact that $b = ac$, we also see that $b^G = \{c^{2n+1}a \mid n \in \mathbb{Z}\}$. □

Let π be the outer automorphism of order 2 which interchanges the elements a and b of $G = \langle a \rangle * \langle b \rangle$.

LEMMA 1.4.2. $\mathrm{Aut}(G) = \langle \pi, i_a \rangle$.

PROOF. First notice that $\pi i_a \pi^{-1} = i_b \in \langle \pi, i_a \rangle$ and so $\mathrm{Inn}\, G \leqslant \langle \pi, i_a \rangle$. Now let $\varphi \in \mathrm{Aut}(G)$ be an arbitrary automorphism. Then φ must permute the set $\Delta = \{a^G, b^G\}$ of conjugacy classes of involutions of G. Replacing φ by $\pi\varphi$ if necessary, we can suppose that $\varphi\left[a^G\right] = a^G$. Thus $\varphi(a) = c^{2n}a$ for some $n \in \mathbb{Z}$. Let $g = c^{-n}$ and $\psi = i_g \varphi$. Then $\psi(a) = a$. Since C is a characteristic subgroup of G, we must have that $\psi[C] = C$. Hence either $\psi(c) = c$ or $\psi(c) = c^{-1}$. In the former case, $\psi = id_G$; and in the latter case, $\psi = i_a$. Thus $\varphi \in \langle \pi, i_a \rangle$. □

We have now shown that $\mathrm{Aut}(G)$ is also an infinite dihedral group. It follows that G_n is a proper subgroup of G_{n+1} for all $n \in \omega$. In order to understand the group $G_\omega = \bigcup_{n\in\omega} G_n$, we must first describe the embedding, $\mathrm{Inn}(G) < \mathrm{Aut}(G)$, in a little more detail.

LEMMA 1.4.3. (a) $[\mathrm{Aut}(G) : \mathrm{Inn}(G)] = 2$.

(b) $(i_a\pi)^2 = i_c$.

(c) $\langle i_c \rangle = [\mathrm{Aut}(G), \mathrm{Aut}(G)]$.

(d) *Any two involutions of* $\mathrm{Inn}(G)$ *are conjugate in* $\mathrm{Aut}(G)$.

PROOF. (a) We have already noted that each automorphism of G permutes the elements of the set $\Delta = \{a^G, b^G\}$. This yields a surjective homomorphism

$\theta : \mathrm{Aut}(G) \to \mathrm{Sym}(\Delta)$. The proof of Lemma 1.4.2 shows that $\ker\theta = \mathrm{Inn}\, G$. Hence $[\mathrm{Aut}(G) : \mathrm{Inn}(G)] = |\mathrm{Sym}(\Delta)| = 2$.

(b) It is easily checked that

$$(i_a\pi i_a\pi)(a) = ababa = (ab)a(ab)^{-1} = cac^{-1}$$

and that

$$(i_a\pi i_a\pi)(b) = aba = (ab)b(ab)^{-1} = cbc^{-1}.$$

Thus $(i_a\pi)^2 = i_c$.

(c) This is an immediate consequence of Lemmas 1.4.1(a), 1.4.2 and 1.4.3(b).

(d) This follows from the fact that $\pi i_a \pi^{-1} = i_b$. □

Hence for each $n \in \omega$, the nth group in the automorphism tower of G is $G_n = C_n \rtimes \langle a \rangle$, where

(1) $C_n = \langle c_n \rangle$ is infinite cyclic;
(2) $ac_na^{-1} = c_n^{-1}$; and
(3) $c_{n+1}^2 = c_n$.

Thus $G_\omega = C_\omega \rtimes \langle a \rangle$, where

(i) $C_\omega = \bigcup_{n\in\omega} C_n$ is isomorphic to the additive group of dyadic rationals $\mathbb{Z}[1/2] = \{m/2^n \mid m \in \mathbb{Z}, n \in \mathbb{N}\}$; and
(ii) the automorphism induced by a on C_ω corresponds to the automorphism $\sigma \in \mathrm{Aut}(\mathbb{Z}[1/2])$ such that $\sigma(x) = -x$ for all $x \in \mathbb{Z}[1/2]$.

In order to simplify notation, we will identify $G_\omega = C_\omega \rtimes \langle a \rangle$ with $\mathbb{Z}[1/2] \rtimes \langle \sigma \rangle$.

LEMMA 1.4.4. (a) *G_ω has a single conjugacy class of involutions.*
(b) *$[G_\omega, G_\omega] = \mathbb{Z}[1/2]$.*

PROOF. (a) By Lemma 1.4.3(d), any two involutions of G_n are conjugate in G_{n+1}. Thus G_ω has a single conjugacy class of involutions.

(b) By Lemma 1.4.3(c), $[G_{n+1}, G_{n+1}] = C_n$. Hence $[G_\omega, G_\omega] = C_\omega = \mathbb{Z}[1/2]$. □

Since $\mathbb{Z}[1/2] = [G_\omega, G_\omega]$ is a characteristic subgroup of G_ω, each $\varphi \in \mathrm{Aut}(G_\omega)$ induces an automorphism of $\mathbb{Z}[1/2]$. Now it is easy to see that any automorphism of $\mathbb{Z}[1/2]$ is just multiplication by some element $u \in U = \{\pm 2^n \mid n \in \mathbb{Z}\}$, the group of multiplicative units of the ring of dyadic rationals. Let $\mathrm{Hol}\,\mathbb{Z}[1/2]$ be the *holomorph* of $\mathbb{Z}[1/2]$; i.e. the semidirect product $\mathbb{Z}[1/2] \rtimes \operatorname{Aut}(\mathbb{Z}[1/2])$, where $\theta x \theta^{-1} = \theta(x)$ for all $\theta \in \operatorname{Aut}(\mathbb{Z}[1/2])$ and $x \in \mathbb{Z}[1/2]$. Since $\operatorname{Aut}(\mathbb{Z}[1/2]) \simeq U$ is abelian, it follows that

$$G_\omega = \mathbb{Z}[1/2] \rtimes \langle \sigma \rangle \trianglelefteq \operatorname{Hol} \mathbb{Z}[1/2]$$

and so each element of $\operatorname{Hol} \mathbb{Z}[1/2]$ induces an automorphism of G_ω via conjugation. We claim that every automorphism of G_ω arises in this fashion, so that $G_{\omega+1} = \operatorname{Hol} \mathbb{Z}[1/2]$. To see this, let $\varphi \in \operatorname{Aut}(G_\omega)$ be an arbitrary automorphism. By Lemma 1.4.4(a), the involutions σ and $\varphi(\sigma)$ are conjugate in G_ω. Thus after replacing φ by $i_g \varphi$ for a suitably chosen $g \in G_\omega$, we can suppose that $\varphi(\sigma) = \sigma$. Let $\theta = \varphi \restriction \mathbb{Z}[1/2] \in \operatorname{Aut}(\mathbb{Z}[1/2])$. Then $\theta x \theta^{-1} = \theta(x) = \varphi(x)$ for all $x \in \mathbb{Z}[1/2]$ and $\theta \sigma \theta^{-1} = \sigma = \varphi(\sigma)$. Thus $\theta y \theta^{-1} = \varphi(y)$ for all $y \in G_\omega$. This completes the proof that $G_{\omega+1} = \operatorname{Hol} \mathbb{Z}[1/2]$.

Finally we shall prove that $G_{\omega+1}$ is a complete group.

Lemma 1.4.5. (a) $[G_{\omega+1}, G_{\omega+1}] = \mathbb{Z}[1/2]$.

(b) *G_ω is a characteristic subgroup of $G_{\omega+1}$.*

Proof. (a) Lemma 1.4.4(b) implies that $\mathbb{Z}[1/2] \leqslant [G_{\omega+1}, G_{\omega+1}]$. Since the quotient $G_{\omega+1}/\mathbb{Z}[1/2]$ is abelian, it follows that $[G_{\omega+1}, G_{\omega+1}] = \mathbb{Z}[1/2]$.

(b) We have just seen that $\mathbb{Z}[1/2]$ is a characteristic subgroup of $G_{\omega+1}$. Let $\psi \in \operatorname{Aut}(\mathbb{Z}[1/2])$ be the automorphism such that $\psi(x) = 2x$ for all $x \in \mathbb{Z}[1/2]$. Then

$$G_{\omega+1}/\mathbb{Z}[1/2] \simeq \operatorname{Aut}(\mathbb{Z}[1/2]) = \langle \sigma \rangle \oplus \langle \psi \rangle$$

is the direct sum of the cyclic group $\langle \sigma \rangle$ of order 2 and the infinite cyclic group $\langle \psi \rangle$. Clearly $\langle \sigma \rangle$ is a characteristic subgroup of $\langle \sigma \rangle \oplus \langle \psi \rangle$. It follows that $G_\omega = \mathbb{Z}[1/2] \rtimes \langle \sigma \rangle$ is a characteristic subgroup of $G_{\omega+1}$. □

Since G_ω is a characteristic subgroup of $G_{\omega+1}$, Theorem 1.2.7 yields that $G_{\omega+1}$ is a complete group. Summing up, we have now proved the following result.

Theorem 1.4.6 (Hulse [**17**]). *The automorphism tower of the infinite dihedral group D_∞ terminates after exactly $\omega + 1$ steps.*

□

1.5. Notes

The notion of the automorphism tower of a centreless group dates back to at least 1937, when Zassenhaus [**53**] asked whether the automorphism tower of every finite centreless group terminates after finitely many steps. In his classical paper [**52**], Wielandt proved that this is indeed the case. One of the key ingredients of Wielandt's proof was his result that if

$$G = G_0 \trianglelefteq G_1 \trianglelefteq \cdots \trianglelefteq G_n \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq \cdots$$

is the automorphism tower of the centreless group G, then $C_{G_n}(G_0) = 1$ for all integers $n \geq 0$. This was later generalised by Hulse [**17**] to the statement that $C_{G_\alpha}(G_0) = 1$ for all ordinals $\alpha \geq 0$. The analysis of the automorphism tower of the infinite dihedral group D_∞ in Section 1.4 is also due to Hulse [**17**].

Plotkin [**36**] contains an interesting account of the automorphism tower problem, including the results of Section 1.2 and a proof of Wielandt's Theorem.

CHAPTER 2

Wielandt's Theorem

In this chapter, we shall present a proof of Wielandt's theorem that the automorphism tower of a finite centreless group terminates after finitely many steps. Wielandt's theorem is by far the deepest result in this book. Its proof involves an intricate analysis of the subnormal subgroups of a finite centreless group, together with some very ingenious commutator calculations. Section 2.1 contains an outline of the proof, modulo two technical results which are proved in Sections 2.2 and 2.3. The reader should not feel too discouraged if he finds some of the material in this chapter rather difficult, especially that in Section 2.3. Nothing from this chapter will be used in the later chapters, and so the reader will not experience any disadvantage in understanding the rest of the book if he simply skips the difficult parts of the proof. I have included a complete proof of Wielandt's theorem for two reasons. Firstly, it provides a striking contrast to the extremely simple proof of the automorphism theorem for infinite groups, which will be proved in Chapter 3. But more importantly, sixty years after its original publication, Wielandt's theorem remains the high point in the study of automorphism towers.

2.1. Automorphism towers of finite groups

In this section, we shall present an outline of the proof of Wielandt's theorem on the automorphism towers of finite centreless groups.

THEOREM 2.1.1 (Wielandt [**52**]). *If G is a finite centreless group, then the automorphism tower of G terminates after finitely many steps.*

Theorem 2.1.1 is an easy consequence of the following important theorem on finite subnormal subgroups.

DEFINITION 2.1.2. Let H be a subgroup of the group G. Then H is said to be a *subnormal subgroup* of G, written H sn G, if there exists a finite series of subgroups $H_0, H_1, \dots, H_n$ such that

$$H = H_0 \trianglelefteq H_1 \trianglelefteq \cdots \trianglelefteq H_n = G.$$

The *defect* $s(G : H)$ of the subnormal subgroup H is the least integer $n \geq 0$ for which such a series exists.

In particular, $s(G : H) = 0$ iff $H = G$, and $s(G : H) = 1$ iff H is a proper normal subgroup of G.

THEOREM 2.1.3 (Wielandt [**52**]). *There exists a function $f : \omega \to \omega$ such that whenever H is a finite subnormal subgroup of a group G with $C_G(H) = 1$, then $|G| \leq f(|H|)$.*

Most of this chapter will be devoted to the proof of Theorem 2.1.3. But first we shall show how to complete the proof of Theorem 2.1.1. So let G be a finite centreless group and let

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \cdots$$

be the automorphism tower of G. Notice that for each $n < \omega$, G is a subnormal subgroup of G_n such that $C_{G_n}(G) = 1$. So Theorem 2.1.3 says that there is an integer $\ell = f(|G|)$ such that $|G_n| \leq \ell$ for all $n < \omega$. Consequently there is an integer n such that $G_n = G_{n+1}$. This completes the proof of Theorem 2.1.1.

Of course, this argument does nothing to clear up the mystery of why Theorem 2.1.1 is true. So in the rest of this section, we shall present the main points of the proof of Theorem 2.1.3. (Towards the end of the proof, we shall appeal to two technical results on finite subnormal subgroups. These will be proved in Sections 2.2 and 2.3.) In order to understand the main difficulty, let us reconsider the statement of Theorem 2.1.1. Suppose that H is a finite subnormal subgroup of the group G and that $C_G(H) = 1$. Let

$$H = H_0 \trianglelefteq H_1 \trianglelefteq \cdots \trianglelefteq H_n = G$$

be a finite series from H to G. For each $0 \leq \ell < n - 1$, we can define a homomorphism $\varphi_\ell : H_{\ell+1} \to \mathrm{Aut}(H_\ell)$ by $\varphi_\ell(g) = i_g \restriction H_\ell$. Since $C_{H_{\ell+1}}(H_\ell) \leqslant C_G(H) = 1$, it follows that each φ_ℓ is an embedding and so $|H_{\ell+1}| \leq |H_\ell|!$. Thus it is obvious that

G is a finite group and that $|G|$ can be bounded in terms of $|H|$ and n. However, it is far from obvious that n can be bounded in terms of $|H|$.

Our strategy in the proof of Theorem 2.1.3 will be to first construct a strictly increasing series of characteristic subgroups in H

$$1 = S_0 < S_1 < \cdots < S_t = H$$

and then construct a corresponding series of characteristic subgroups in G

$$1 = R_0 \leqslant R_1 \leqslant \cdots \leqslant R_t \leqslant G$$

such that $S_i \leqslant R_i$ for each $0 \leq i \leq t$. In particular, $H = S_t \leqslant R_t$.

Claim 2.1.4. *It is enough to bound $|R_t|$ in terms of $|H|$.*

Proof. Since R_t is a characteristic subgroup of G, we can define a homomorphism $\varphi : G \to \mathrm{Aut}(R_t)$ by $\varphi(g) = i_g \restriction R_t$. Furthermore, since $H \leqslant R_t$, it follows that $C_G(R_t) \leqslant C_G(H) = 1$ and so φ is an embedding. Hence

$$|G| \leq |\mathrm{Aut}(R_t)| \leq |R_t|!.$$

□

Also notice that $t \leq |H|$. Hence it is enough to bound $|R_{i+1}|$ inductively in terms of $|R_i|$ and $|H|$.

Before we can define the series $\langle S_i \mid 0 \leq i \leq t\rangle$, we need to introduce some important characteristic subgroups. Let G be any group and let p be a prime. Suppose that H and K are normal p-subgroups of G. Then $HK/K \simeq H/H \cap K$ is a p-group and hence HK is also a normal p-subgroup of G. It follows that G has a unique maximal normal p-subgroup, namely the subgroup generated by all the normal p-subgroups of G.

Definition 2.1.5. If G is a group and p is a prime, then $O_p(G)$ is the unique maximal normal p-subgroup of G.

Clearly $O_p(G)$ is a characteristic subgroup of G. If G is a finite group, then Theorem 2.1.7 gives a useful criterion for the existence of a prime p such that $O_p(G) \neq 1$.

LEMMA 2.1.6. *Let G be a group and let p be a prime. If P is a subnormal p-subgroup of G, then $P \leqslant O_p(G)$.*

PROOF. We shall prove the result by induction on the defect $s = s(G : P)$ of P. If $s \leq 1$, then $P \trianglelefteq G$ and so $P \leqslant O_p(G)$. Suppose that $s > 1$ and let

$$P = H_0 \trianglelefteq \cdots \trianglelefteq H_{s-1} \trianglelefteq H_s = G$$

be a series of finite length from P to G. Then $s(H_{s-1} : P) = s-1$ and by induction hypothesis, we have that $P \leqslant O_p(H_{s-1})$. But $O_p(H_{s-1})$ is a characteristic subgroup of H_{s-1} and hence $O_p(H_{s-1})$ is a normal subgroup of $H_s = G$. This means that $O_p(H_{s-1}) \leqslant O_p(G)$ and so $P \leqslant O_p(G)$. $\square$

THEOREM 2.1.7. *If G is a finite group, then the following statements are equivalent.*

(a) *G has a nontrivial normal abelian subgroup.*
(b) *G has a nontrivial subnormal abelian subgroup.*
(c) *There exists a prime p such that $O_p(G) \neq 1$.*

PROOF. (a) $\Rightarrow$ (b). This is clear.
(b) $\Rightarrow$ (c). Suppose that A is a nontrivial subnormal abelian subgroup of G. Let p be any prime divisor of $|A|$. Then A contains a cyclic subgroup P of order p. Since P is a subnormal subgroup of G, Lemma 2.1.6 implies that $O_p(G) \neq 1$.
(c) $\Rightarrow$ (a). Suppose that there exists a prime p such that $P = O_p(G) \neq 1$. As P is a nontrivial finite p-group, $Z(P) \neq 1$. Since $Z(P)$ is a characteristic subgroup of $P = O_p(G)$, it follows that $Z(P)$ is a normal subgroup of G. Thus $Z(P)$ is a nontrivial normal abelian subgroup of G. $\square$

DEFINITION 2.1.8. A finite group G is said to be *semisimple* if G has no nontrivial normal abelian subgroups.

Thus if G is a finite semisimple group, then every minimal nontrivial subnormal subgroup of G is simple nonabelian.

DEFINITION 2.1.9. If G is a group, then $\sigma(G)$ is the subgroup generated by the simple nonabelian subnormal subgroups of G. If G has no simple nonabelian subnormal subgroups, then we set $\sigma(G) = 1$.

Clearly $\sigma(G)$ is a characteristic subgroup of G.

THEOREM 2.1.10. *If G is any group, then $\sigma(G)$ is the internal direct sum of the simple nonabelian subnormal subgroups of G.*

Theorem 2.1.10 is a straightforward consequence of the following theorem, which we shall prove in Section 2.2.

THEOREM 2.1.11. *If H is a simple nonabelian subnormal subgroup of the group G, then H normalises every subnormal subgroup of G.*

PROOF OF THEOREM 2.1.10. Let $\{S_i \mid i \in I\}$ be the set of simple nonabelian subnormal subgroups of G. Let i, j be distinct elements of I. By Theorem 2.1.11, S_i and S_j normalise each other. Hence $S_i \cap S_j$ is a normal subgroup of S_i and so $S_i \cap S_j = 1$. Applying Lemma 1.2.3(b), we obtain that

$$[S_i, S_j] \leqslant S_i \cap S_j = 1.$$

It follows that $\langle S_j \mid j \in I \smallsetminus \{i\}\rangle$ centralises S_i and hence

$$S_i \cap \langle S_j \mid j \in I \smallsetminus \{i\}\rangle = 1$$

for all $i \in I$. Thus $\sigma(G) = \langle S_i \mid i \in I\rangle$ is the internal direct sum of the subgroups $\{S_i \mid i \in I\}$. □

Now we are ready to begin the proof of Theorem 2.1.3. Suppose that H is a finite subnormal subgroup of the group G such that $C_G(H) = 1$. Then G is also a finite group and we shall find a bound for $|G|$ which depends only on the integer $|H|$.

First we construct a strictly increasing series of characteristic subgroups in H

$$1 = S_0 < S_1 < \cdots < S_t = H$$

as follows. Suppose inductively that S_i has been constructed. If H/S_i is semisimple, then $S_{i+1}/S_i = \sigma(H/S_i)$ is the product of all the simple nonabelian subnormal subgroups of H/S_i. If H/S_i is not semisimple, then there exists a prime p_i such that $O_{p_i}(H/S_i) \neq 1$ and we define $S_{i+1}/S_i = O_{p_i}(H/S_i)$.

Next we construct a corresponding series of characteristic subgroups in G

$$1 = R_0 \leqslant R_1 \leqslant \cdots \leqslant R_t \leqslant G$$

as follows. If H/S_i is semisimple, then $R_{i+1}/R_i = \sigma(G/R_i)$ is the product of all the simple nonabelian subnormal subgroups of G/R_i. Otherwise, $R_{i+1}/R_i = O_{p_i}(G/R_i)$, where p_i is the prime such that $S_{i+1}/S_i = O_{p_i}(H/S_i)$.

Lemma 2.1.12. *$S_i \leqslant R_i$ for each $0 \leq i \leq t$.*

Proof. We argue by induction on $i \geq 0$. The result is clearly true when $i = 0$. Suppose that $i < t$ and that $S_i \leqslant R_i$. If $S_{i+1}R_i/R_i = 1$, then $S_{i+1} \leqslant R_i \leqslant R_{i+1}$. So we can assume that $S_{i+1}R_i/R_i \neq 1$. First notice that

$$S_{i+1}R_i/R_i \simeq S_{i+1}/S_{i+1} \cap R_i$$

and so $S_{i+1}R_i/R_i$ is a homomorphic image of S_{i+1}/S_i. Thus $S_{i+1}R_i/R_i$ is either a product of simple nonabelian groups or else a p_i-group, depending on whether H/S_i is semisimple or not. Next note that

$$S_{i+1}R_i/R_i \trianglelefteq \cdots \trianglelefteq S_tR_i/R_i = HR_i/R_i;$$

and that HR_i/R_i sn G/R_i, because H sn G. Hence $S_{i+1}R_i/R_i$ is a subnormal subgroup of G/R_i.

Thus if H/S_i is semisimple, then $S_{i+1}R_i/R_i$ is a product of simple nonabelian subnormal subgroups of G/R_i and so $S_{i+1}R_i/R_i \leqslant R_{i+1}/R_i = \sigma(G/R_i)$. And if $S_{i+1}/S_i = O_{p_i}(H/S_i)$, then $S_{i+1}R_i/R_i$ is a subnormal p_i-subgroup of G/R_i; and so by Lemma 2.1.6,

$$S_{i+1}R_i/R_i \leqslant O_{p_i}(G/R_i) = R_{i+1}/R_i.$$

In both cases, we have obtained that $S_{i+1} \leqslant R_{i+1}$. □

As we explained earlier in Claim 2.1.4, in order to bound $|G|$ in terms of $|H|$, it is enough to bound $|R_t|$ in terms of $|H|$. And since $t \leq |H|$, we need only bound $|R_{i+1}|$ inductively in terms of $|R_i|$ and $|H|$. The following two results will enable us to accomplish this.

Lemma 2.1.13. *Suppose that H is a subnormal subgroup of the finite group G and that $C_G(H) = 1$. If $N \trianglelefteq G$ and $R/N = \sigma(G/N)$, then $|R| \leq |HN|!$.*

Proof. We can suppose that $\sigma(G/N) \neq 1$. So R/N is the product of all the simple nonabelian subnormal subgroups of G/N. By Theorem 2.1.11, each simple nonabelian subnormal subgroup of G/N normalises the subnormal subgroup HN/N. Hence R/N normalises HN/N and this implies that R normalises HN. Let $\varphi : R \to \mathrm{Aut}(HN)$ be the homomorphism defined by $\varphi(g) = i_g \restriction HN$. Since $C_R(HN) \leqslant C_G(H) = 1$, it follows that φ is an embedding. Thus

$$|R| \leq |\mathrm{Aut}(HN)| \leq |HN|!.$$

$\square$

The next result will be proved in Section 2.3.

LEMMA 2.1.14. *Let p be a prime. Suppose that H is a subnormal subgroup of the finite group G and that $C_G(H) = 1$. If $N \trianglelefteq G$ and $R/N = O_p(G/N)$, then $|R| \leq |N|.|HN|!$.*

Let $m_i = |R_i|$ for each $0 \leq i \leq t$. Then $m_0 = 1$, and we can bound m_{i+1} in terms of m_i and $h = |H|$ as follows. If H/S_i is semisimple, then $R_{i+1}/R_i = \sigma(G/R_i)$ and so, applying Lemma 2.1.13, we obtain that

$$m_{i+1} = |R_{i+1}| \leq (hm_i)!.$$

If H/S_i is not semisimple, then $R_{i+1}/R_i = O_{p_i}(G/R_i)$ and Lemma 2.1.14 yields that

$$m_{i+1} = |R_{i+1}| \leq m_i(hm_i)!.$$

In both cases, $m_{i+1} \leq m_i(hm_i)!$. This completes the proof of Theorem 2.1.3.

2.2. Subnormal subgroups

In this section, we shall prove Theorem 2.1.11. But first we need to develop some of the basic theory of subnormal subgroups.

DEFINITION 2.2.1. If X is a nonempty subset of the group G, then the *normal closure* of X in G is the subgroup

$$X^G = \langle gXg^{-1} \mid g \in G \rangle.$$

Thus X^G is the smallest normal subgroup of G which contains X. Notice that if H is a subgroup of G, then H is a normal subgroup iff $H^G = H$. The following lemma characterises subnormal subgroups in terms of iterated normal closures.

Definition 2.2.2. If X is a nonempty subset of the group G, then the *ith normal closure* of X in G is the subgroup $X^{G,i}$ defined inductively by

(a) $X^{G,0} = G$.
(b) $X^{G,i+1} = X^{X^{G,i}}$.

An easy induction shows that $X \subseteq X^{G,i}$ for all $i \geq 0$ and hence $X^{G,i+1} \trianglelefteq X^{G,i}$.

Lemma 2.2.3. *Let H sn G and suppose that*

$$H = H_n \trianglelefteq H_{n-1} \trianglelefteq \cdots \trianglelefteq H_0 = G$$

is a series of finite length from H to G. Then $H^{G,i} \leqslant H_i$ for all $0 \leq i \leq n$ and so $H = H^{G,n}$.

Proof. We shall prove that $H^{G,i} \leqslant H_i$ by induction on $i \leq n$. Clearly the result is true for $i = 0$. Suppose that $i < n$ and that $H^{G,i} \leqslant H_i$. Since $H_{i+1} \trianglelefteq H_i$, we have that $H_{i+1}^{H_i} = H_{i+1}$. Thus

$$H^{G,i+1} = H^{H^{G,i}} \leqslant H_{i+1}^{H_i} = H_{i+1}.$$

□

Hence if H is a subnormal subgroup of G, then $s(G : H)$ is the least integer n such that $H^{G,n} = H$; and

$$H = H^{G,n} \trianglelefteq H^{G,n-1} \trianglelefteq \cdots \trianglelefteq H^{G,1} \trianglelefteq H^{G,0} = G$$

is a series of minimal length from H to G. Consequently, if $s(G : H) = n > 0$, then

$$s(H^G : H) = s(H^{G,1} : H) = n - 1.$$

This observation is often useful in arguments which proceed by induction on the defect $s = s(G : H)$. In the remainder of this chapter, we shall make repeated use of the following two results.

Theorem 2.2.4. *Suppose that H and K are subgroups of the group G and that $J = \langle H, K \rangle$. Then*

(a) $[H, K] \trianglelefteq J$*; and*
(b) $K^J = [H, K]K$.

PROOF. (a) Let h, $g \in H$ and $k \in K$. By Lemma 1.2.2(b),

$$[h,k]^g = [hg,k][g,k]^{-1} \in [H,K].$$

It follows that H normalises $[H,K]$. By Lemma 1.2.3(a), $[H,K] = [K,H]$. Hence K also normalises $[H,K]$ and so $[H,K] \trianglelefteq J$.

(b) We must show that $[H,K]K$ is the smallest normal subgroup of J which contains K. First we shall prove that $[H,K]K \trianglelefteq J$. By (a), $[H,K]^g = [H,K]$ for all $g \in J$. Hence if $g \in K$, then

$$([H,K]K)^g = [H,K]^g K^g = [H,K]K.$$

Thus K normalises $[H,K]K$. Now suppose that $g \in H$ and $k \in K$. Then

$$g^{-1}k^{-1}g = (g^{-1}k^{-1}gk)k^{-1} = [g,k]k^{-1} \in [H,K]K.$$

Hence if $g \in H$, then $K^g \leqslant [H,K]K$ and so

$$([H,K]K)^g = [H,K]^g K^g \leqslant [H,K]K.$$

This implies that H also normalises $[H,K]K$. Hence $[H,K]K \trianglelefteq J$.

Finally suppose that N is a normal subgroup of J such that $K \leqslant N \trianglelefteq J$. Then for each $h \in H$ and $k \in K$, we have that $[h,k] = (h^{-1}k^{-1}h)k \in N$. It follows that $[H,K] \leqslant N$ and so $[H,K]K \leqslant N$. □

LEMMA 2.2.5. *If K sn G and $H \leqslant G$, then $H \cap K$ sn H. Furthermore, $s(H : H \cap K) \leq s(G : K)$.*

PROOF. Let $n = s(G : K)$ and let

$$K = K_0 \trianglelefteq K_1 \trianglelefteq \cdots \trianglelefteq K_n = G$$

be a series of finite length from K to G. Intersecting this series with H, we obtain the series

$$H \cap K = H \cap K_0 \trianglelefteq \cdots \trianglelefteq H \cap K_n = H.$$

Thus $H \cap K$ is a subnormal subgroup of H and $s(H : H \cap K) \leq n = s(G : K)$. □

Theorem 2.1.11 is an easy consequence of the following result.

LEMMA 2.2.6. *Suppose that H and K are subnormal subgroups of the group G and that $H \cap K = 1$. If H is a simple nonabelian group, then $[H,K] = 1$.*

Proof. Let H be a simple nonabelian subnormal subgroup of G. Suppose that K is a subnormal subgroup of G such that $H \cap K = 1$ and $[H, K] \neq 1$. We suppose that K has been chosen so that $s = s(\langle H, K\rangle : H)$ is minimal. Let $J = \langle H, K\rangle$. First suppose that $s \leq 1$. Then $H \trianglelefteq J$. In particular, K normalises H and hence Lemma 1.2.3 implies that $[H, K] \leqslant H$. By Theorem 2.2.4(a), $[H, K] \trianglelefteq J$. It follows that $[H, K] \trianglelefteq H$. Since H is a simple group and $[H, K] \neq 1$, we must have that $[H, K] = H$. Applying Theorem 2.2.4(b), we find that $H \leqslant [H, K]K = K^J$ and hence $K^J = J$. But K is a subnormal subgroup of J and so Lemma 2.2.3 implies that $K = J$. But then $H = H \cap J = H \cap K = 1$, which is a contradiction.

Now suppose that $s > 1$. Then H is not a normal subgroup of J and so there exists an element $g \in K$ such that $H \neq H^g$. Note that H^g is also a subnormal subgroup of G and so $H \cap H^g$ is a subnormal subgroup of H. Since H is simple, we must have that $H \cap H^g = 1$. Remember that

$$s\left(H^J : H\right) = s(J : H) - 1 = s - 1.$$

Since $\langle H, H^g\rangle \leqslant H^J$, it follows that $s(\langle H, H^g\rangle : H) \leq s - 1$. By the minimality of s, we have that $[H, H^g] = 1$. Therefore if $h_1, h_2 \in H$, then

$$1 = [h_1, h_2^g] = [h_1, h_2[h_2, g]] = [h_1, [h_2, g]][h_1, h_2]^{[h_2, g]}.$$

(The last equality follows from Lemma 1.2.2(c).) By Theorem 2.2.4(a), H normalises $[H, K]$ and so $[H, [H, K]] \leqslant [H, K]$. Thus

$$[h_1, h_2]^{[h_2, g]} = [h_1, [h_2, g]]^{-1} \in [H, K]$$

and so $[h_1, h_2] \in [H, K]$ for all $h_1, h_2 \in H$. Thus $[H, H] \leqslant [H, K]$. Since H is a simple nonabelian group, it follows that $[H, H] = H$. Hence we have shown that $H \leqslant [H, K]$. Arguing as in the first paragraph of this proof, we now find that $K^J = J$; and again this leads to the contradiction that $H = 1$. □

Proof of Theorem 2.1.11. Suppose that H is a simple nonabelian subnormal subgroup of the group G and let K be an arbitrary subnormal subgroup of G. By Lemma 2.2.5, $H \cap K$ is a subnormal subgroup of H. Since H is simple, either $H \cap K = H$ or $H \cap K = 1$. In the former case, $H \leqslant K$; and in the latter case, Lemma 2.2.6 yields that $[H, K] = 1$. So in both cases, we find that H normalises K. □

2.3. Finite p-groups

In this section, we shall prove Lemma 2.1.14. We shall make use of the following two results. The first result is well-known and a proof can be found in any of the standard textbooks in group theory. (For example, see Section 2.1 of Suzuki [**48**].)

THEOREM 2.3.1. *If P is a finite p-group, then P is nilpotent. Hence if N is a nontrivial normal subgroup of P, then $N \cap Z(P) \neq 1$.*

□

The second result is more technical; and requires some preliminary explanation. Let H be a finite group and let p be a prime. Suppose that N_1, N_2 are normal subgroups of H such that H/N_1 and H/N_2 are p-groups. Consider the homomorphism

$$\pi : H \to H/N_1 \times H/N_2$$

defined by $\pi(h) = (hN_1, hN_2)$. Then $\ker \pi = N_1 \cap N_2$ and so $H/N_1 \cap N_2$ is isomorphic to a subgroup of $H/N_1 \times H/N_2$. In particular, $H/N_1 \cap N_2$ is also a p-group. This implies that there exists a smallest normal subgroup N of H such that H/N is a p-group.

DEFINITION 2.3.2. Let H be a finite group and let p be a prime. Then $O^p(H)$ is the smallest normal subgroup of H such that $H/O^p(H)$ is a p-group.

PROPOSITION 2.3.3. *Suppose that H is a subnormal subgroup of the finite group G and that p be a prime. Then $O_p(G)$ normalises $O^p(H)$.*

Before proving Proposition 2.3.3, we shall complete the proof of Lemma 2.1.14. Suppose that H is a subnormal subgroup of the finite group G and that $C_G(H) = 1$. Let $N \trianglelefteq G$ and let $R/N = O_p(G/N)$.

CLAIM 2.3.4. $O^p(HN/N) = O^p(H)N/N$.

PROOF. Let M be the subgroup such that $N \trianglelefteq M$ and $M/N = O^p(HN/N)$. Since

$$(HN/N) \,/\, (O^p(H)N/N) \simeq HN/O^p(H)N$$

is a p-group, it follows that $M \leqslant O^p(H)N$. Now consider the homomorphism

$$\psi : H \to HN/M$$

defined by $\psi(h) = hM$. Then $\ker\psi = H \cap M$ and so $H/H \cap M$ is isomorphic to a subgroup of the p-group HN/M. This implies that $O^p(H) \leqslant H \cap M$ and hence $O^p(H)N \leqslant M$. Thus $M = O^p(H)N$. $\square$

It is easily checked that HN/N is a subnormal subgroup of G/N. Thus Proposition 2.3.3 yields that $R/N = O_p(G/N)$ normalises $O^p(H)N/N = O^p(HN/N)$; and this implies that R normalises $O^p(H)N$. Let $\varphi : R \to \mathrm{Aut}(O^p(H)N)$ be the homomorphism defined by $\varphi(g) = i_g \upharpoonright O^p(H)N$ and let $C = C_R(O^p(H)N)$. Then $\ker\varphi = C$ and so

$$[R : C] \leq |\mathrm{Aut}(O^p(H)N)| \leq |HN|!.$$

We shall prove that $C \leqslant N$. This implies that

$$|R| \leq |C|.|HN|! \leq |N|.|HN|!,$$

as required.

First notice that H normalises $O^p(H)$, N and R; and hence H normalises $C = C_R(O^p(H)N)$. Now let P be a Sylow p-subgroup of H and let Q be a Sylow p-subgroup of HC such that $P \leqslant Q$. Let $T = C \cap Q$. Clearly P normalises T.

Claim 2.3.5. $T = 1$.

Proof. Suppose that $T \neq 1$. We shall derive a contradiction by finding an element $1 \neq z \in T \cap C_G(H)$. First remember that if S is a Sylow p-subgroup of a group K and $M \trianglelefteq K$, then SM/M is a Sylow p-subgroup of K/M. In particular, since $H/O^p(H)$ is a p-group, we must have that $H/O^p(H) = PO^p(H)/O^p(H)$ and so $H = PO^p(H)$.

Since P normalises T, it follows that PT is a p-group. By Theorem 2.3.1, there exists an element $1 \neq z \in T \cap Z(PT)$. But since $z \in C \leqslant C_G(O^p(H))$ and $H = PO^p(H)$, this means that $z \in C_G(H)$. $\square$

Now recall that if S is a Sylow p-subgroup of a group K and $M \trianglelefteq K$, then $M \cap S$ is a Sylow p-subgroup of M. In particular, since $C \trianglelefteq HC$, we have that $T = C \cap Q$

is a Sylow p-subgroup of C. Thus p does not divide $|C|$. Since $R/N = O_p(G/N)$ is a p-group, we must have that $C \leqslant N$. This completes the proof of Lemma 2.1.14.

The proof of Proposition 2.3.3 will involve some intricate commutator calculations. In particular, we shall make use of the following characterisation of $H^{G,i}$ in terms of iterated commutator subgroups.

Definition 2.3.6. If $n \geq 2$ and $X_0,\ X_1, \ldots, X_n$ are nonempty subsets of the group G, then we inductively define

$$[X_0, \ldots, X_n] = \left[[X_0, \ldots, X_{n-1}], X_n\right].$$

For each $n \geq 1$, we also define

$$[X,\, {}_nY] = [X, \underbrace{Y, \ldots, Y}_{n}].$$

Proposition 2.3.7. *If H is a subgroup of the group G, then $H^{G,i} = [G,\, {}_iH]\, H$ for each $i \geq 1$.*

Proof. It is easily checked that H normalises $[G,\, {}_iH]\, H$ for each $i \geq 1$. Thus $\langle [G,\, {}_iH], H\rangle = [G,\, {}_iH]\, H$. By Theorem 2.2.4, the proposition holds when $i = 1$. Assuming that the result holds for $i \geq 1$, then applying Theorem 2.2.4 once again, we obtain that

$$H^{G,i+1} = H^{H^{G,i}} = H^{[G,\, {}_iH]H} = \left[[G,\, {}_iH], H\right] H = [G,\, {}_{i+1}H]\, H.$$

□

We shall also make use of the following result.

Lemma 2.3.8. *Suppose that H and K are subgroups of the group G, and that K normalises H. Then $[HK, K] = [H, K][K, K]$.*

Proof. First note that

$$[H, K][K, K] \leqslant [HK, K][HK, K] = [HK, K].$$

Now let $h \in H$ and $x,\, y \in K$. Applying Lemma 1.2.2(b), we obtain that

$$[hx, y] = [h, y]^x[x, y] = [h^x, y^x][x, y] \in [H, K][K, K].$$

Thus $[HK, K] \leqslant [H, K][K, K]$. □

Finally we are ready to begin the proof of Proposition 2.3.3. So suppose that H is a subnormal subgroup of the finite group G and that p is a prime.

CLAIM 2.3.9. $[O_p(G), O^p(H)] = [O_p(G), O^p(H), O^p(H)]$.

Before proving Claim 2.3.9, we shall show how to complete the proof of Proposition 2.3.3. Let $T = \langle O_p(G), O^p(H)\rangle = O_p(G)O^p(H)$. Since $O^p(H) \trianglelefteq H$ and H sn G, it follows that $O^p(H)$ sn G. Hence $O^p(H)$ sn T. Applying Proposition 2.3.7 and Lemma 2.3.8, we see that

$$\begin{aligned} O^p(H)^{T,1} &= [T, O^p(H)]O^p(H) \\ &= [O_p(G)O^p(H), O^p(H)]O^p(H) \\ &= [O_p(G), O^p(H)]O^p(H). \end{aligned}$$

Similarly we have that $O^p(H)^{T,2} = [O_p(G), O^p(H), O^p(H)]O^p(H)$. Applying Claim 2.3.9, we obtain that $O^p(H)^{T,1} = O^p(H)^{T,2}$ and so $O^p(H)^{T,1} = O^p(H)^{T,i}$ for all $i \geq 1$. Since $O^p(H)$ sn T, Lemma 2.2.3 implies that there exists an integer $n \geq 1$ such that $O^p(H) = O^p(H)^{T,n}$. But this means that

$$O^p(H) = O^p(H)^{T,n} = O^p(H)^{T,1} = O^p(H)^T$$

and so $O_p(G)$ normalises $O^p(H)$.

PROOF OF CLAIM 2.3.9. Clearly both $O_p(G)$ and $[O_p(G), O^p(H)]$ are normalised by $O^p(H)$. Hence $[O_p(G), O^p(H)] \leqslant O_p(G)$ and

$$[O_p(G), O^p(H), O^p(H)] = [[O_p(G), O^p(H)], O^p(H)] \leqslant [O_p(G), O^p(H)].$$

The proof of the opposite inclusion is more involved. First note that, by Theorem 2.2.4, $[O_p(G), O^p(H), O^p(H)] \trianglelefteq [O_p(G), O^p(H)]$. Fix some element $g \in O_p(G)$ and let x, $y \in O^p(H)$. Using the definition of the commutator, it is easily verified that

$$[xy, g] = [x, g]\,[[x, g]y]\,[y, g].$$

Thus

$$[xy, g] = [x, g][y, g] \mod [O_p(G), O^p(H), O^p(H)],$$

since $[[x, g], y] = [[g, x]^{-1}, y] \in [O_p(G), O^p(H), O^p(H)]$. Thus we can define a homomorphism

$$\varphi : O^p(H) \to [O_p(G), O^p(H)] \,/\, [O_p(G), O^p(H), O^p(H)]$$

by

$$\varphi(x) = [x, g]\,[O_p(G), O^p(H), O^p(H)]\,.$$

Let $K = \ker\varphi$. As $[O_p(G), O^p(H)] \leqslant O_p(G)$, it follows that $O^p(H)/K$ is a p-group. Let $L = \bigcap_{h\in H} K^h$. Then L is a normal subgroup of H such that $L \leqslant K \leqslant O^p(H)$. Let $H = \{h_1, \ldots, h_t\}$. Then we can define a homomorphism

$$\psi : O^p(H) \to O^p(H)/K^{h_1} \times \cdots \times O^p(H)/K^{h_t}$$

by $\psi(x) = \left(xK^{h_1}, \ldots, xK^{h_t}\right)$. Clearly $\ker\psi = L$ and so $O^p(H)/L$ is a p-group. This implies that H/L is also a p-group and so $L = K = O^p(H)$. Consequently $[x, g] \in [O_p(G), O^p(H), O^p(H)]$ for all $x \in O^p(H)$ and $g \in O_p(G)$; and hence $[O_p(G), O^p(H)] \leqslant [O_p(G), O^p(H), O^p(H)]$. This completes the proof of Claim 2.3.9.

$\square$

2.4. Notes

My account of Wielandt's Theorem is closely based on that in Robinson's textbook [**39**]. The same material can also be found in Plotkin [**36**]. Alternative proofs of Theorem 2.1.3 can be found in Pettet [**35**] and Schenkman [**40**]. Pettet [**35**] also solves the problem of finding an explicit bound in the statement of Theorem 2.1.3 as follows. Suppose that H is a finite subnormal subgroup of the group G and that $C_G(H) = 1$. Let H^∞ be the *nilpotent residual* of H; i.e. the smallest normal subgroup of H such that H/H^∞ is nilpotent. Then

$$|G| \leq (|Z(H^\infty)|\,|\operatorname{Aut}(H^\infty)|)!.$$

CHAPTER 3

The Automorphism Tower Theorem

In this chapter, we shall present two proofs of the automorphism tower theorem. The first proof, which is given in Section 3.1, uses only the most elementary properties of regular cardinals; and shows that the automorphism tower of an infinite centreless group G terminates after at most $\left(2^{|G|}\right)^+$ steps. The second proof, which is given in Section 3.3, makes use of Fodor's Lemma on regressive functions defined on stationary sets; and shows that actually the automorphism tower of G always terminates after strictly less than $\left(2^{|G|}\right)^+$ steps. Most of Sections 3.1 and 3.3 will be devoted to a development of the basic theory of regular cardinals and stationary sets.

In Section 3.2, we shall prove that if $\omega \leq \lambda < \kappa$ and H is a centreless group of cardinality λ whose automorphism tower terminates after exactly $\alpha \geq 1$ steps, then there exists a centreless group G of cardinality κ whose automorphism tower also terminates after exactly α steps. Finally, in Section 3.4, we shall discuss the question of finding better upper bounds for the heights of automorphism towers.

3.1. The automorphism tower theorem

In this section, we shall prove that the automorphism tower of an arbitrary infinite centreless group G eventually terminates. As this theorem turns out to be an easy consequence of the result that every successor cardinal is regular, most of this section will be devoted to an account of some of the basic properties of regular cardinals. First we need to introduce the notion of cofinality.

Definition 3.1.1. Let α and β be ordinals. Then the mapping $f : \alpha \to \beta$ is *cofinal* iff $\operatorname{ran} f$ is unbounded in β.

Definition 3.1.2. Let β be an ordinal. Then the *cofinality* of β, written $\operatorname{cf}(\beta)$, is the least ordinal α such that there exists a cofinal mapping $f : \alpha \to \beta$.

For example, it is clear that $\mathrm{cf}(\omega) = \omega$. The notion of cofinality is only interesting when β is a limit ordinal. For suppose that β is a successor ordinal; say, $\beta = \gamma + 1 = \gamma \cup \{\gamma\}$. Then we can define a cofinal map $f : 1 \to \beta$ by $f(0) = \gamma$ and so $\mathrm{cf}(\beta) = 1$.

LEMMA 3.1.3. *There exists a cofinal map $f : \mathrm{cf}(\beta) \to \beta$ such that f is strictly increasing.*

PROOF. Clearly we can suppose that β is a limit ordinal. Let $g : \mathrm{cf}(\beta) \to \beta$ be any cofinal map, and define $f : \mathrm{cf}(\beta) \to \beta$ recursively on $\xi < \beta$ by

$$f(\xi) = \max\{g(\xi), \sup\{(f(\eta) + 1) \mid \eta < \xi\}\}.$$

Notice that if $\xi < \mathrm{cf}(\beta)$, then $\{f(\eta) + 1 \mid \eta < \xi\}$ must be bounded in β. Thus the map f is well-defined. Since $f(\xi) \geq g(\xi)$ for all $\xi < \mathrm{cf}(\beta)$, it follows that f is a cofinal map; and it is clear that f is strictly increasing. □

LEMMA 3.1.4. *If α is a limit ordinal and $f : \alpha \to \beta$ is a strictly increasing cofinal map, then $\mathrm{cf}(\alpha) = \mathrm{cf}(\beta)$.*

PROOF. Let $g : \mathrm{cf}(\alpha) \to \alpha$ be a cofinal map. Since f is strictly increasing and cofinal, it follows that $f \circ g : \mathrm{cf}(\alpha) \to \beta$ is also cofinal. Thus $\mathrm{cf}(\beta) \leq \mathrm{cf}(\alpha)$.

To see that $\mathrm{cf}(\alpha) \leq \mathrm{cf}(\beta)$, let $h : \mathrm{cf}(\beta) \to \beta$ be a cofinal map and define $\varphi : \mathrm{cf}(\beta) \to \alpha$ by

$$\varphi(\xi) = \text{ the least } \eta < \alpha \text{ such that } f(\eta) > h(\xi).$$

Then it is easily checked that $\varphi : \mathrm{cf}(\beta) \to \alpha$ is a cofinal map. □

For example, we can define a strictly increasing cofinal map $f : \omega \to \aleph_\omega$ by $f(n) = \aleph_n$. Hence $\mathrm{cf}(\aleph_\omega) = \omega$.

DEFINITION 3.1.5. (a) An infinite cardinal κ is *regular* if $\mathrm{cf}(\kappa) = \kappa$.
(b) An infinite cardinal κ is *singular* if $\mathrm{cf}(\kappa) < \kappa$.

If λ is an infinite cardinal, then the least cardinal greater than λ is denoted by λ^+. Thus $\aleph_\alpha^+ = \aleph_{\alpha+1}$ for every ordinal α. Cardinals of the form λ^+ are called successor cardinals.

THEOREM 3.1.6. *If λ is an infinite cardinal, then λ^+ is regular.*

PROOF. Suppose not. Then there exists a cofinal map $f : \lambda \to \lambda^+$. For each $\alpha < \lambda$, let $B_\alpha = \{\beta \mid \beta < f(\alpha)\}$. Then $\lambda^+ = \bigcup_{\alpha<\lambda} B_\alpha$. But $|B_\alpha| \leq \lambda$ for all $\alpha < \lambda$ and so $\left|\bigcup_{\alpha<\lambda} B_\alpha\right| \leq \lambda$, which is a contradiction. $\square$

Notice that if κ is a cardinal and X is a set of cardinality κ, then $\mathrm{cf}(\kappa)$ is the least cardinal θ such that we can express $X = \bigcup_{\alpha<\theta} X_\alpha$, where $|X_\alpha| < \kappa$ for each $\alpha < \theta$. We shall make use of this observation in the proof of the next theorem, which will be needed in Section 3.3.

THEOREM 3.1.7 (König). *If λ is an infinite cardinal, then* $\mathrm{cf}(2^\lambda) > \lambda$.

PROOF. Let $\kappa = \mathrm{cf}(2^\lambda)$ and suppose that $\kappa \leq \lambda$. Since $|\mathcal{P}(\lambda)| = 2^\lambda$, this implies that there exist sets $S_\alpha \subset \mathcal{P}(\lambda)$ for $\alpha < \kappa$ such that

(a) $|S_\alpha| < 2^\lambda$ for all $\alpha < \kappa$; and
(b) $\mathcal{P}(\lambda) = \bigcup_{\alpha<\kappa} S_\alpha$.

Since $\kappa \leq \lambda$, we can express $\lambda = \bigsqcup_{\alpha<\kappa} X_\alpha$ as the disjoint union of κ subsets such that $|X_\alpha| = \lambda$ for all $\alpha < \kappa$. For each $\alpha < \kappa$, let

$$P_\alpha = \{Y \in \mathcal{P}(X_\alpha) \mid \text{ There exists } Z \in S_\alpha \text{ such that } Y = Z \cap X_\alpha\}.$$

Then $|P_\alpha| \leq |S_\alpha| < 2^\lambda$ and hence there exists a subset $A_\alpha \in \mathcal{P}(X_\alpha) \smallsetminus P_\alpha$. Let $A = \bigsqcup_{\alpha<\kappa} A_\alpha$. Then $A \in S_\alpha$ for some $\alpha < \kappa$. But since $A_\alpha = A \cap X_\alpha$, this means that $A_\alpha \in P_\alpha$, which is a contradiction. $\square$

DEFINITION 3.1.8. Let κ be a regular uncountable cardinal.

(a) A subset $C \subseteq \kappa$ is *closed* in κ if for every limit ordinal $\delta < \kappa$ and every increasing δ-sequence

$$\alpha_0 < \alpha_1 < \cdots < \alpha_\xi < \cdots$$

of elements of C, we have that $\sup\{\alpha_\xi \mid \xi < \delta\} \in C$.

(b) A subset $C \subseteq \kappa$ is a *club* if C is both closed and unbounded in κ.

The notion of a club is of central importance in many applications of set theory to algebra; and in Section 3.3, we shall present a detailed account of clubs and the related notion of stationary sets. In this section, we shall only prove Theorem 3.1.10, which is the set-theoretic key to the automorphism tower theorem.

DEFINITION 3.1.9. Let X be a set and let κ be an infinite cardinal. We say that $X = \bigcup_{\alpha<\kappa} X_\alpha$ is a *smooth strictly increasing union* if the following conditions are satisfied.

(a) If $\alpha < \beta < \kappa$, then $X_\alpha \subsetneqq X_\beta$.
(b) If α is a limit ordinal such that $\alpha < \kappa$, then $X_\alpha = \bigcup_{\xi<\alpha} X_\xi$.

THEOREM 3.1.10. *Let κ be a regular uncountable cardinal. Suppose that $X = \bigcup_{\alpha<\kappa} X_\alpha$ and $Y = \bigcup_{\alpha<\kappa} Y_\alpha$ are smooth strictly increasing unions such that $|X_\alpha| < \kappa$ and $|Y_\alpha| < \kappa$ for all $\alpha < \kappa$. If $f : X \to Y$ is a bijection, then*

$$C = \{\alpha < \kappa \mid f[X_\alpha] = Y_\alpha\}$$

is a club of κ.

PROOF. It is clear that C is closed. The main point is to show that C is unbounded in κ.

CLAIM 3.1.11. *If $\alpha < \kappa$, then there exists $\beta < \kappa$ such that $f[X_\alpha] \subseteq Y_\beta$.*

PROOF OF CLAIM 3.1.11. Let $|X_\alpha| = \lambda$ and write $X_\alpha = \{x_\xi \mid \xi < \lambda\}$. Let $\varphi : \lambda \to \kappa$ be the map defined by

$$\varphi(\xi) = \text{ the least } \gamma < \kappa \text{ such that } f(x_\xi) \in Y_\gamma.$$

Since $\lambda < \kappa = \mathrm{cf}(\kappa)$, there exists $\beta < \kappa$ such that $\varphi(\xi) \leq \beta$ for all $\xi < \lambda$. Thus $f[X_\alpha] \subseteq Y_\beta$. □

CLAIM 3.1.12. *If $\beta < \kappa$, then there exists $\gamma < \kappa$ such that $Y_\beta \subseteq f[X_\gamma]$.*

PROOF OF CLAIM 3.1.12. By applying Claim 3.1.11 to the inverse function $f^{-1} : Y \to X$, we see that there exists $\gamma < \kappa$ such that $f^{-1}[Y_\beta] \subseteq X_\gamma$. The result follows. □

Let $\alpha < \kappa$. We shall find an element $\beta \in C$ such that $\alpha < \beta < \kappa$. Using Claims 3.1.11 and 3.1.12, we can inductively define a strictly increasing sequence

$$\alpha = \alpha_0 < \alpha_1 < \cdots < \alpha_n < \cdots < \kappa$$

for $n \in \omega$ such that

$$f[X_{\alpha_0}] \subseteq Y_{\alpha_1} \subseteq f[X_{\alpha_2}] \subseteq \cdots \subseteq f[X_{\alpha_{2n}}] \subseteq Y_{\alpha_{2n+1}} \subseteq \cdots$$

Let $\beta = \sup\{\alpha_n \mid n \in \omega\}$. Since $\omega < \kappa = \mathrm{cf}(\kappa)$, it follows that $\beta < \kappa$. Furthermore,

$$f[X_\beta] = \bigcup_{n\in\omega} f[X_{\alpha_{2n}}] = \bigcup_{n\in\omega} Y_{\alpha_{2n+1}} = Y_\beta.$$

Thus $\beta \in C$. $\square$

THEOREM 3.1.13. *If G is an infinite centreless group, then the automorphism tower of G terminates after at most $\left(2^{|G|}\right)^+$ steps.*

PROOF. Let G be an infinite centreless group and let

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \cdots$$

be the automorphism tower of G. By Theorem 1.1.10, $C_{G_\alpha}(G) = 1$ for all ordinals α. Suppose that $\alpha \geq 1$ and that $g \in N_{G_\alpha}(G)$. Then g induces an automorphism of G via conjugation, and so there exists an element $h \in G_1 = \mathrm{Aut}(G)$ such that $gxg^{-1} = hxh^{-1}$ for all $x \in G$. Thus $h^{-1}g \in C_{G_\alpha}(G) = 1$. Therefore $g = h \in G_1$ and we have shown that $N_{G_\alpha}(G) = G_1$ for all $\alpha \geq 1$. It follows that

$$[G_{\alpha+1} : G_1] = [G_{\alpha+1} : N_{G_\alpha}(G)]$$

equals the cardinality of the set of all conjugates of G in $G_{\alpha+1}$. Since $G \leqslant G_\alpha$ and $G_\alpha \trianglelefteq G_{\alpha+1}$, each such conjugate is contained in G_α and so $[G_{\alpha+1} : G_1] \leq |G_\alpha|^{|G|}$. Hence

$$|G_{\alpha+1}| = |G_1|\,[G_{\alpha+1} : G_1] \leq |G_1|\,|G_\alpha|^{|G|} = |G_\alpha|^{|G|}$$

and an easy transfinite induction shows that $|G_\alpha| \leq 2^{|G|}$ for all $\alpha < \left(2^{|G|}\right)^+$.

Let $\kappa = \left(2^{|G|}\right)^+$. Suppose that the automorphism tower of G does not terminate in less than κ steps. Thus $G_\alpha \subsetneq G_\beta$ for all $\alpha < \beta < \kappa$ and $G_\kappa = \bigcup_{\alpha<\kappa} G_\alpha$ is a smooth strictly increasing union. We shall show that $G_\kappa = G_{\kappa+1}$. Let $\pi \in \mathrm{Aut}(G_\kappa)$ be an arbitrary automorphism. Since κ is a regular uncountable cardinal and $|G_\alpha| < \kappa$ for all $\alpha < \kappa$, Theorem 3.1.10 yields that

$$C = \{\alpha < \kappa \mid \pi[G_\alpha] = G_\alpha\}$$

is a club of κ. If $\alpha \in C$, then there exists an element $g_\alpha \in G_{\alpha+1} = \mathrm{Aut}(G_\alpha)$ such that $g_\alpha x g_\alpha^{-1} = \pi(x)$ for all $x \in G_\alpha$. Notice that if α, $\beta \in C$ and $\alpha < \beta$, then $g_\alpha^{-1} g_\beta \in C_{G_\kappa}(G_\alpha) \leqslant C_{G_\kappa}(G) = 1$ and so $g_\alpha = g_\beta$. Thus there is a fixed element $g \in G_\kappa$ such that $\pi \restriction G_\alpha = i_g \restriction G_\alpha$ for each $\alpha \in C$. Since C is unbounded in κ, this implies that $\pi = i_g \in \mathrm{Inn}(G_\kappa)$. We have now shown that $\mathrm{Inn}(G_\kappa) = \mathrm{Aut}(G_\kappa)$; i.e. that $G_\kappa = G_{\kappa+1}$. □

DEFINITION 3.1.14. If G is a centreless group, then the *height* $\tau(G)$ of the automorphism tower of G is the least ordinal α such that $G_{\alpha+1} = G_\alpha$.

Most of this book will be concerned with the problem of finding an optimal upper bound for $\tau(G)$ in terms of the cardinality κ of G.

DEFINITION 3.1.15. If κ is an infinite cardinal, then τ_κ is the least ordinal such that $\tau(G) < \tau_\kappa$ for every centreless group G of cardinality κ.

Theorem 3.1.13 says that $\tau_\kappa \leq (2^\kappa)^+ + 1$. In Section 3.3, we shall use Fodor's Lemma to prove that $\tau_\kappa < (2^\kappa)^+$. In the other direction, in Section 4.1, we shall show that for every $\alpha < \kappa^+$, there exists a centreless group G of cardinality κ such that $\tau(G) = \alpha$. Thus $\kappa^+ \leq \tau_\kappa < (2^\kappa)^+$. It is natural to ask whether better upper and lower bounds for τ_κ can be proved in ZFC. In Chapter 7, we shall show that no such upper bounds can be proved in ZFC. It remains an open problem whether a better lower bound can be proved.

QUESTION 3.1.16. Let $\kappa \geq \omega$. Does there exist a centreless group G of cardinality κ such that $\tau(G) \geq \kappa^+$?

In Chapter 7, we shall see that if $\kappa > \omega$, then it is consistent that such a group exists. However, it is unknown whether it is consistent that there exists a countable centreless group G such that $\tau(G) \geq \omega_1$.

Nothing is known concerning the exact value of τ_κ in any model of ZFC.

PROBLEM 3.1.17. Find a model M of ZFC and an infinite cardinal $\kappa \in M$ such that it is possible to compute the exact value of τ_κ in M.

In Chapter 7, we shall see that it is consistent that τ_κ is not a cardinal. But it remains open whether τ_κ can be a successor ordinal. This seems a remote possibility, for it would mean that there was a centreless group G of cardinality κ with an automorphism tower of *maximum* height $\tau(G) = \tau_\kappa - 1$. However, it seems likely that given a centreless group G of cardinality κ such that $\tau(G) = \alpha$, it should be possible to construct a related group H of cardinality κ such that $\tau(H) = \alpha + 1$.

Thus I expect that τ_κ must be a limit ordinal. Similar considerations lead me to expect that $\mathrm{cf}(\tau_\kappa) > \kappa$. For suppose that $\alpha = \sup_{i<\kappa} \alpha_i$ and that for each $i < \kappa$, there exists a centreless group G_i of cardinality κ such that $\tau(G_i) = \alpha_i$. Then it should be possible to construct a centreless group G of cardinality κ such that $\tau(G) = \alpha$. A first approximation to such a group G might be the direct sum $\oplus_{i<\kappa} G_i$.

Conjecture 3.1.18. If κ is an infinite cardinal, then τ_κ is a limit ordinal such that $\mathrm{cf}(\tau_\kappa) > \kappa$.

It is also natural to ask which ordinals $\alpha < \tau_\kappa$ are actually realised as the heights of automorphism towers of centreless groups of cardinality κ.

Conjecture 3.1.19. For every $\alpha < \tau_\kappa$, there exists a centreless group G of cardinality κ such that $\tau(G) = \alpha$.

3.2. τ_κ is increasing

In this section, we shall prove that the function, $\kappa \mapsto \tau_\kappa$, is increasing. This result is an immediate consequence of the following theorem, which will also be needed in Chapter 7.

Theorem 3.2.1. *Let $\omega \leq \lambda < \kappa$. If H is a centreless group of cardinality λ such that $\tau(H) \geq 1$, then $\tau(H \times \mathrm{Alt}(\kappa)) = \tau(H)$.*

Corollary 3.2.2. *If $\omega \leq \lambda < \kappa$, then $\tau_\lambda \leq \tau_\kappa$.*

□

The remainder of this section will be devoted to the proof of Theorem 3.2.1. Let $G = H \times \mathrm{Alt}(\kappa)$ and let H_β, G_β be the βth groups in the automorphism towers of H, G respectively. We shall prove by induction that $G_\beta = H_\beta \times \mathrm{Sym}(\kappa)$ for all $\beta \geq 1$. To accomplish this, we need to keep track of $\varphi[\mathrm{Alt}(\kappa)]$ for each automorphism φ of G_β. The next lemma shows that for all $\varphi \in \mathrm{Aut}(G_\beta)$, either $\varphi[\mathrm{Alt}(\kappa)] \leqslant H_\beta$ or $\varphi[\mathrm{Alt}(\kappa)] \leqslant \mathrm{Sym}(\kappa)$. The main point will be to eliminate the possibility that $\varphi[\mathrm{Alt}(\kappa)] \leqslant H_\beta$. This will be straightforward when β is a successor ordinal. To deal with the case when β is a limit ordinal, we shall make use of the result that $\mathrm{Alt}(\kappa)$ is strictly simple.

LEMMA 3.2.3. *Suppose that A is a simple nonabelian normal subgroup of the direct product $H \times S$. Then either $A \leqslant H$ or $A \leqslant S$.*

PROOF. Let $1 \neq g = xy \in A$, where $x \in H$ and $y \in S$. If $y = 1$, then the conjugacy class $g^A = x^A$ is contained in H and so $A = \langle g^A \rangle \leqslant H$. So suppose that $y \neq 1$. Let $\pi : H \times S \to S$ be the canonical projection map. Then $1 \neq y \in \pi[A] \leqslant S$ and $\pi[A] \simeq A$. Hence there exists an element $z \in \pi[A] \leqslant S$ such that $zyz^{-1} \neq y$. Since $A \trianglelefteq H \times S$, it follows that

$$zgz^{-1}g^{-1} = zxyz^{-1}y^{-1}x^{-1} = zyz^{-1}y^{-1} \neq 1$$

is an element of $A \cap S$. Arguing as above, we now obtain that $A \leqslant S$. □

The notion of an ascendant subgroup generalises that of a subnormal subgroup.

DEFINITION 3.2.4. Let H be a subgroup of the group G. Then H is said to be an *ascendant subgroup* of G if there exist an ordinal β and a strictly increasing chain of subgroups $\{H_\alpha \mid \alpha \leq \beta\}$ such that the following conditions are satisfied.

(a) $H_0 = H$ and $H_\beta = G$.
(b) If $\alpha < \beta$, then $H_\alpha \trianglelefteq H_{\alpha+1}$.
(c) If δ is a limit ordinal such that $\delta \leq \beta$, then $H_\delta = \bigcup_{\alpha<\delta} H_\alpha$.

For example, let G be a centreless group and let $\tau = \tau(G)$. If $\langle G_\alpha \mid \alpha \leq \tau \rangle$ is the automorphism tower of G, then G is an ascendant subgroup of G_α for each $\alpha \leq \tau$.

DEFINITION 3.2.5. A group A is *strictly simple* if it has no nontrivial proper ascendant subgroups.

Clearly if A is strictly simple, then A is simple. However, Hall [**13**] has shown that the converse does not hold.

THEOREM 3.2.6 (Macpherson and Neumann [**28**]). *For each $\kappa \geq \omega$, the alternating group* $\mathrm{Alt}(\kappa)$ *is strictly simple.*

PROOF. Suppose that H is a nontrivial proper ascendant subgroup of $\mathrm{Alt}(\kappa)$. Then there exists an ordinal β and a strictly increasing chain $\{H_\alpha \mid \alpha \leq \beta\}$ of subgroups such that the following conditions are satisfied.

(a) $H_0 = H$ and $H_\beta = \mathrm{Alt}(\kappa)$.

(b) If $\alpha < \beta$, then $H_\alpha \trianglelefteq H_{\alpha+1}$.

(c) If δ is a limit ordinal such that $\delta \leq \beta$, then $H_\delta = \bigcup_{\alpha<\delta} H_\alpha$.

Since $\mathrm{Alt}(\kappa)$ is simple, it follows that β is a limit ordinal. Let γ be the least ordinal such that H_γ contains a 3-cycle π. Then H_γ is also a nontrivial proper ascendant subgroup of $\mathrm{Alt}(\kappa)$ and so we can suppose that $\pi \in H_0$. Let $\Delta = \mathrm{supp}(\pi)$. Let $\sigma \in \mathrm{Alt}(\kappa)$ be any 3-cycle such that $\mathrm{supp}(\sigma) \cap \Delta = \emptyset$ and let $\Phi = \mathrm{supp}(\sigma)$. Then $A_\sigma = \mathrm{Alt}(\Delta \cup \Phi)$ is a finite simple group. Let $\alpha < \beta$ be the least ordinal such that $A_\sigma \leqslant H_\alpha$. Clearly α is not a limit ordinal. Suppose that $\alpha = \gamma + 1$ is a successor ordinal. Then

$$\pi \in A_\sigma \cap H_\gamma \trianglelefteq A_\sigma \cap H_{\gamma+1} = A_\sigma,$$

and so $A_\sigma \cap H_\gamma$ is a nontrivial proper normal subgroup of A_σ, which is a contradiction. Thus $\alpha = 0$. We have now shown that $\sigma \in H_0$ for every 3-cycle $\sigma \in \mathrm{Alt}(\kappa)$ such that $\mathrm{supp}(\sigma) \cap \Delta = \emptyset$. Since each alternating group is generated by its 3-cycles, it follows that $\mathrm{Alt}(\kappa \smallsetminus \Delta) \leqslant H_0$. Finally note that if $\sigma \in \mathrm{Alt}(\kappa)$ is any 3-cycle such that $\mathrm{supp}(\sigma) \cap \Delta = \emptyset$, then $\mathrm{Alt}(\kappa) = \langle \mathrm{Alt}(\kappa \smallsetminus \Delta), A_\sigma \rangle$. Hence $\mathrm{Alt}(\kappa) \leqslant H_0$, which is a contradiction. □

LEMMA 3.2.7. *Suppose that H is a centreless group. Let $\tau = \tau(H)$ and let $\langle H_\alpha \mid \alpha \leq \tau \rangle$ be the automorphism tower of H. If A is a strictly simple normal subgroup of H_τ, then $A \leqslant H_0$.*

PROOF. Let $\alpha \leq \tau$ be the least ordinal such that $A \leqslant H_\alpha$. First suppose that α is a limit ordinal. Then $A = \bigcup_{\beta<\alpha} (A \cap H_\beta)$ and $A \cap H_\beta \trianglelefteq A \cap H_{\beta+1}$ for all $\beta < \alpha$. It follows that if $\gamma < \alpha$ is the least ordinal such that $A \cap H_\gamma \neq 1$, then $A \cap H_\gamma$ is a nontrivial proper ascendant subgroup of A. But this contradicts the assumption that A is strictly simple.

Next suppose that $\alpha = \beta + 1$ is a successor ordinal. Since $A \cap H_\beta$ is a proper normal subgroup of A, it follows that $A \cap H_\beta = 1$. Since $A \leqslant H_{\beta+1} \leqslant N_{H_\tau}(H_\beta)$ and $H_\beta \leqslant N_{H_\tau}(A) = H_\tau$, Lemma 1.2.3 implies that $[A, H_\beta] \leqslant A \cap H_\beta = 1$. But then $A \leqslant C_{H_{\beta+1}}(H_\beta) = 1$, which is a contradiction. The only remaining possibility is that $\alpha = 0$. □

Proof of Theorem 3.2.1. Let $\tau = \tau(H)$ and let

$$H = H_0 \trianglelefteq H_1 \trianglelefteq \cdots \trianglelefteq H_\alpha \trianglelefteq H_{\alpha+1} \trianglelefteq \cdots \trianglelefteq H_\tau = H_{\tau+1} = \cdots$$

be the automorphism tower of H. Let $G = H \times \mathrm{Alt}(\kappa)$ and let

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \cdots$$

be the automorphism tower of G. We shall prove by induction on $\alpha \geq 1$ that $G_\alpha = H_\alpha \times \mathrm{Sym}(\kappa)$.

First consider the case when $\alpha = 1$. Let $\varphi \in \mathrm{Aut}(G)$ be any automorphism. Then $\varphi[\mathrm{Alt}(\kappa)]$ is a simple nonabelian normal subgroup of $G = H \times \mathrm{Alt}(\kappa)$. By Lemma 3.2.3, either $\varphi[\mathrm{Alt}(\kappa)] \leqslant H$ or $\varphi[\mathrm{Alt}(\kappa)] \leqslant \mathrm{Alt}(\kappa)$. Since

$$|\varphi[\mathrm{Alt}(\kappa)]| = \kappa > \lambda = |H|,$$

it follows that $\varphi[\mathrm{Alt}(\kappa)] \leqslant \mathrm{Alt}(\kappa)$. As $\varphi[\mathrm{Alt}(\kappa)]$ is a normal subgroup of G, we must have that $\varphi[\mathrm{Alt}(\kappa)] = \mathrm{Alt}(\kappa)$. Note that $C_G(\mathrm{Alt}(\kappa)) = H$. Hence we must also have that $\varphi[H] = H$. By Theorem 1.3.12, $\mathrm{Aut}(\mathrm{Alt}(\kappa)) = \mathrm{Sym}(\kappa)$. It follows that

$$G_1 = \mathrm{Aut}(G) = \mathrm{Aut}(H) \times \mathrm{Aut}(\mathrm{Alt}(\kappa)) = H_1 \times \mathrm{Sym}(\kappa).$$

Next suppose that $\alpha = \beta + 1$ and that $G_\beta = H_\beta \times \mathrm{Sym}(\kappa)$. Let $\varphi \in \mathrm{Aut}(G_\beta)$ be any automorphism. By Lemma 3.2.3, either $\varphi[\mathrm{Alt}(\kappa)] \leqslant H_\beta$ or $\varphi[\mathrm{Alt}(\kappa)] \leqslant \mathrm{Sym}(\kappa)$. As $\mathrm{Alt}(\kappa)$ is a strictly simple group, Lemma 3.2.7 implies that $\varphi[\mathrm{Alt}(\kappa)] \leqslant \mathrm{Sym}(\kappa)$. Since $\varphi[\mathrm{Alt}(\kappa)]$ is a normal subgroup of G_β, it follows that $\varphi[\mathrm{Alt}(\kappa)] = \mathrm{Alt}(\kappa)$. Using the facts that $C_{G_\beta}(\mathrm{Alt}(\kappa)) = H_\beta$ and $C_{G_\beta}(H_\beta) = \mathrm{Sym}(\kappa)$, we now see that $\varphi[H_\beta] = H_\beta$ and $\varphi[\mathrm{Sym}(\kappa)] = \mathrm{Sym}(\kappa)$. Hence

$$G_{\beta+1} = \mathrm{Aut}(G_\beta) = \mathrm{Aut}(H_\beta) \times \mathrm{Aut}(\mathrm{Sym}(\kappa)) = H_{\beta+1} \times \mathrm{Sym}(\kappa).$$

Finally no difficulties arise when α is a limit ordinal. $\square$

Question 3.2.8. Is τ_κ a strictly increasing function of κ?

In Chapter 7, we shall prove that it is consistent that τ_κ is a strictly increasing function of κ.

3.3. A better bound

In Section 3.1, we proved Wielandt's theorem that if H is a finite subnormal subgroup of a group G with $C_G(H) = 1$, then there is an upper bound for $|G|$ depending only on $|H|$. In [**7**], Faber proved the following analogous result for infinite ascendant subgroups.

Theorem 3.3.1. *If H is an infinite ascendant subgroup of the group G and $C_G(H) = 1$, then $|G| \leq 2^{|H|}$.*

Clearly Theorem 3.3.1 gives the best possible bound. (For example, consider the inclusion $\mathrm{Alt}(\kappa) \trianglelefteq \mathrm{Sym}(\kappa)$.) It also yields the following improvement on Theorem 3.1.13.

Corollary 3.3.2. *If G is an infinite centreless group, then $\tau(G) < \left(2^{|G|}\right)^+$.*

Proof. Let G be an infinite centreless group and let

$$G = G_0 \trianglelefteq G_1 \trianglelefteq G_2 \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq G_{\alpha+1} \trianglelefteq \cdots$$

be the automorphism tower of G. Let $\tau = \tau(G)$. Then G is an ascendant subgroup of G_τ and $C_{G_\tau}(G) = 1$. Thus $|G_\tau| \leq 2^{|G|}$. Since $G_\alpha \subsetneq G_{\alpha+1}$ for all $\alpha < \tau$, it follows that $|G_\tau| \geq |\tau|$. Hence $\tau < \left(2^{|G|}\right)^+$. $\square$

Theorem 3.3.1 is a straightforward consequence of Fodor's Lemma on regressive functions. Before we can prove Fodor's Lemma, we must first present a little more of the theory of clubs. As Kunen points out in [**25**], there is a suggestive analogy between measure theory and the theory of clubs; namely, clubs can be imagined as being large or of probability measure 1. The following lemma is an analogue of the result that in a probability space, the intersection of countably many sets of measure 1 also has measure 1.

Lemma 3.3.3. *Let κ be a regular uncountable cardinal and let $\lambda < \kappa$. If C_α is a club of κ for each $\alpha < \lambda$, then $C = \bigcap_{\alpha<\lambda} C_\alpha$ is also a club of κ.*

Proof. It is easily checked that C is closed. The main point is to show that C is unbounded in κ. To this end, for each $\alpha < \lambda$, define the function $f_\alpha : \kappa \to \kappa$ by

$$f_\alpha(\beta) = \text{ the least } \gamma \in C_\alpha \text{ such that } \gamma > \beta.$$

Let $g : \kappa \to \kappa$ be the function defined by $g(\beta) = \sup\{f_\alpha(\beta) \mid \alpha < \lambda\}$. Since $\lambda < \mathrm{cf}(\kappa) = \kappa$, we have that $g(\beta) < \kappa$. Thus g is well-defined. For each $1 \leq n < \omega$, let

$$g^n = \underbrace{g \circ \cdots \circ g}_{n \text{ times}} : \kappa \to \kappa$$

be the n-fold composition of g; and define $g^\omega : \kappa \to \kappa$ by

$$g^\omega(\beta) = \sup\{g^n(\beta) \mid 1 \leq n < \omega\}.$$

Clearly $\beta < g^\omega(\beta) < \kappa$. We claim that $g^\omega(\beta) \in C = \bigcap_{\alpha<\lambda} C_\alpha$. To see this , fix some ordinal $\alpha < \lambda$. Then

$$f_\alpha(\beta) \leq g(\beta) < f_\alpha(g(\beta)) \leq g^2(\beta) < \cdots < f_\alpha(g^n(\beta)) \leq g^{n+1}(\beta) < \cdots$$

Thus $g^\omega(\beta) = \sup\{f_\alpha(g^n(\beta)) \mid 1 \leq n < \omega\}$. Since $f_\alpha(g^n(\beta)) \in C_\alpha$ for all $1 \leq n < \omega$, it follows that $g^\omega(\beta) \in C_\alpha$. $\square$

Definition 3.3.4. Let κ be a regular uncountable cardinal. A subset $S \subseteq \kappa$ is *stationary* in κ if $S \cap C \neq \emptyset$ for every club C of κ.

If $T \subseteq \kappa$ is nonstationary, then there exists a club C such that $T \subseteq \kappa \smallsetminus C$. So continuing with our measure theoretic analogy, nonstationary sets can be imagined as being small or of probability measure 0. Consequently, stationary sets can be imagined as having positive measure. By Lemma 3.3.3, every club is stationary. If $\kappa > \lambda^+$, then the next result provides an example of a stationary set which is not a club. (To see that T is not a club, for each $\xi < \lambda^+$, let β_ξ be the ξth element of T. Then $\sup\{\beta_\xi \mid \xi < \lambda^+\} \notin T$.)

Lemma 3.3.5. *Let κ be a regular uncountable cardinal. If λ is a regular cardinal such that $\omega \leq \lambda < \kappa$, then*

$$T = \{\alpha < \kappa \mid \mathrm{cf}(\alpha) = \lambda\}$$

is a stationary subset of κ.

Proof. Let C be a club of κ and let $\{\alpha_\xi \mid \xi < \kappa\}$ be the increasing enumeration of C. Then $\alpha_\lambda = \sup\{\alpha_\xi \mid \xi < \lambda\}$ and so $\mathrm{cf}(\alpha_\lambda) = \mathrm{cf}(\lambda) = \lambda$. Thus $\alpha_\lambda \in T \cap C$. $\square$

Algebraists often experience more difficulty in grasping the significance of the notion of a stationary set than that of a club. So before giving the definition of a regressive function, I will give a simple example which illustrates the natural manner in which stationary sets arise in the construction of uncountable structures.

EXAMPLE 3.3.6. A linear ordering L is said to be a DLO if L is dense and has no endpoints. It is well-known that every countable DLO is isomorphic to the ordered set $\mathbb{Q}$ of rational numbers. Now consider the problem of constructing nonisomorphic DLOs of cardinality ω_1. For each $A \subseteq \omega_1$, we shall construct a DLO

$$D^A = \bigcup_{\alpha<\omega_1} D^A_\alpha$$

as the smooth strictly increasing union of a suitably chosen sequence of countable DLOs. In the successor steps of our construction, we shall make use of two kinds of embeddings $D^A_\alpha \subset D^A_{\alpha+1}$.

DEFINITION 3.3.7. An embedding $D_1 \subset D_2$ of countable DLOs is said to be a *rational embedding* if

(a) D_1 is an initial segment of D_2; and
(b) D_1 has a least upper bound in D_2.

In this case, the embedding $D_1 \subset D_2$ is isomorphic to $(-\infty, 1) \cap \mathbb{Q} \subset \mathbb{Q}$.

DEFINITION 3.3.8. An embedding $D_1 \subset D_2$ of countable DLOs is said to be an *irrational embedding* if

(a) D_1 is an initial segment of D_2; and
(b) D_1 has *no* least upper bound in D_2.

In this case, the embedding $D_1 \subset D_2$ is isomorphic to $(-\infty, \sqrt{2}) \cap \mathbb{Q} \subset \mathbb{Q}$.

For each $A \subseteq \omega_1$, let $D^A = \bigcup_{\alpha<\omega_1} D^A_\alpha$ be the smooth strictly increasing union of the countable DLOs D^A_α chosen so that

(i) if $\alpha \in A$, then $D^A_\alpha \subset D^A_{\alpha+1}$ is a rational embedding; and
(ii) if $\alpha \notin A$, then $D^A_\alpha \subset D^A_{\alpha+1}$ is an irrational embedding.

It is natural to expect that if A and B are "sufficiently different" subsets of ω_1, then $D^A \not\cong D^B$. As the next result shows, stationary sets provide the correct notion of "sufficiently different".

PROPOSITION 3.3.9. *$D^A \not\cong D^B$ iff the symmetric difference $A \triangle B$ is a stationary subset of ω_1.*

PROOF. First assume that $S = A \triangle B$ is a stationary subset of ω_1. Suppose that $f : D^A \to D^B$ is an isomorphism. By Theorem 3.1.10, the set

$$C = \{\alpha < \omega_1 \mid f[D^A_\alpha] = D^B_\alpha\}$$

is a club of ω_1. Hence there exists an ordinal $\alpha \in S \cap C$. Without loss of generality, we can suppose that $\alpha \in A \smallsetminus B$. But then D^A_α has a least upper bound in D^A and $f[D^A_\alpha] = D^B_\alpha$ has no least upper bound in D^B, which is a contradiction.

Now assume that $S = A \triangle B$ is a nonstationary subset of ω_1. Then there exists a club C such that $S \cap C = \emptyset$. Let $\{\alpha_\xi \mid \xi < \omega_1\}$ be the increasing enumeration of C. Note that for each $\xi < \omega_1$, $D^A_{\alpha_\xi}$ has a least upper bound in $D^A_{\alpha_{\xi+1}}$ iff $D^B_{\alpha_\xi}$ has a least upper bound in $D^B_{\alpha_{\xi+1}}$. Hence it is easy to inductively define a sequence of isomorphisms $f_\xi : D^A_{\alpha_\xi} \to D^B_{\alpha_\xi}$ such that $f_\eta \subset f_\xi$ for all $\eta < \xi < \omega_1$. Then $f = \bigcup_{\xi<\omega_1} f_\xi$ is an isomorphism from D^A onto D^B. □

It is well-known that for every regular uncountable cardinal κ, there exists a family $\{A_\xi \mid \xi < 2^\kappa\}$ of subsets of κ such that $A_\xi \triangle A_\eta$ is stationary for all $\xi < \eta < 2^\kappa$. (Since we shall not require this result in the main body of this book, we shall not give a proof here. A clear treatment of this result can be found in Section II.4 of [**6**].) Hence our construction yields a family of 2^{ω_1} pairwise nonisomorphic DLOs of cardinality ω_1.

DEFINITION 3.3.10. Let κ be a regular uncountable cardinal and let $T \subseteq \kappa$. The function $f : T \to \kappa$ is *regressive* if $f(\alpha) < \alpha$ for all $\alpha \in T$.

For example, suppose that T is a nonstationary subset of κ such that $0 \notin T$. Then there exists a club C such that $T \cap C = \emptyset$ and $0 \in C$. So we can define a regressive function $f : T \to \kappa$ by

$$\begin{aligned} f(\alpha) &= \sup\{\beta < \alpha \mid \beta \in C\} \\ &= \max\{\beta < \alpha \mid \beta \in C\}. \end{aligned}$$

Notice that for each $\gamma < \kappa$, the set $\{\alpha \in T \mid f(\alpha) = \gamma\}$ has cardinality less than κ. In contrast, Fodor's Lemma implies that every regressive function defined on a stationary subset of κ is constant on a set of cardinality κ.

THEOREM 3.3.11 (Fodor's Lemma). *Let κ be a regular uncountable cardinal and let T be a stationary subset of κ. If $f : T \to \kappa$ is a regressive function, then there exists an ordinal $\gamma < \kappa$ such that*

$$S = \{\alpha \in T \mid f(\alpha) = \gamma\}$$

is a stationary subset of κ.

Fodor's Lemma is an easy consequence of the following somewhat technical lemma.

LEMMA 3.3.12. *Let κ be a regular uncountable cardinal and let C_γ be a club of κ for each $\gamma < \kappa$. Then the* diagonal intersection

$$D = \{\beta < \kappa \mid \beta \in \bigcap_{\gamma<\beta} C_\gamma\}$$

is also a club of κ.

PROOF. First we shall show that D is closed in κ. So suppose that

$$\beta_0 < \beta_1 < \cdots < \beta_\xi < \ldots$$

is an increasing δ-sequence of elements of D for some limit ordinal $\delta < \kappa$. Let $\beta = \sup\{\beta_\xi \mid \xi < \delta\}$. If $\gamma < \beta$, then there exists $\xi_0 < \delta$ such that $\gamma < \beta_\xi$ for all $\xi_0 \leq \xi < \delta$. Thus $\beta_\xi \in C_\gamma$ for all $\xi_0 \leq \xi < \delta$ and so $\beta = \sup\{\beta_\xi \mid \xi_0 \leq \xi < \delta\} \in C_\gamma$. Hence $\beta \in D$.

To see that D is unbounded in κ, define the function $g : \kappa \to \kappa$ by

$$g(\beta) = \text{ the least } \alpha \in \bigcap_{\gamma<\beta} C_\gamma \text{ such that } \alpha > \beta.$$

By Lemma 3.3.3, $\bigcap_{\gamma<\beta} C_\gamma$ is also a club of κ. Thus g is well-defined. Once again, for each $1 \leq n < \omega$, let $g^n : \kappa \to \kappa$ be the n-fold composition of g; and let $g^\omega(\beta) = \sup\{g^n(\beta) \mid 1 \leq n < \omega\}$. Clearly $\beta < g^\omega(\beta) < \kappa$. We claim that $g^\omega(\beta) \in D$. To see this, fix some $\gamma < g^\omega(\beta)$. Then there exists an integer m such that $\gamma < g^n(\beta)$ for all $m \leq n < \omega$. Thus $g^n(\beta) \in C_\gamma$ for all $m < n < \omega$ and so $g^\omega(\beta) = \sup\{g^n(\beta) \mid m < n < \omega\} \in C_\gamma$. Hence $g^\omega(\beta) \in D$. □

PROOF OF FODOR'S LEMMA. Suppose that the regressive function $f : T \to \kappa$ is a counterexample. Then for every $\gamma < \kappa$, there exists a club C_γ of κ such that $f(\beta) \neq \gamma$ for all $\beta \in C_\gamma \cap T$. By Lemma 3.3.12, the diagonal intersection $D = \{\beta < \kappa \mid \beta \in \bigcap_{\gamma<\beta} C_\gamma\}$ is also a club of κ. Hence there exists an ordinal $\beta \in T \cap D$. Since $\beta \in D$, we have that $f(\beta) \neq \gamma$ for all $\gamma < \beta$. But this contradicts the fact that $f(\beta) < \beta$. □

We are now ready to prove Faber's theorem on infinite ascendant subgroups.

Proof of Theorem 3.3.1. Suppose that H is an infinite ascendant subgroup of the group G and that $C_G(H) = 1$. Then there exist an ordinal β and a strictly increasing chain of subgroups $\{H_\alpha \mid \alpha \leq \beta\}$ such that the following conditions are satisfied.

(a) $H_0 = H$ and $H_\beta = G$.
(b) If $\alpha < \beta$, then $H_\alpha \trianglelefteq H_{\alpha+1}$.
(c) If δ is a limit ordinal such that $\delta \leq \beta$, then $H_\delta = \bigcup_{\alpha<\delta} H_\alpha$.

We shall prove inductively that $|H_\alpha| \leq 2^{|H|}$ for all $\alpha \leq \beta$. This is certainly true when $\alpha = 0$. Next consider the case when α is a successor ordinal; say, $\alpha = \gamma + 1$. For each $h \in H_\alpha$, let $\varphi_h : H \to H_\alpha$ be the embedding such that $\varphi_h(x) = hxh^{-1}$ for all $x \in H$. Since $H \leqslant H_\gamma \trianglelefteq H_\alpha$, we have that $\varphi_h[H] \leqslant H_\gamma$ for all $h \in H_\alpha$. Thus

$$|\{\varphi_h \mid h \in H_\alpha\}| \leq |H_\gamma|^{|H|} \leq 2^{|H|\cdot|H|} = 2^{|H|}.$$

If $\varphi_h = \varphi_g$, then $h^{-1}g \in C_{H_\alpha}(H) = 1$ and so $h = g$. Hence $|H_\alpha| \leq 2^{|H|}$.

Finally suppose that α is a limit ordinal. If $\alpha < \left(2^{|H|}\right)^+$, then it is clear that $|H_\alpha| \leq 2^{|H|}$. Thus to complete the proof of Theorem 3.3.1, we need only show that $\beta < \left(2^{|H|}\right)^+$. Suppose not. Let $\kappa = \left(2^{|H|}\right)^+$ and $\lambda = \operatorname{cf}\left(2^{|H|}\right)$. By Lemma 3.3.5, $T = \{\alpha < \kappa \mid \operatorname{cf}(\alpha) = \lambda\}$ is a stationary subset of κ. For each $\alpha \in T$, choose an element $h_\alpha \in H_{\alpha+1} \smallsetminus H_\alpha$. Since $H \leqslant H_\alpha \trianglelefteq H_{\alpha+1}$, we have that

$$h_\alpha H h_\alpha^{-1} \leqslant H_\alpha = \bigcup_{\gamma<\alpha} H_\gamma.$$

By Lemma 3.1.4, $|H| < \operatorname{cf}\left(2^{|H|}\right) = \lambda = \operatorname{cf}(\alpha)$. Hence there exists an ordinal $f(\alpha) < \alpha$ such that $h_\alpha H h_\alpha^{-1} \leqslant H_{f(\alpha)}$. Applying Fodor's Lemma to the regressive function $f : T \to \kappa$, we see that there exists an ordinal $\gamma < \kappa$ such that

$$S = \{\alpha \in T \mid f(\alpha) = \gamma\}$$

is a stationary subset of κ. For each $\alpha \in S$, let $\varphi_\alpha : H \to H_\gamma$ be the embedding such that $\varphi_\alpha(x) = h_\alpha x h_\alpha^{-1}$ for all $x \in H$. Then

$$|\{\varphi_\alpha \mid \alpha \in S\}| \leq |H_\gamma|^{|H|} \leq 2^{|H|} < \kappa.$$

Since $|S| = \kappa$, there exist distinct ordinals α_1, $\alpha_2 \in S$ such that $\varphi_{\alpha_1} = \varphi_{\alpha_2}$. But then $1 \neq h_{\alpha_1}^{-1} h_{\alpha_2} \in C_{H_\kappa}(H)$, which is a contradiction. □

We shall end this section with the following easy observation, which shows that Corollary 3.3.2 does not give the best possible upper bound for τ_κ.

THEOREM 3.3.13. *If κ is an infinite cardinal, then $\tau_\kappa < (2^\kappa)^+$.*

PROOF. By Corollary 3.3.2, if G is a centreless group of cardinality κ, then $\tau(G) < (2^\kappa)^+$. Since there are only 2^κ centreless groups of cardinality κ up to isomorphism, it follows that

$$\sup\{\tau(G) \mid G \text{ is a centreless group of cardinality } \kappa\} < (2^\kappa)^+ .$$

□

3.4. The automorphism tower problem revisited

We have just seen that Corollary 3.3.2 does not give the best possible upper bound for τ_κ. So it is natural to ask whether a better upper bound on τ_κ can be proved in ZFC, preferably one which does not involve cardinal exponentiation. Since the proofs in this chapter are extremely simple and use only the most basic results in group theory, together with some elementary properties of the infinite cardinal numbers, it is not really surprising that Corollary 3.3.2 does not give the best possible upper bound for τ_κ. In contrast, the proof of Wielandt's theorem is much deeper and involves an intricate analysis of the subnormal subgroups of a finite centreless group. The real question behind the search for better upper bounds for τ_κ is whether there exists a subtler, more informative, group-theoretic proof of the automorphism tower theorem for infinite groups. The main result of Chapter 7 says that no such bounds can be proved in ZFC, and thus can be interpreted as saying that no such proof exists.

Of course, it is still possible that such proofs exist for various restricted classes of infinite centreless groups. For example, Rae and Roseblade presented such a proof for the class of Černikov groups in [**38**], where they showed that the automorphism tower of a centreless Černikov group terminates after finitely many steps. (The definition of a Černikov group can be found in Section 5.4 of Robinson [**39**].) And in [**17**], Hulse gave an intricate proof that the automorphism tower of a centreless polycyclic group terminates in countably many steps. (Once again, the definition of a polycyclic group can be found in Section 5.4 of Robinson [**39**].) It turns out that there is actually a much easier proof of Hulse's theorem. Every polycyclic group is finitely generated and so Hulse's theorem is a consequence of the following stronger theorem. However, a careful examination of Hulse's proof should yield a better upper bound for the height of the automorphism tower of a centreless polycyclic group.

THEOREM 3.4.1. *If G is a finitely generated centreless group, then $\tau(G) < \omega_1$.*

Once again, this result is an immediate consequence of the appropriate analogue of Wielandt's theorem on finite subnormal subgroups.

THEOREM 3.4.2. *If H is a finitely generated ascendant subgroup of the group G and $C_G(H) = 1$, then $|G| \leq \omega$.*

PROOF. Since the proof of Theorem 3.4.2 is almost identical to that of Theorem 3.3.1, we shall just mention the main points. Suppose that $H = \langle h_1, \dots h_n \rangle$ is a finitely generated ascendant subgroup of the group G and that $C_G(H) = 1$. Then there exist an ordinal β and a strictly increasing chain of subgroups $\{H_\alpha \mid \alpha \leq \beta\}$ such that the following conditions are satisfied.

(a) $H_0 = H$ and $H_\beta = G$.

(b) If $\alpha < \beta$, then $H_\alpha \trianglelefteq H_{\alpha+1}$.

(c) If δ is a limit ordinal such that $\delta \leq \beta$, then $H_\delta = \bigcup_{\alpha<\delta} H_\alpha$.

By Theorem 2.1.3, we can suppose that H is an infinite group. We can now prove inductively that $|H_\alpha| = \omega$ for all $\alpha \leq \beta$. First consider the case when $\alpha = \gamma + 1$ is a successor ordinal. For each $h \in H_\alpha$, let $\varphi_h : H \to H_\alpha$ be the embedding such that $\varphi_h(x) = hxh^{-1}$ for all $x \in H$. Then $\varphi_h[H] \leqslant H_\gamma$ for all $h \in H_\alpha$ and each embedding φ_h is uniquely determined by its restriction $\varphi_h \restriction \{h_1, \dots, h_n\}$. Thus

$$\omega \leq |H_{\alpha+1}| = |\{\varphi_h \mid h \in H_{\alpha+1}\}| \leq |H_\gamma|^n = \omega^n = \omega.$$

Finally suppose that α is a limit ordinal. If $\alpha < \omega_1$, then it is clear that $|H_\alpha| = \omega$. Thus we need only show that $\beta < \omega_1$. Suppose not. Let $T = \{\alpha < \omega_1 \mid \mathrm{cf}(\alpha) = \omega\}$ and for each $\alpha \in T$, choose an element $h_\alpha \in H_{\alpha+1} \smallsetminus H_\alpha$. Then we have that $h_\alpha H h_\alpha^{-1} \leqslant H_\alpha = \bigcup_{\gamma<\alpha} H_\gamma$. Since H is a finitely generated group, there exists an ordinal $f(\alpha) < \alpha$ such that $h_\alpha H h_\alpha^{-1} \leqslant H_{f(\alpha)}$. Applying Fodor's Lemma to the regressive function $f : T \to \omega_1$, we see that there exists an ordinal $\gamma < \omega_1$ such that $S = \{\alpha \in T \mid f(\alpha) = \gamma\}$ is a stationary subset of ω_1. Arguing as in the proof of Theorem 3.3.1, we easily obtain a contradiction. □

Definition 3.4.3. τ_{fg} is the least ordinal such that $\tau(G) < \tau_{fg}$ for every finitely generated centreless group G.

B. H. Neumann [**33**] has shown that there are 2^ω two-generator centreless groups up to isomorphism. So it is conceivable that Theorem 3.4.1 gives the best possible bound.

Question 3.4.4. Is $\tau_{fg} = \omega_1$?

We obtain an even more interesting problem when we restrict our attention to the class of finitely presented groups.

Definition 3.4.5. τ_{fp} is the least ordinal such that $\tau(G) < \tau_{fp}$ for every finitely presented centreless group G.

Since there are only countably many finitely presented groups up to isomorphism, it follows that τ_{fp} is a countable ordinal.

Problem 3.4.6. Compute the exact value of τ_{fp}.

Problem 3.4.6 is probably very difficult. But there is a special case that appears to be much more managable. It is well-known that every polycyclic group is finitely presented; and a careful reading of Hulse's paper [**17**] may be enough to solve the following problem.

Definition 3.4.7. τ_{pc} is the least ordinal such that $\tau(G) < \tau_{pc}$ for every centreless polycyclic group G.

Problem 3.4.8. Compute the exact value of τ_{pc}.

Finally it should be pointed out that the argument of Theorem 3.4.2 gives an easy proof of the fact that if H is a finite ascendant subgroup of a group G such that $C_G(H) = 1$, then G is necessarily a countable group. (Using the notation of the proof of Theorem 3.4.2, it is clear that H_n is finite for each $n \in \omega$. If H_ω is infinite, then the argument of Theorem 3.4.2 shows that $|H_\alpha| = \omega$ for all $\omega \leq \alpha \leq \beta$.) Of course, this implies that the automorphism tower of a finite centreless group terminates in countably many steps. However, there does not seem to be an easy reduction from countable to finite; and it appears that some form of Wielandt's analysis is necessary.

3.5. Notes

My account of clubs and stationary sets is closely based on that in Kunen's textbook [**26**]. Theorem 3.1.13 first appeared in my 1985 paper [**49**]. Soon afterwards, Ulrich Felgner and I noticed independently that Fodor's Lemma implied that the automorphism tower of an infinite centreless group actually terminates after strictly less than $\left(2^{|G|}\right)^+$ steps. While I gave a direct proof, which later appeared in [**50**], Felgner realised that Corollary 3.3.2 was an immediate consequence of Faber's work [**7**] on infinite ascendent subgroups. I am very grateful to Felgner for pointing out the connection between Faber's work and the automorphism tower problem. (Of course, I am even more grateful that nobody noticed this connection before the publication of [**49**].) Theorem 3.4.1 was also noticed independently by Felgner. Theorem 3.2.1 first appeared in Just-Shelah-Thomas [**22**].

CHAPTER 4

The Normaliser Tower Technique

Much of this book will be concerned with the problem of constructing centreless groups with extremely long automorphism towers. Unfortunately it is usually very difficult to compute the successive groups in the automorphism tower of a centreless group. For example, we saw in Section 1.4 that it is already a nontrivial task just to compute the automorphism tower of the infinite dihedral group. In this chapter, we shall introduce the normaliser tower technique, which will enable us to entirely bypass this problem. Instead, throughout this book, we shall only have to deal with the much easier problem of computing the successive normalisers of a subgroup H of a group G.

In Section 4.1, after defining the notion of a normaliser tower, we shall reduce the problem of constructing centreless groups with long automorphism towers to that of finding long normaliser towers within the automorphism groups of first-order structures. Our reduction will make use of two "coding theorems", which will be proved in the remaining sections of this chapter. In Section 4.2, we shall prove that if $\mathcal{M}$ is an arbitrary first-order structure, then there exists a graph Γ such that $\mathrm{Aut}(\Gamma) \simeq \mathrm{Aut}(\mathcal{M})$; and in Section 4.3, we shall prove that if Γ is any graph, then there exists a field K_Γ such that $\mathrm{Aut}(K_\Gamma) \simeq \mathrm{Aut}(\Gamma)$. Section 4.4 contains the proof of a technical field-theoretic lemma which is needed in Section 4.3. In some of the later chapters of this book, we shall make use of the observation that the construction of the field K_Γ in Section 4.3 is upwards absolute. (The notion of "upwards absoluteness" will be defined in Chapter 6.) So it is important that the reader should at least read the definition of the field K_Γ in Section 4.3. However, the reader will not experience any disadvantage in understanding the rest of this book if he simply skips Section 4.4.

4.1. Normaliser towers

In this section, we shall reduce the problem of constructing centreless groups with long automorphism towers to that of finding long normaliser towers within the automorphism groups of first-order structures.

DEFINITION 4.1.1. If H is a subgroup of the group G, then the *normaliser tower* of H in G is defined inductively as follows.

(a) $N_0(H) = H$.
(b) If $\alpha = \beta + 1$, then $N_\alpha(H) = N_G(N_\beta(H))$.
(c) If α is a limit ordinal, then $N_\alpha(H) = \bigcup_{\beta<\alpha} N_\beta(H)$.

The *height* of the normaliser tower of H in G is the least ordinal α such that $N_\alpha(H) = N_{\alpha+1}(H)$. When it is necessary for the notation to include an explicit reference to the ambient group G, we shall write $N_\alpha(H) = N_\alpha(H, G)$.

The definition of the normaliser tower is motivated by the following observation, which says that automorphism towers can be regarded as special cases of normaliser towers.

PROPOSITION 4.1.2. *Let G be a centreless group and let $\tau = \tau(G)$. Then $N_{G_\tau}(G_\alpha) = G_{\alpha+1}$ for all $\alpha < \tau$. Hence the automorphism tower of G coincides with the normaliser tower of G in $G_{\tau(G)}$.*

PROOF. Let $\alpha < \tau$. Since the inclusion $G_\alpha \leqslant G_{\alpha+1}$ is isomorphic to the inclusion $\mathrm{Inn}(G_\alpha) \leqslant \mathrm{Aut}(G_\alpha)$, it follows that $G_{\alpha+1} \leqslant N_{G_\tau}(G_\alpha)$. Conversely, suppose that $g \in N_{G_\tau}(G_\alpha)$. Then g induces an automorphism of G_α via conjugation, and so there exists $h \in G_{\alpha+1}$ such that $hxh^{-1} = gxg^{-1}$ for all $x \in G_\alpha$. Thus $h^{-1}g \in C_{G_\tau}(G_\alpha) = 1$ and so $g = h \in G_{\alpha+1}$. Hence $N_{G_\tau}(G_\alpha) = G_{\alpha+1}$. □

As we shall see later in this section, if α is any ordinal, then it is easy to construct examples of pairs of groups, $H \leqslant G$, such that the normaliser tower of H in G terminates after exactly α steps. The following results will enable us to convert arbitrary normaliser towers into corresponding automorphism towers.

THEOREM 4.1.3. *Let S be a simple nonabelian group and let G, H be groups such that* $\mathrm{Inn}(S) \leqslant G, H \leqslant \mathrm{Aut}(S)$. *If $\pi : G \to H$ is an isomorphism, then there* *exists* $\varphi \in \mathrm{Aut}(S)$ *such that*

$$\pi(g) = \varphi g \varphi^{-1}$$

for all $g \in G$.

PROOF. By Theorem 1.2.8, $\mathrm{Inn}(S)$ is the unique minimal nontrivial normal subgroup of both G and H. It follows that $\pi[\mathrm{Inn}(S)] = \mathrm{Inn}(S)$ and hence there exists $\varphi \in \mathrm{Aut}(S)$ such that

$$\pi(c) = \varphi c \varphi^{-1}$$

for all $c \in \mathrm{Inn}(S)$. Now let $g \in G$ be an arbitrary element. Then for all $c \in \mathrm{Inn}(S)$, we have that

$$\begin{aligned}
(\varphi g)c(\varphi g)^{-1} &= \varphi(gcg^{-1})\varphi^{-1} \\
&= \pi(gcg^{-1}) \quad \text{ since } gcg^{-1} \in \mathrm{Inn}(S), \\
&= \pi(g)\pi(c)\pi(g)^{-1} \\
&= \pi(g)\varphi c \varphi^{-1}\pi(g)^{-1} \\
&= (\pi(g)\varphi)c(\pi(g)\varphi)^{-1}.
\end{aligned}$$

Since $C_{\mathrm{Aut}(S)}(\mathrm{Inn}(S)) = 1$, it follows that $\varphi g = \pi(g)\varphi$ and hence $\pi(g) = \varphi g \varphi^{-1}$. $\square$

THEOREM 4.1.4. *Let* S *be a simple nonabelian group and let* G *be a group such that* $\mathrm{Inn}(S) \leqslant G \leqslant \mathrm{Aut}(S)$. *Then the automorphism tower of* G *coincides with the normaliser tower of* G *in* $\mathrm{Aut}(S)$.

PROOF. Clearly it is enough to show that if G is an arbitrary group such that $\mathrm{Inn}(S) \leqslant G \leqslant \mathrm{Aut}(S)$, then the inclusion $\mathrm{Inn}(G) \trianglelefteq \mathrm{Aut}(G)$ is naturally isomorphic to the inclusion $G \trianglelefteq N_{\mathrm{Aut}(S)}(G)$.

Applying Theorem 4.1.3 in the special case when $G = H$, we see that for every automorphism $\pi \in \mathrm{Aut}(G)$, there exists a corresponding element $\varphi_\pi \in \mathrm{Aut}(S)$ such that

$$\pi(g) = \varphi_\pi g \varphi_\pi^{-1}$$

for all $g \in G$. Furthermore, since $C_{\mathrm{Aut}(S)}(\mathrm{Inn}(S)) = 1$, it follows that there exists a unique such element φ_π. Consider the homomorphism $\theta : \mathrm{Aut}(G) \to \mathrm{Aut}(S)$ defined by $\theta(\pi) = \varphi_\pi$. For each $a \in G$, let $i_a \in \mathrm{Inn}(G)$ be the corresponding inner automorphism. Then it is clear that $\theta(i_a) = a$ for all $a \in G$. In particular,

$\ker\theta \cap \mathrm{Inn}(G) = 1$ and hence Lemma 1.3.8 implies that θ is an embedding. Since $\mathrm{Inn}(G) \trianglelefteq \mathrm{Aut}(G)$ and $\theta[\mathrm{Inn}(G)] = G$, it follows that $\theta[\mathrm{Aut}(G)] \leqslant N_{\mathrm{Aut}(S)}(G)$. But every element of $N_{\mathrm{Aut}(S)}(G)$ induces an automorphism of G via conjugation and hence $\theta[\mathrm{Aut}(G)] = N_{\mathrm{Aut}(S)}(G)$. □

COROLLARY 4.1.5. *Let κ be an infinite cardinal and let G be a group such that $\mathrm{Alt}(\kappa) \leqslant G \leqslant \mathrm{Sym}(\kappa)$. Then G is a centreless group and the automorphism tower of G coincides with the normaliser tower of G in $\mathrm{Sym}(\kappa)$.*

PROOF. In Example 1.3.6, we showed that the inclusion

$$\mathrm{Inn}(\mathrm{Alt}(\kappa)) \trianglelefteq \mathrm{Aut}(\mathrm{Alt}(\kappa))$$

is naturally isomorphic to the inclusion $\mathrm{Alt}(\kappa) \trianglelefteq \mathrm{Sym}(\kappa)$. So the result follows from Theorem 4.1.4. □

While Corollary 4.1.5 is certainly suggestive, it appears to be difficult to use it to construct examples of long automorphism towers. Instead it is more useful to apply Theorem 4.1.4 to the simple groups $S = PSL(2, K)$ for suitably chosen fields K.

THEOREM 4.1.6. *Let K be a field such that $|K| > 3$ and let H be a subgroup of $\mathrm{Aut}(K)$. Let*

$$G = PGL(2, K) \rtimes H \leqslant P\mathit{\Gamma}L(2, K) = PGL(2, K) \rtimes \mathrm{Aut}(K).$$

Then G is a centreless group and for each ordinal α,

$$G_\alpha = PGL(2, K) \rtimes N_\alpha(H),$$

where $N_\alpha(H)$ is the αth group in the normaliser tower of H in $\mathrm{Aut}(K)$.

PROOF. Since $|K| > 3$, it follows that $PSL(2, K)$ is a simple group. (For example, see Section 1.9 of Suzuki [**48**].) By a well-known theorem of Schreier and van der Waerden [**41**], every automorphism of $PSL(2, K)$ is induced via conjugation by an element of $P\mathit{\Gamma}L(2, K)$; and so the inclusion

$$\mathrm{Inn}(PSL(2, K)) \trianglelefteq \mathrm{Aut}(PSL(2, K))$$

is naturally isomorphic to the inclusion $PSL(2,K) \trianglelefteq P\Gamma L(2,K)$. Hence Lemma 1.1.2 implies that $C_{P\Gamma L(2,K)}(PSL(2,K)) = 1$ and it follows that G is centreless. Finally it is easily checked that if $H \leqslant \mathrm{Aut}(K)$, then

$$N_{P\Gamma L(2,K)}(PGL(2,K) \rtimes H) = PGL(2,K) \rtimes N_{\mathrm{Aut}(K)}(H).$$

So the result follows from Theorem 4.1.4. □

It is well-known that every group G can realised as the automorphism group of a suitable graph Γ. (We shall prove this result in Section 4.2.) So the following result implies that every group G can also be realised as the automorphism group of a suitable field K.

THEOREM 4.1.7 (Fried and Kollár [**10**]). *Let $\Gamma = \langle X, E\rangle$ be any graph. Then there exists a field K_Γ of cardinality* $\max\{|X|, \omega\}$ *which satisfies the following conditions.*

(a) *X is an $\mathrm{Aut}(K_\Gamma)$-invariant subset of K_Γ.*
(b) *The restriction mapping, $\pi \mapsto \pi \upharpoonright X$, is an isomorphism from $\mathrm{Aut}(K_\Gamma)$ onto $\mathrm{Aut}(\Gamma)$.*

We shall prove Theorem 4.1.7 in Section 4.3. Combining Theorems 4.1.6 and 4.1.7, we have now reduced our problem to that of finding long normaliser towers within the automorphism groups of suitably chosen graphs. In fact, we can do slightly better than this. The following theorem, which will be proved in Section 4.2, reduces our problem to that of finding long normaliser towers within the automorphism groups of arbitrary first-order structures.

THEOREM 4.1.8. *Let $\mathcal{M}$ be a structure for the first-order language L and suppose that $\kappa \geq \max\{|\mathcal{M}|, |L|, \omega\}$. Then there exists a graph Γ of cardinality κ such that $\mathrm{Aut}(\mathcal{M}) \simeq \mathrm{Aut}(\Gamma)$.*

THEOREM 4.1.9. *Let $\mathcal{M}$ be a structure for the first-order language L and let H be a subgroup of $\mathrm{Aut}(\mathcal{M})$. Suppose that*

(a) *$\kappa \geq \max\{|\mathcal{M}|, |L|, |H|, \omega\}$; and*
(b) *the normaliser tower of H in $\mathrm{Aut}(\mathcal{M})$ terminates after exactly α steps.*

Then there exists a centreless group G of cardinality κ such that $\tau(G) = \alpha$.

Proof of Theorem 4.1.9. By Theorem 4.1.8, we can assume that $\mathcal{M}$ is a graph Γ of cardinality κ. By Theorem 4.1.7, there exists a field K_Γ of cardinality κ such that $\mathrm{Aut}(K_\Gamma) \simeq \mathrm{Aut}(\Gamma)$. To simplify notation, identify $\mathrm{Aut}(K_\Gamma)$ with $\mathrm{Aut}(\Gamma)$ and let $G = PGL(2, K_\Gamma) \rtimes H$. By Theorem 4.1.6, G is a centreless group of cardinality κ such that $\tau(G) = \alpha$. □

Next for each ordinal α, we shall construct a pair of groups, $H \leqslant G$, such that the normaliser tower of H in G terminates after exactly α steps.

Definition 4.1.10. The ascending chain of groups

$$W_0 \leqslant W_1 \leqslant \cdots \leqslant W_\alpha \leqslant W_{\alpha+1} \leqslant \dots$$

is defined inductively as follows.

(a) $W_0 = C_2$, the cyclic group of order 2.

(b) Suppose that $\alpha = \beta + 1$. Then

$$W_\beta = W_\beta \oplus 1 \leqslant \left[W_\beta \oplus W_\beta^*\right] \rtimes \langle \sigma_{\beta+1} \rangle = W_{\beta+1}.$$

Here W_β^* is an isomorphic copy of W_β; and $\sigma_{\beta+1}$ is an element of order 2 which interchanges the factors $W_\beta \oplus 1$ and $1 \oplus W_\beta^*$ of the direct sum $W_\beta \oplus W_\beta^*$ via conjugation. Thus $W_{\beta+1}$ is isomorphic to the wreath product $W_\beta \mathrm{\ Wr\ } C_2$.

(c) If α is a limit ordinal, then $W_\alpha = \bigcup_{\beta<\alpha} W_\beta$.

Lemma 4.1.11. $|W_\alpha| \leq \max\{|\alpha|, \omega\}$ *for all ordinals* α.

Proof. This follows by an easy induction on α. □

Lemma 4.1.12. (a) *If* $1 \leq n < \omega$, *then the normaliser tower of* W_0 *in* W_n *terminates after exactly* $n + 1$ *steps.*

(b) *If* $\alpha \geq \omega$, *then the normaliser tower of* W_0 *in* W_α *terminates after exactly* α *steps.*

Proof. (a) It is easily checked that

$$N_1(W_0, W_n) = W_0 \oplus W_0^* \oplus W_1^* \oplus \cdots \oplus W_{n-1}^*$$

and that

$$N_2(W_0, W_n) = W_1 \oplus W_1^* \oplus \cdots \oplus W_{n-1}^*;$$

and that, in general, for each $0 \le \ell \le n-1$,

$$N_{\ell+1}(W_0, W_n) = W_\ell \oplus \bigoplus_{\ell \le m < n} W_m^*.$$

(b) For example, consider the case when $\alpha > \omega$. Then for each $\ell \in \omega$,

$$N_{\ell+1}(W_0, W_\alpha) = W_\ell \oplus \bigoplus_{\ell \le \beta < \alpha} W_\beta^*;$$

and for each γ such that $\omega \le \gamma < \alpha$,

$$N_\gamma(W_0, W_\alpha) = W_\gamma \oplus \bigoplus_{\gamma \le \beta < \alpha} W_\beta^*.$$

□

THEOREM 4.1.13. *Suppose that α is any ordinal and that $\kappa \ge \max\{|\alpha|, \omega\}$. Then there exists a centreless group G of cardinality κ such that $\tau(G) = \alpha$.*

PROOF. First we shall find a pair of groups, $H \leqslant W$, such that

(i) $|W| \le \max\{|\alpha|, \omega\}$, and
(ii) the normaliser tower of H in W terminates after exactly α steps.

Applying Lemmas 4.1.11 and 4.1.12, we see that if $\alpha \ge \omega$, then we can take $H = W_0$ and $W = W_\alpha$; and if $2 \le \alpha < \omega$, then we can take $H = W_0$ and $W = W_{\alpha-1}$. This just leaves the cases when $\alpha = 0$, 1. When $\alpha = 0$, then we can take $H = W = C_2$; and when $\alpha = 1$, we can take $H = \mathrm{Alt}(3)$ and $W = \mathrm{Sym}(3)$.

Next we shall construct a first-order structure $\mathcal{M}$ such that $\mathrm{Aut}(\mathcal{M}) \simeq W$. For each $w \in W$, let R_w be a binary relation symbol and let L be the first-order language $\{R_w \mid w \in W\}$. Let

$$\mathcal{M} = \left(W; R_w^{\mathcal{M}}\right)_{w \in W}$$

be the L-structure such that $R_w^{\mathcal{M}} = \{(x, xw) \mid x \in W\}$ for each $w \in W$. For each $g \in W$, let $\lambda_g \in \mathrm{Sym}(W)$ be the left multiplication map, defined by $\lambda_g(x) = gx$ for all $x \in W$. Then it is easily checked that $\lambda_g \in \mathrm{Aut}(\mathcal{M})$ for all $g \in W$. We claim that

$$\mathrm{Aut}(\mathcal{M}) = \{\lambda_g \mid g \in W\}.$$

To see this, let $\pi \in \mathrm{Aut}(\mathcal{M})$ be any automorphism and let $g = \pi(1)$. If $x \in W$, then $(1, x) \in R_x$ and so $(g, \pi(x)) \in R_x$. Thus $\pi(x) = gx$ for all $x \in W$ and so $\pi = \lambda_g$.

By Theorem 4.1.9, there exists a centreless group G of cardinality κ such that $\tau(G) = \alpha$. $\square$

COROLLARY 4.1.14. *If κ is an infinite cardinal, then $\tau_\kappa \geq \kappa^+$.*

$\square$

4.2. Coding structures in graphs

In this section, we shall prove Theorem 4.1.8. We shall begin by showing how an arbitrary structure $\mathcal{M}$ can be coded within a structure $\mathcal{M}'$ for a suitable countable language.

LEMMA 4.2.1. *Let $\mathcal{M} = \langle M, \dots \rangle$ be a structure for the first-order language L and let $\kappa \geq \max\{|\mathcal{M}|, |L|, \omega\}$. Then there exists a countable first-order language L' and a structure $\mathcal{M}'$ for L' of cardinality κ such that the following conditions are satisfied.*

(a) *M is an $\mathrm{Aut}(\mathcal{M}')$-invariant subset of $\mathcal{M}'$.*
(b) *The restriction map, $\pi \mapsto \pi \restriction M$, is an isomorphism from $\mathrm{Aut}(\mathcal{M}')$ onto $\mathrm{Aut}(\mathcal{M})$.*

PROOF. We can suppose that L is a relational first-order language. (If f is an n-ary function symbol, then we replace it by an $(n+1)$-ary relation symbol R which represents the graph of the associated function $f^{\mathcal{M}}$; and if c is a constant symbol, then we replace it by a unary relation symbol C which represents the subset $\{c^{\mathcal{M}}\}$.) Let $L = \bigcup_{n\geq 1} L_n$, where $L_n = \{R_{n,\alpha} \mid \alpha < \lambda_n\}$ is the set of n-ary relations symbols of L. Let $L' = \{P, <\} \cup \{S_n \mid n \geq 1\}$ be a first-order language such that

(i) P is a unary relation symbol;
(ii) $<$ is a binary relation symbol; and
(iii) for each $n \geq 1$, S_n is an $(n+1)$-ary relation symbol.

Let $\mathcal{M}'$ be the structure for the language L' defined as follows.

(1) The universe of $\mathcal{M}'$ is the disjoint union $M \sqcup \kappa$.
(2) $P^{\mathcal{M}'} = \kappa$.
(3) $<^{\mathcal{M}'}$ is the usual well-ordering of κ.
(4) For each $n \geq 1$, $\langle a_1, \dots, a_n, \alpha\rangle \in S_n^{\mathcal{M}'}$ iff $a_1, \dots, a_n \in M$, $\alpha < \lambda_n$ and $\langle a_1, \dots, a_n\rangle \in R_{n,\alpha}^{\mathcal{M}}$.

Clearly M is an $\mathrm{Aut}(\mathcal{M}')$-invariant subset of $\mathcal{M}'$; and since the well-ordered set κ is rigid, $\pi \restriction \kappa = id_\kappa$ for each $\pi \in \mathrm{Aut}(\mathcal{M}')$. It follows easily that $\mathcal{M}'$ satisfies our requirements. □

Next we shall show how a structure $\mathcal{M}$ for a countable language can be coded within a suitably chosen *connected* graph Γ. (In Section 4.3, we shall present Fried and Kollár's coding of graphs within fields. As we shall see, their argument requires the additional hypothesis that the encoded graph Γ has no isolated vertices. However, if the graph Γ happens to have isolated vertices, then we can first use Lemma 4.2.2 to code Γ within a connected graph Γ^+ and then apply the construction of Fried and Kollár to Γ^+.) The proof of the following lemma is essentially a proper subset of that of Hodges [**16**, Theorem 5.5.1].

LEMMA 4.2.2. *Let $\mathcal{M} = \langle M, \dots \rangle$ be a structure for the countable first-order language L and let $\kappa \geq \max\{|\mathcal{M}|, \omega\}$. Then there exists a connected graph $\Gamma_{\mathcal{M}}$ of cardinality κ such that the following conditions are satisfied.*

(a) *M is an $\mathrm{Aut}(\Gamma_{\mathcal{M}})$-invariant subset of $\Gamma_{\mathcal{M}}$.*
(b) *The restriction map, $\pi \mapsto \pi \restriction M$, is an isomorphism from $\mathrm{Aut}(\Gamma_{\mathcal{M}})$ onto $\mathrm{Aut}(\mathcal{M})$.*

PROOF. We can suppose that $\mathcal{M}$ has cardinality κ. (If $|\mathcal{M}| < \kappa$, then we can first use Lemma 4.2.1 to code $\mathcal{M}$ within a structure $\mathcal{M}'$ of cardinality κ.) As in the proof of Lemma 4.2.1, we can suppose that L is a relational first-order language; say $L = \{S_n \mid 1 \leq n < \lambda\}$, where $1 \leq \lambda \leq \omega$ and each S_n is an r_n-ary relation symbol. In order to ensure that our construction always produces a connected graph, we shall also assume that S_1 is a binary relation symbol which represents the "trivial connected relation" $S_1^{\mathcal{M}} = \{\langle a, b\rangle \in M^2 \mid a \neq b\}$.

We shall begin by describing the main ingredients of our coding construction. Let Γ be a graph and let $n \geq 3$. Then a vertex v of Γ is said to be *n-tagged* if Γ

contains an induced subgraph of the form:

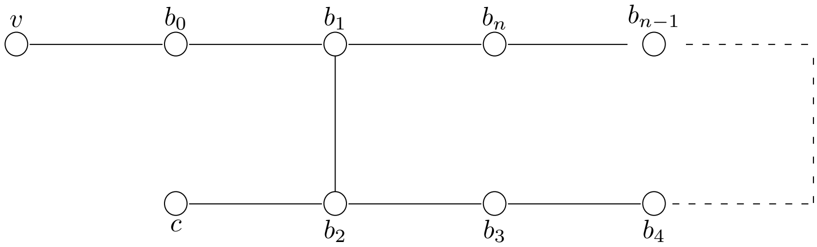

In this case, we say that the induced subgraph $\{c, b_0, \dots, b_n\}$ is an *n-tag* of v. Note that the induced subgraph $\{v, c, b_0, \dots, b_n\}$ is rigid.

Next we shall describe the construction of the graph $\Gamma_{\mathcal{M}}$. First we regard M as a null graph. Now for each $1 \leq n < \lambda$ and each tuple $\langle a_1, \dots a_{r_n} \rangle \in S_n^{\mathcal{M}}$, we add a corresponding finite graph $\Gamma_n(a_1, \dots, a_{r_n})$ as follows.

(i) First we adjoin a vertex d, together with an $(n+2)$-tag T of d.
(ii) Then we adjoin pairwise disjoint paths p_ℓ of length $\ell + 1$ from a_ℓ to d for each $1 \leq \ell \leq r_n$.

For example, if $a_1 \neq a_2 \in M$, then the graph $\Gamma_1(a_1, a_2)$ has the form:

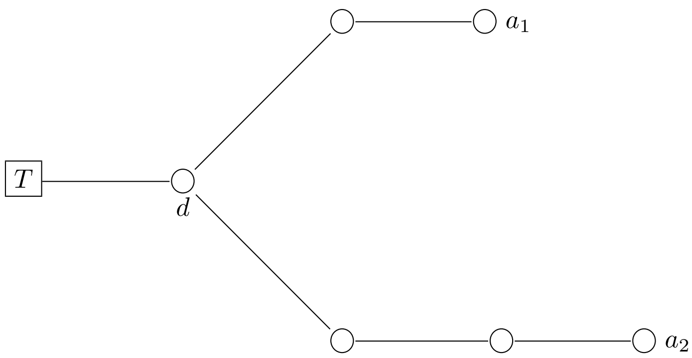

Note that each of the induced subgraphs $\Gamma_n(a_1, \dots, a_{r_n})$ is also rigid. This completes the construction of $\Gamma_{\mathcal{M}}$.

Finally we shall check that $\Gamma_{\mathcal{M}}$ satisfies our requirements. First notice that since $S_1^{\mathcal{M}} = \{\langle a, b \rangle \in M^2 \mid a \neq b\}$, it follows both that $\Gamma_{\mathcal{M}}$ is connected and that each vertex $m \in M$ has infinite valency in $\Gamma_{\mathcal{M}}$. By construction, each vertex $v \notin M$ has finite valency. Hence M is an $\mathrm{Aut}(\Gamma_{\mathcal{M}})$-invariant subset of $\Gamma_{\mathcal{M}}$. It is clear that every automorphism $\varphi \in \mathrm{Aut}(\mathcal{M})$ extends to an automorphism $\pi \in \mathrm{Aut}(\Gamma_{\mathcal{M}})$. Thus it only remains to show that if $\pi \in \mathrm{Aut}(\Gamma_{\mathcal{M}})$ satisfies $\pi \restriction M = id_M$, then $\pi = id_{\Gamma_{\mathcal{M}}}$. So consider any vertex $v \notin M$. Then v lies in a π-invariant induced subgraph Γ_0 of the form $\Gamma_n(a_1, \dots, a_{r_n})$ for some $n \geq 1$ and $a_1, \dots, a_{r_n} \in M$. Since Γ_0 is rigid, it follows that $\pi(v) = v$. $\square$

Clearly Theorem 4.1.8 is an immediate consequence of Lemmas 4.2.1 and 4.2.2. We shall end this section with some observations which will be needed in Chapter 8. Suppose that $L = \{S_n \mid 1 \leq n < \lambda\}$ is a countable relational first-order language. Let $\mathcal{C}$ be the category of infinite structures for L and let $\mathcal{G}$ be the category of graphs. For each $\mathcal{M} \in \mathcal{C}$, let $\Gamma_{\mathcal{M}}$ be the corresponding graph of cardinality $|\mathcal{M}|$, constructed as in the proof of Lemma 4.2.2. Then if $\mathcal{M}_1$, $\mathcal{M}_2 \in \mathcal{C}$,

(a) $\mathcal{M}_1 \simeq \mathcal{M}_2$ iff $\Gamma_{\mathcal{M}_1} \simeq \Gamma_{\mathcal{M}_2}$; and

(b) $\mathrm{Aut}(\mathcal{M}_1) \simeq \mathrm{Aut}(\Gamma_{\mathcal{M}_1})$.

In Chapter 8, we shall also make use of the fact that the construction of $\Gamma_{\mathcal{M}}$ from $\mathcal{M}$ is upwards absolute. (The notion of "upwards absoluteness" will be defined in Chapter 6.) Hence properties (a) and (b) will continue to hold in any generic extension of the set-theoretic universe.

4.3. Coding graphs in fields

In this section, we shall prove Theorem 4.1.7. Let $\Gamma = \langle X, E \rangle$ be any graph. We shall construct a field K_Γ of cardinality $\max\{|X|, \omega\}$ which satisfies the following conditions.

(a) X is an $\mathrm{Aut}(K_\Gamma)$-invariant subset of K_Γ.

(b) The restriction mapping, $\pi \mapsto \pi \restriction X$, is an isomorphism from $\mathrm{Aut}(K_\Gamma)$ onto $\mathrm{Aut}(\Gamma)$.

If the Γ happens to have isolated vertices, then we first use Lemma 4.2.2 to code Γ within a connected graph Γ^+. To simplify notation, we shall assume that Γ has no isolated vertices.

Definition 4.3.1. Let F be a field of characteristic 0 and let p be a prime. Let S be a set of algebraically independent elements over F and let $F(S)$ be the corresponding purely transcendental extension of F. Then $F(S)(S, p)$ denotes the field which is obtained by adjoining the elements $\{s(\ell) \mid s \in S, \ell \in \omega\}$ to $F(S)$, where

(a) $s(0) = s$, and

(b) $s(\ell+1)^p = s(\ell)$.

Note that the field $F(S)(S,p)$ is uniquely defined up to isomorphism. Let $(p_n \mid n \in \omega)$ be the increasing sequence of odd primes. We shall construct K_Γ as the union of an increasing chain $\{K_n(\Gamma) \mid n \in \omega\}$ of fields. At each stage of the construction, we shall also define a distinguished subset $H_n(\Gamma)$ of $K_n(\Gamma)$. To begin the construction, regard the set X of vertices of Γ as a set of algebraically independent elements over the field $\mathbb{Q}$ of rational numbers; and let

$$K_0(\Gamma) = \mathbb{Q}(X)(X, p_0)$$

and

$$H_0(\Gamma) = \{u + v \mid (u,v) \in E\}.$$

Suppose inductively that we have defined $K_n(\Gamma)$ and $H_n(\Gamma)$. First let t_n be a transcendental element over $K_n(\Gamma)$ and let $L_n = K_n(\Gamma)(t_n)(\{t_n\}, p_{n+1})$. Then let $K_{n+1}(\Gamma)$ be the splitting field over L_n of the set of polynomials

$$P_n = \{y^2 - (t_n - a) \mid a \in H_n(\Gamma)\}$$

and let $H_{n+1}(\Gamma)$ be a set which contains exactly one root of each of the polynomials in P_n. Finally let $K_\Gamma = \bigcup_{n\in\omega} K_n(\Gamma)$.

The idea behind this construction is easily explained. In the first two stages of the construction, we attempt to encode the graph Γ within the field $K_1(\Gamma)$. Later we shall see that $K_1(\Gamma)$ satisfies the following conditions.

(1) Both X and $\{t_0\}$ are $\mathrm{Aut}(K_1(\Gamma))$-invariant subsets of $K_1(\Gamma)$.
(2) The restriction mapping, $\pi \mapsto \pi \upharpoonright X$, is a surjective homomorphism from $\mathrm{Aut}(K_1(\Gamma))$ onto $\mathrm{Aut}(\Gamma)$.

However, the restriction map is *not* injective. While each automorphism $\varphi \in \mathrm{Aut}(\Gamma)$ can be extended uniquely to an automorphism $\theta \in \mathrm{Aut}(L_0)$ satisfying $\theta(t_0) = t_0$, there are then many ways to extend θ to an automorphism of the splitting field $K_1(\Gamma)$ of the set of polynomials

$$P_0 = \{y^2 - (t_0 - a) \mid a \in H_0(\Gamma)\}.$$

So at the next stage of the construction, we use $K_2(\Gamma)$ to pick out a distinguished root of each of the polynomials in P_0. Iterating this procedure ω times, we obtain a field K_Γ which satisfies all of our requirements.

It is relatively straightforward to show that every automorphism $\varphi \in \mathrm{Aut}(\Gamma)$ lifts inductively through automorphisms of the $K_n(\Gamma)$ to an automorphism of K_Γ. The only thing that needs to be checked is that we cannot accidently block an automorphism by an "unfortunate choice" of the roots in some $H_{n+1}(\Gamma)$; and Lemma 4.3.2 will deal with this point. On the other hand, it will require a substantial effort to show that all automorphisms of K_Γ arise in this manner; and Lemma 4.3.2 will also play an important role in the proof of this result.

LEMMA 4.3.2. *Let F be a field of characteristic 0 and let t be a transcendental element over F. Suppose that $T_1, \dots, T_n$ are mutually prime nonconstant polynomials in $F[t]$, each having no multiple factors; and for each $1 \leq i \leq n$, let ϑ_i be an element such that $\vartheta_i^2 = T_i$. Let $F_0 = F(t)$ and for each $1 \leq i \leq n$, let $F_i = F(t, \vartheta_1, \dots, \vartheta_i)$. Then the following statements hold for each $1 \leq i \leq n$.*

(a) $\vartheta_i \notin F_{i-1}$.
(b) *If $\eta \in F_i$ satisfies $\eta^2 \in F_0$, then there exist an element $c \in F_0$ and a subset Z of $\{1, \dots, i\}$ such that $\eta = c \prod_{\ell \in Z} \vartheta_\ell$.*
(c) *If $\eta \in F_i$ and η is algebraic over F, then $\eta \in F$.*

PROOF. First we shall prove by induction on $i \geq 1$ that both (a) and (b) hold. Clearly (a) holds when $i = 1$. Assume inductively that (a) holds for all $k \leq i$ and that (b) holds for all $k < i$. Suppose that $\eta \in F_i \smallsetminus F_{i-1}$ and that $\eta^2 = a \in F_0$. Let $\eta = b + c\vartheta_i$, where $b, c \in F_{i-1}$. Since $\vartheta_i \notin F_{i-1}$, there exists an automorphism $\pi \in \mathrm{Gal}(F_i/F_{i-1})$ such that $\pi(\vartheta_i) = -\vartheta_i$. Since $\pi(\eta) = b - c\vartheta_i$ is also a root of the polynomial $x^2 - a \in F_0[x]$, it follows that

$$b - c\vartheta_i = \pi(\eta) = -\eta = -b - c\vartheta_i.$$

Hence $\eta = c\vartheta_i$. Notice that $\eta/\vartheta_i = c \in F_{i-1}$ and that $(\eta/\vartheta_i)^2 = a/T_i \in F_0$. Hence by inductive hypothesis, there exists an element $d \in F_0$ and a subset Y of $\{1, \dots, i-1\}$ such that $\eta/\vartheta_i = d \prod_{\ell \in Y} \vartheta_\ell$. Thus $\eta = d\vartheta_i \prod_{\ell \in Y} \vartheta_\ell$ and (b) also holds for i.

Now suppose that $\vartheta_{i+1} \in F_i$. Since $\theta_{i+1}^2 = T_{i+1} \in F_0$, (b) implies that there exists relatively prime polynomials $f(t)$, $g(t) \in F[t]$ and a subset Z of $\{1, \dots, i\}$

such that $\vartheta_{i+1} = f(t)/g(t) \prod_{\ell \in Z} \vartheta_\ell$ and so

$$g^2(t)T_{i+1} = f^2(t) \prod_{\ell \in Z} T_\ell.$$

Since the polynomials $\{T_\ell \mid \ell \in Z\}$ are mutually prime and have no multiple factors, it follows that $g(t) \in F$. Similary we find that $f(t) \in F$ and then that $Z = \emptyset$. But this is a contradiction. Hence $\vartheta_{i+1} \notin F_i$ and (a) also holds for $i + 1$.

Finally we shall prove that (c) holds by induction on $i \geq 0$. Clearly (c) holds when $i = 0$. Assume inductively that (c) holds for all $k < i$. Suppose that $\eta = b + c\vartheta_i$ is algebraic over F, where b, $c \in F_{i-1}$. Then clearly t remains transcendental over $\widetilde{F} = F(\eta)$. Furthermore, the polynomials $T_1, \ldots, T_n$ are mutually prime nonconstant polynomials in $\widetilde{F}[t]$, each having no multiple factors. Let $\widetilde{F}_0 = \widetilde{F}(t)$ and for each $1 \leq i \leq n$, let $\widetilde{F}_i = \widetilde{F}(t, \vartheta_1, \ldots, \vartheta_i)$. Applying (a), we obtain that $\vartheta_i \notin \widetilde{F}_{i-1}$. Since $\eta = b + c\vartheta_i \in \widetilde{F} \subseteq \widetilde{F}_{i-1}$, this implies that $c = 0$. Thus $\eta = b \in F_{i-1}$. By inductive hypothesis, $\eta \in F$ and so (c) also holds for i. □

Corollary 4.3.3. *Suppose that $n \in \omega$ and $b \in K_\Gamma$. If b is algebraic over $K_n(\Gamma)$, then $b \in K_n(\Gamma)$.*

Proof. We shall prove that for each $n \in \omega$, if $b \in K_{n+1}(\Gamma)$ is algebraic over $K_n(\Gamma)$, then $b \in K_n(\Gamma)$. Then an easy inductive argument yields Corollary 4.3.3. Fix some $n \in \omega$. Let $p = p_{n+1}$ and $F = K_n(\Gamma)$. Suppose that $b \in K_{n+1}(\Gamma)$ is algebraic over $K_n(\Gamma)$. There exists a finite subset $\{a_i \mid 1 \leq i \leq s\}$ of $H_n(\Gamma)$ such that b is in the splitting field over L_n of the set of polynomials

$$\{y^2 - (t_n - a_i) \mid 1 \leq i \leq s\}.$$

For each $1 \leq i \leq s$, let $\vartheta_i \in K_{n+1}(\Gamma)$ be an element such that $\vartheta_i^2 = t_n - a_i$. Choose an integer k and an element $t \in L_n$ such that

(a) $t^{p^k} = t_n$, and

(b) $b \in F(t, \vartheta_1, \ldots, \vartheta_s)$.

For each $1 \leq i \leq s$, let $T_i = t^{p^k} - a_i \in F[t]$. Then $T_1, \ldots, T_s$ satisfy the hypotheses of Lemma 4.3.2. So applying Lemma 4.3.2(c), we obtain that $b \in F = K_n(\Gamma)$. □

In the proof of Lemma 4.3.8, we shall make use of the fact that -1 has no square root in K_Γ. This result is an immediate consequence of Corollary 4.3.3, together with the observation that -1 has no square root in $K_0(\Gamma)$.

We are now ready to prove that every automorphism $\varphi \in \mathrm{Aut}(\Gamma)$ extends to an automorphism of K_Γ.

LEMMA 4.3.4. *For each $\varphi \in \mathrm{Aut}(\Gamma)$, there exists a unique $\pi \in \mathrm{Aut}(K_\Gamma)$ such that the following conditions are satisfied.*

(a) $\pi \restriction X = \varphi$.
(b) $\pi(t_n) = t_n$ *for all* $n \in \omega$.
(c) $\pi[H_n(\Gamma)] = H_n(\Gamma)$ *for all* $n \in \omega$.

PROOF. We shall define the sequence of approximations $\pi_n = \pi \restriction K_n(\Gamma)$ by induction on $n \geq 0$. First note that there exists a unique $\pi_0 \in \mathrm{Aut}(K_0(\Gamma))$ such that $\pi_0 \restriction X = \varphi$. Now suppose that we have defined $\pi_n \in \mathrm{Aut}(K_n(\Gamma))$ such that that $\pi_n(t_k) = t_k$ for all $k < n$ and that $\pi_n[H_k(\Gamma)] = H_k(\Gamma)$ for all $k \leq n$. Clearly there is a unique extension π'_n to L_n such that $\pi'_n(t_n) = t_n$. For each $a \in H_n(\Gamma)$, let $r_a \in H_{n+1}(\Gamma)$ be the corresponding root of $y^2 - (t_n - a)$. Arguing as in the proof of Corollary 4.3.3, we see that $r_a \notin L_n(\{r_b \mid b \in H_n(\Gamma) \smallsetminus \{a\}\})$. It follows that there is a unique extension $\pi_{n+1} \in \mathrm{Aut}(K_{n+1}(\Gamma))$ such that $\pi_{n+1}(t_n) = t_n$ and $\pi_{n+1}(r_a) = r_{\pi_n(a)}$ for each $a \in H_n(\Gamma)$. $\square$

The rest of this section will be devoted to the proof that every automorphism of K_Γ arises in the above manner.

DEFINITION 4.3.5. Let F be a field and let p be a prime. Then a nonzero element $u \in F$ is said to be a *p-high* element of F if the equation $y^{p^n} = u$ is solvable in F for all $n \in \omega$.

In the hope of distinguishing these elements, each of the vertices of Γ and each of the transcendal elements $\{t_n \mid n \in \omega\}$ has been made p-high in K_Γ for a suitably chosen odd prime p. Of course, these will not be the only p-high elements of K_Γ. For example, since each vertex x of Γ is 3-high, it follows that $\pm x^{m/3^\ell}$ is also 3-high for each $m \in \mathbb{Z}$ and $\ell \in \mathbb{N}$. The next two lemmas tell us that every p-high element of K_Γ arises in essentially this manner. The rather technical proof of Lemma 4.3.6 will be given in Section 4.4.

LEMMA 4.3.6. *Let F be a field of characteristic 0 and let t be a transcendental element over F. Let p be an odd prime, and let $\{t(\ell) \mid \ell \in \mathbb{N}\}$ be a set of elements* *such that* $t(0) = t$ *and* $t(\ell+1)^p = t(\ell)$*. Let*

$$L = F(t(0), t(1), \dots, t(\ell), \dots) = F(t)(\{t\}, p).$$

Let $\{T_i \mid i \in I\}$ *be a set of mutually prime polynomials in* $F[t]$*, none of which is divisible by* t *or has a multiple factor; and for each* $i \in I$*, let* ϑ_i *be an element such that* $\vartheta_i^2 = T_i$*. Let* $M = L(\dots, \vartheta_i, \dots)$*.*

(a) *If* u *is a* p*-high element of* M*, then* $u = ct(\ell)^m$ *for some* $\ell \in \mathbb{N}$*,* $m \in \mathbb{Z}$ *and some* p*-high element* c *of* F*.*

(b) *If* p' *is an odd prime such that* $p' \neq p$ *and* u *is a* p'*-high element of* M*, then* u *is a* p'*-high element of* F*.*

Lemma 4.3.7. *Let* p *be an odd prime, and suppose that* u *is a* p*-high element of* K_Γ*.*

(a) *If* $p = p_0$*, then either* u *or* $-u$ *is a product of elements of the form* x^{m/p^ℓ}*, where* x *is a vertex of* Γ*,* $m \in \mathbb{Z}$ *and* $\ell \in \mathbb{N}$*.*

(b) *If* $p = p_{n+1}$ *for some* $n \in \omega$*, then either* u *or* $-u$ *is of the form* t_n^{m/p^ℓ}*, where* $m \in \mathbb{Z}$ *and* $\ell \in \mathbb{N}$*.*

Proof of Lemma 4.3.7. The result is an easy consequence of Lemma 4.3.6, together with the observation that ± 1 are the only p-high elements of $\mathbb{Q}$. $\square$

Lemma 4.3.8. *Let* $n \in \omega$ *and suppose that* $b \in K_\Gamma$ *satifies the equation*

$$b^2 = e(t_n^r - a),$$

where $e^2 = 1$*,* $a \in K_n(\Gamma) \smallsetminus \{0\}$*, and* $r = m/p_{n+1}^\ell$ *for some* $m \in \mathbb{Z} \smallsetminus \{0\}$ *and* $\ell \in \mathbb{N}$*. Then* $e = r = 1$ *and* $a \in H_n(\Gamma)$*.*

Proof. Let $p = p_{n+1}$ and $F = K_n(\Gamma)$. By Corollary 4.3.3, $b \in K_{n+1}(\Gamma)$. Thus there exists a finite subset $\{a_i \mid 1 \leq i \leq s\}$ of $H_n(\Gamma)$ such that b is in the splitting field over L_n of the set of polynomials $\{y^2 - (t_n - a_i) \mid 1 \leq i \leq s\}$. For each $1 \leq i \leq s$, let $\vartheta_i \in K_{n+1}(\Gamma)$ be an element such that $\vartheta_i^2 = t_n - a_i$. Choose an integer $k > \ell$ and an element $t \in L_n$ such that

(a) $t^{p^k} = t_n$, and

(b) $b \in F(t, \vartheta_1, \dots, \vartheta_s)$.

For each $1 \leq i \leq s$, let $T_i = t^{p^k} - a_i \in F[t]$. Then $T_1, \dots, T_s$ satisfy the hypotheses of Lemma 4.3.2. Note that

$$b^2 = e(t_n^r - a) = e(t^{mp^{k-\ell}} - a) \in F(t).$$

By Lemma 4.3.2, there exist polynomials $f(t)$, $g(t) \in F[t]$ and a subset Z of $\{1, \dots, s\}$ such that

$$b = \frac{f(t)}{g(t)} \prod_{i \in Z} \vartheta_i.$$

It follows that

$$g^2(t)e(t^{mp^{k-\ell}} - a) = f^2(t) \prod_{i \in Z} (t^{p^k} - a_i).$$

First suppose that $m > 0$. Then we must have that $eg^2(t) = f^2(t)$, since the other factors do not have multiple roots. Hence

$$(t^{mp^{k-\ell}} - a) = \prod_{i \in Z} (t^{p^k} - a_i).$$

We claim that $|Z| = 1$. To see this, suppose that $|Z| = z \geq 2$; say, $Z = \{i_1, \dots, i_z\}$. By considering the coefficient of $t^{(z-1)p^k}$ in the expansion of $\prod_{i \in Z}(t^{p^k} - a_i)$, we find that $\sum_{i \in Z} a_i = 0$. But this is impossible. For if $n = 0$, then each a_i has the form $u + v$ for some u, $v \in X$; and if $n > 0$, then we have already noted that the proof of Corollary 4.3.3 shows that $a_{i_z} \notin L_{n-1}(a_{i_1}, \dots, a_{i_{z-1}})$. Thus $|Z| = 1$ and so $t^{mp^{k-\ell}} - a = t^{p^k} - a_i$ for some $1 \leq i \leq s$. Consequently, $r = m/p^\ell = 1$ and $a = a_i \in H_n(\Gamma)$. Also since $e = f^2(t)/g^2(t)$ and -1 has no square root in K_Γ, it follows that $e = 1$.

Now suppose that $m < 0$. Then multiplying by $t^{-mp^{k-\ell}}$, we obtain that

$$g^2(t)e(1 - at^{-mp^{k-\ell}}) = f^2(t) \prod_{i \in Z} (t^{p^k} - a_i) t^{-mp^{k-\ell}}$$

and so

$$-eag^2(t)(t^{-mp^{k-\ell}} - a^{-1}) = f^2(t) \prod_{i \in Z} (t^{p^k} - a_i) t^{-mp^{k-\ell}}.$$

Arguing as above, we must have that

$$-eag^2(t) = f^2(t)t^{2c}$$

for some integer c such that $0 \leq -mp^{k-\ell} - 2c \leq 1$ and

$$t^{-mp^{k-\ell}} - a^{-1} = t^d \prod_{i \in Z} (t^{p^k} - a_i),$$

where $d = -mp^{k-\ell} - 2c$. It follows that $d = 0$ and that $-mp^{k-\ell} = p^k$. But then $p^k = 2c$, which is impossible since p is odd. $\square$

Finally Theorem 4.1.7 is an immediate consequence of Lemmas 4.3.4 and 4.3.9.

LEMMA 4.3.9. *If $\pi \in \operatorname{Aut}(K_\Gamma)$, then*

(a) $\pi[X] = X$;

(b) $\pi(t_n) = t_n$ *for all* $n \in \omega$;

(c) $\pi[H_n(\Gamma)] = H_n(\Gamma)$ *for all* $n \in \omega$; *and*

(d) *there exists* $\varphi \in \operatorname{Aut}(\Gamma)$ *such that* $\pi \restriction X = \varphi$.

PROOF. If $v \in X$, then $\pi(v)$ is a p_0-high element of K_Γ and so Lemma 4.3.7(a) yields that $\pi(v) \in K_0(\Gamma)$. Similarly, if $n \in \omega$, then $\pi(t_n)$ is a p_{n+1}-high element of K_Γ, and so $\pi(t_n) \in K_{n+1}(\Gamma)$. Applying Corollary 4.3.3, we obtain inductively that $\pi[K_n(\Gamma)] \subseteq K_n(\Gamma)$ for all $n \in \omega$.

Now fix some integer $n \in \omega$. By Lemma 4.3.7(b), we have that $\pi(t_n) = et_n^r$, where $e^2 = 1$ and $r = m/p_{n+1}^\ell$ for some $m \in \mathbb{Z} \setminus \{0\}$ and $\ell \in \mathbb{N}$. Let $a \in H_n(\Gamma)$. Then there exists an element $c \in K_{n+1}(\Gamma)$ such that $c^2 = t_n - a$. Hence there exists $b \in K_{n+1}(\Gamma)$ such that

$$b^2 = e(t_n^r - \pi(ea)).$$

Since $\pi(ea) \in K_n(\Gamma)$, Lemma 4.3.8 yields that $e = r = 1$ and $\pi(a) = \pi(ea) \in H_n(\Gamma)$. Thus $\pi(t_n) = t_n$ and $\pi[H_n(\Gamma)] \subseteq H_n(\Gamma)$. Similarly, $\pi^{-1}[H_n(\Gamma)] \subseteq H_n(\Gamma)$ and so $\pi[H_n(\Gamma)] = H_n(\Gamma)$.

In particular, if $(u, v) \in E$, then $\pi(u) + \pi(v) = \pi(u + v) \in H_0(\Gamma)$. Thus there exists an edge $(u', v') \in E$ such that $\pi(u) + \pi(v) = u' + v'$. Since $\pi(u)$ is a p_0-high element of K_Γ, either $\pi(u)$ or $-\pi(u)$ is a product of elements of the form x^{m/p_0^ℓ}, where $x \in X$ is a vertex of Γ, $m \in \mathbb{Z} \setminus \{0\}$ and $\ell \in \mathbb{N}$. The same is true of $\pi(v)$. This easily implies that $\{\pi(u), \pi(v)\} = \{u', v'\}$. Since Γ has no isolated vertices, it follows that $\pi(u) \in X$ for all $u \in X$. It is now clear that $\pi \restriction X \in \operatorname{Aut}(\Gamma)$. $\square$

4.4. A technical lemma

In this section, we shall prove Lemma 4.3.6. So let F be a field of characteristic 0 and let t be a transcendental element over F. Let p be an odd prime, and let

$\{t(\ell) \mid \ell \in \mathbb{N}\}$ be a set of elements such that $t(0) = t$ and $t(\ell+1)^p = t(\ell)$. Let

$$L = F(t(0), t(1), \dots, t(\ell), \dots) = F(t)(\{t\}, p).$$

Let $\{T_i \mid i \in I\}$ be a set of mutually prime polynomials in $F[t]$, none of which is divisible by t or has a multiple factor; and for each $i \in I$, let ϑ_i be an element such that $\vartheta_i^2 = T_i$. Let $M = L(\dots, \vartheta_i, \dots)$. Suppose that q is an odd prime and that u is a q-high element of M. Then we must prove that u is one of the "obvious" q-high elements; namely:

(a) if $q = p$, then $u = ct(\ell)^m$ for some $\ell \in \mathbb{N}$, $m \in \mathbb{Z}$ and some p-high element c of F;

(b) if $q \neq p$, then u is a q-high element of F.

At various points during the proof of Lemma 4.3.6, it will be helpful if F contains various primitive roots of unity. Fortunately, it is enough to prove Lemma 4.3.6 in the special case when F is algebraically closed. To see this, suppose that Lemma 4.3.6 is true whenever F is an algebraically closed field. Now let F be an arbitrary field of characteristic 0 and let $\overline{F}$ be the algebraic closure of F. Let $\overline{L} = \overline{F}(t(0), t(1), \dots, t(\ell), \dots)$ and let $\overline{M} = \overline{L}(\dots, \vartheta_i, \dots)$.

First suppose that u is a p-high element of M. Then u is also a p-high element of $\overline{M}$. Hence $u = ct(\ell)^m$ for some $\ell \in \mathbb{N}$, $m \in \mathbb{Z}$ and some p-high element c of $\overline{F}$. Since $t(\ell)^m$ is p-high in M, it follows that $c = u/t(\ell)^m$ is also p-high in M. Consequently, c is a p-high element of $M \cap \overline{F}$. But the proof of Corollary 4.3.3 shows that $M \cap \overline{F} = F$. Hence c is a p-high element of F. A similar argument shows that if u is a q-high element of M for some odd prime $q \neq p$, then u is a q-high element of F.

So from now on, we shall assume that F is an algebraically closed field. In particular, every nonzero element of F is q-high.

Suppose that u is a q-high element of M. Then there exists an integer $\ell \geq 0$ and a finite subset $\{i_1, \dots, i_m\}$ of I such that $u \in F(t(\ell), \vartheta_{i_1}, \dots, \vartheta_{i_m})$. There are essentially two possibilities to consider.

Possibility (I): u is already a q-high element of $F(t(\ell), \vartheta_{i_1}, \dots, \vartheta_{i_m})$.

Possibility (II): As n gets larger, in order to obtain solutions to the equation $y^{q^n} = u$, it is necessary to pass to progressively larger algebraic extensions of $F(t(\ell), \vartheta_{i_1}, \dots, \vartheta_{i_m})$.

First we shall use valuation theory to show that if possibility (I) holds, then $u \in F$. Afterwards we shall use some basic Kummer theory to deal with possibility (II).

Definition 4.4.1. Let K be any field. Then a *valuation* ν of K is a map $\nu : K \to \mathbb{R}$ such that for every a, $b \in K$, the following properties hold:

(a) $\nu(a) \geq 0$, and $\nu(a) = 0$ iff $a = 0$;
(b) $\nu(ab) = \nu(a)\nu(b)$; and
(c) $\nu(a + b) \leq \nu(a) + \nu(b)$.

The subgroup $\nu[K^*]$ of $\mathbb{R}$ is called the *value group* of ν; and the valuation ν is said to be *discrete* if the value group $\nu[K^*]$ is an infinite cyclic group. Two valuations ν_1, ν_2 of the field K are said to be *equivalent* iff there exists a positive $r \in \mathbb{R}$ such that $\nu_2(a) = \nu_1(a)^r$ for all $a \in K$.

Example 4.4.2. Let K_0 be any field and let x be a transcendental element over K_0.Let $K = K_0(x)$ be the corresponding rational function field. Fix some real number r such that $0 < r < 1$. For each irreducible polynomial $f \in K_0[x]$, we can define a discrete valuation ν_f of K by

$$\nu_f(a) = r^{-\ell \deg(f)} \text{ iff } u = (b/c)f^\ell,$$

where $\ell \in \mathbb{Z}$ and b, $c \in K_0[x]$ are polynomials which are relatively prime to f. We can also define a discrete valuation ν_∞ of K by setting

$$\nu_\infty(b/c) = r^{\deg(b) - \deg(c)},$$

for each nonzero quotient of polynomials b, $c \in K_0[x]$. It is easily checked that the set

$$S = \{\nu_\infty\} \cup \{\nu_f \mid f \in K_0[x] \text{ is irreducible }\}$$

of pairwise nonequivalent discrete valuations of K satisfies the *product formula*; i.e. for every $0 \neq a \in K$,

(i) there exist only finitely many $\nu \in S$ such that $\nu(a) \neq 1$; and

(ii) $\prod_{\nu \in S} \nu(a) = 1$.

It is also easily checked that the following property holds.

(iii) $K_0 = \{0\} \cup \{a \in K \mid \nu(a) = 1 \text{ for all } \nu \in S\}$.

In our analysis of possibility (I), we shall make use of the fact that for every finite extension of a rational function field, there exists a corresponding set of discrete valuations satisfying the appropriate analogues of the above properties.

THEOREM 4.4.3. *Let K_0 be any field and let E be a finite extension of the rational function field $K_0(x)$. Then there exists a set $\{\nu_j \mid j \in J\}$ of pairwise nonequivalent discrete valuations of E which satisfies the following properties.*

(i) *If $0 \neq a \in E$, then there exist only finitely many $j \in J$ such that $\nu_j(a) \neq 1$.*

(ii) *If $0 \neq a \in E$, then $\prod_{j \in J} \nu_i(a) = 1$.*

(iii) *$\{0\} \cup \{a \in E \mid \nu_j(a) = 1$ for all $j \in J\}$ is the algebraic closure of K_0 in E.*

PROOF. This is proved in Chapter 12 of Artin [**1**]. $\square$

LEMMA 4.4.4. *Let q be any prime. Let $\ell \geq 0$ be an integer and let $\{i_1, \dots, i_m\}$ be a finite subset of I. If u is a q-high element of $F(t(\ell), \vartheta_{i_1}, \dots, \vartheta_{i_m})$ then $u \in F$.*

PROOF. Let $x = t(\ell)$. Then $E = F(t(\ell), \vartheta_{i_1}, \dots, \vartheta_{i_m})$ is a finite extension of the rational function field $F(x)$. Let $\{\nu_j \mid j \in J\}$ be the set of discrete valuations of E, which is given by Theorem 4.4.3. Consider the abelian group homomorphism

$$\pi : E^* \to \prod_{j \in J} \nu_j[E^*]$$

defined by $\pi(a) = (\nu_j(a))_{j \in J}$. Then $\pi[E^*]$ is a subgroup of the free abelian group $\bigoplus_{j \in J} \nu_j[E^*]$ and so $\pi[E^*]$ is a free abelian group. By Lemma 4.3.2, if $a \in E$ is algebraic over F, then $a \in F$. Thus $\ker \pi = F^*$. Hence if $a \in E \setminus F$, then $\pi(a)$ is a nonidentity element of the free abelian group $\pi[E^*]$ and this implies that a is not a q-high element of E. $\square$

LEMMA 4.4.5. *If u is a q-high element of M for some odd prime $q \neq p$, then $u \in F$.*

PROOF. There exists an integer $\ell \geq 0$ and a finite subset $\{i_1, \ldots, i_m\}$ of I such that $u \in E = F(t(\ell), \vartheta_{i_1}, \ldots, \vartheta_{i_m})$. Suppose that there exists an integer $n \geq 0$ such that the equation $y^{q^{n+1}} = u$ is not solvable in E. Let n be the least such integer and let $v \in E$ be such that $v^{q^n} = u$. Let $w \in M$ satisfy $w^q = v$. Then $[E(w) : E] = q$. But clearly if E' is a finite extension of E such that $E \subseteq E' \subseteq M$, then $[E' : E] = 2^r p^s$ for some r, $s \geq 0$. Thus u must already be q-high in E and so $u \in F$. $\square$

It now only remains for us to analyse the case when u is a p-high element of M such that $u \notin F$. (The above analysis shows that this case correponds exactly to possibility (II).) Once again, there exists an integer $\ell \geq 0$ and a finite subset $\{i_1, \ldots, i_m\}$ of I such that $u \in E = F(t(\ell), \vartheta_{i_1}, \ldots, \vartheta_{i_m})$. Let $E = L(\vartheta_{i_1}, \ldots, \vartheta_{i_m})$. If E' is a finite extension of E such that $E \subseteq E' \subseteq M$, then $[E' : E] = 2^r$ for some $r \geq 0$ and it follows that u is already p-high in E. Thus in order to complete the proof of Lemma 4.3.6, it is enough to prove the following two lemmas.

LEMMA 4.4.6. *If u is a p-high element of $L(\vartheta_{i_1}, \ldots, \vartheta_{i_m})$ for some finite subset $\{i_1, \ldots, i_m\}$ of I, then $u \in L$.*

(Of course, the argument of the previous paragraph implies that if $u \in L$ is a p-high element of M, then u is already p-high in L.)

LEMMA 4.4.7. *If u is a p-high element of L, then $u = ct(\ell)^m$ for some $\ell \in \mathbb{N}$, $m \in \mathbb{Z}$ and $c \in F$.*

In order to prove Lemma 4.4.7, we must check that in passing from $F(t(\ell))$ to $F(t(\ell + 1))$, only the "obvious" elements of $F(t(\ell))$ obtain pth roots. This is one of the issues which is addressed by Kummer theory. More specifically, let n be a positive integer and let K be a field of characteristic 0 containing a primitive nth root of unity. Then an extension L of K is said to be a *Kummer extension* of exponent n iff

$$L = K(\sqrt[n]{\Delta})$$

for some subgroup Δ of K^* containing the group $(K^*)^n$ of nth powers of K^*; i.e. L is the field generated by the roots $\sqrt[n]{a}$ for $a \in \Delta$. The most basic result of Kummer theory is that Δ is uniquely determined via the formula

$$\Delta = (L^*)^n \cap K.$$

(For example, see Section VIII.8 of Lang [**27**].) In other words, in passing from K to the Kummer extension $K(\sqrt[n]{\Delta})$, only the "obvious" elements of K obtain nth roots.

THEOREM 4.4.8. *Let n be a positive integer and let K be a field of characteristic 0 containing a primitive nth root of unity. Let $0 \neq a \in K$ and suppose that z is a root of the equation $y^n = a$. If $0 \neq b \in K(z)$ satisfies $b^n \in K$, then $b = cz^k$ for some $k \in \mathbb{Z}$ and $c \in K$.*

PROOF. Let $L = K(z)$. Then L is the Kummer extension $K(\sqrt[n]{\Delta})$, where Δ is the subgroup of K^* generated by $(K^*)^n \cup \{a\}$. Hence

$$b^n \in (L^*)^n \cap K = \Delta;$$

say, $b^n = c^n a^k$ for some $k \in \mathbb{Z}$ and $c \in K^*$. After multiplying c by a suitable nth root of unity if necessary, we obtain that $b = cz^k$. □

PROOF OF LEMMA 4.4.7. Let $\ell \geq 0$ be an integer such that $u \in F(t(\ell))$. Let $x = t(\ell)$ and let $u = x^m f_1(x)/g_1(x)$, where $m \in \mathbb{Z}$ and $f_1(x)$, $g_1(x) \in F[x]$ are relatively prime polynomials, neither of which is divisible by x. Since x is p-high in L, it follows that $c = u/x^m$ is also p-high in L. Suppose that there exists an integer $n \geq 0$ such that the equation $y^{p^{n+1}} = c$ is not solvable in $F(x)$. Let n be the least such integer and let $b \in F(x)$ be such that $b^{p^n} = c$. Then there exist relatively prime polynomials $f_2(x)$, $g_2(x) \in F[x]$, neither of which is divisible by x, such that $b = f_2(x)/g_2(x)$. After increasing ℓ if necessary, we can suppose that there exists an element $a \in F(t(\ell + 1))$ such that $a^p = b$. By Theorem 4.4.8, $a = t(\ell + 1)^k d$ for some $k \in \mathbb{Z}$ and $d \in F(x)$. Note that since $a \notin F(x)$, it follows that k is not divisible by p. Let $d = x^r f_3(x)/g_3(x)$, where $r \in \mathbb{Z}$ and $f_3(x)$, $g_3(x) \in F[x]$ are relatively prime polynomials, neither of which is divisible by x. Since

$$x^{k+rp} f_3^p(x)/g_3^p(x) = a^p = b = f_2(x)/g_2(x),$$

it follows that $k + rp = 0$. But this contradicts the fact that k is not divisible by p. Hence c is already p-high in $F(x)$ and so by Lemma 4.4.4, $c \in F$. Thus $u = cx^m = ct(\ell)^m$ has the required form. $\square$

Finally we shall prove Lemma 4.4.6. We shall argue by induction on $m \geq 1$ that the following statement holds.

Claim 4.4.9. *Let F be an algebraically closed field of characteristic 0 and let t be a transcendental element over F. Let p be an odd prime, and let $\{t(\ell) \mid \ell \in \mathbb{N}\}$ be a set of elements such that $t(0) = t$ and $t(\ell+1)^p = t(\ell)$. Let*

$$L = F(t(0), t(1), \ldots, t(\ell), \ldots) = F(t)(\{t\}, p).$$

Let $\{T_1, \ldots, T_m\}$ be a set of mutually prime polynomials in $F[t]$, none of which is divisible by t or has a multiple factor; and for each $1 \leq i \leq m$, let ϑ_i be an element such that $\vartheta_i^2 = T_i$. If u is a p-high element of $L(\vartheta_1, \ldots, \vartheta_m)$, then $u \in L$.

We shall begin by giving the relatively straightforward induction step of the argument. So suppose that $m > 1$ and that u is p-high in $L(\vartheta_1, \ldots, \vartheta_m)$. Let $N = L(\vartheta_3, \ldots, \vartheta_m)$. Then Lemma 4.3.2 implies that we can express u uniquely as

$$u = a + b_1\vartheta_1 + b_2\vartheta_2 + c\vartheta_1\vartheta_2$$

where a, b_1, b_2, $c \in N$. Furthermore, there exist elements π_1, π_2 of the Galois group $\mathrm{Gal}(N(\vartheta_1, \vartheta_2)/N)$ such that

- $\pi_1(\vartheta_1) = -\vartheta_1$ and $\pi_1(\vartheta_2) = \vartheta_2$;
- $\pi_2(\vartheta_1) = \vartheta_1$ and $\pi_2(\vartheta_2) = -\vartheta_2$.

Applying π_1 and π_2 to u, we see that the elements

$$u_1 = a - b_1\vartheta_1 + b_2\vartheta_2 - c\vartheta_1\vartheta_2$$

and

$$u_2 = a + b_1\vartheta_1 - b_2\vartheta_2 - c\vartheta_1\vartheta_2$$

are also p-high in $N(\vartheta_1, \vartheta_2)$; and so uu_1 and uu_2 are also p-high in $N(\vartheta_1, \vartheta_2)$. Note that $uu_1 = u\pi_1(u) \in N(\vartheta_2)$ and that $uu_2 = u\pi_2(u) \in N(\vartheta_1)$. It follows that uu_1 and uu_2 are already p-high in $N(\vartheta_2)$ and $N(\vartheta_1)$ respectively. By induction hypothesis, we obtain that uu_1, $uu_2 \in L$. Hence there exists an element $d \in L$ such that $u_1 = du_2$; i.e.

$$a - b_1\vartheta_1 + b_2\vartheta_2 - c\vartheta_1\vartheta_2 = da + db_1\vartheta_1 - db_2\vartheta_2 - dc\vartheta_1\vartheta_2.$$

It follows that either

- $d = 1$ and $b_1 = b_2 = 0$; or
- $d = -1$ and $a = c = 0$.

Consequently, u can be expressed in one of the two following forms:

$$u = a + c\vartheta_1\vartheta_2 \quad \text{or} \quad u = b_1\vartheta_1 + b_2\vartheta_2.$$

In the first case, u is a p-high element of $N(\vartheta_1\vartheta_2) = L(\vartheta_1\vartheta_2, \vartheta_3, \dots, \vartheta_m)$. Since the polynomials T_1T_2, $T_3, \dots, T_m$ also satisfy the hypotheses of Claim 4.4.9, the induction hypothesis yields that $u \in L$. In the second case,

$$u^2 = (b_1^2T_1 + b_2^2T_2) + 2b_1b_2\vartheta_1\vartheta_2$$

is a p-high element of $N(\vartheta_1\vartheta_2)$ and so $u^2 \in L$. Thus $b_1b_2 = 0$ and so either $u = b_2\vartheta_2 \in N(\vartheta_2)$ or $u = b_1\vartheta_1 \in N(\vartheta_1)$. Once again, the induction hypothesis yields that $u \in L$.

Finally we shall deal with the more difficult basis step of the argument. In order to simplify the notation, let $T = T_1$ and $\vartheta = \vartheta_1$. Suppose that $u = a + b\vartheta \in L(\vartheta)$, where a, $b \in L$. Then $\overline{u} = a - b\vartheta$ denotes the image of u under the nontrivial element of the Galois group $\mathrm{Gal}(L(\vartheta)/L)$.

Claim 4.4.10. *If $b \in L^*$, then the element $v = b\vartheta$ is not p-high in $L(\vartheta)$.*

Proof. Suppose that there exists an element $b \in L^*$ such that $v = b\vartheta$ is p-high in $L(\vartheta)$. Let $\ell \geq 0$ be an integer such that $b \in F(t(\ell))$. Let $x = t(\ell)$ and let $f(x)$, $g(x) \in F[x]$ be relatively prime polynomials such that $b = f(x)/g(x)$.

Since -1 is p-high in F, it follows that $v^2 = -v\overline{v}$ is p-high in L. Applying Lemma 4.4.7, we find that either

- $f^2(x)T = cg^2(x)x^n$; or
- $f^2(x)Tx^n = cg^2(x)$

for some $n \geq 0$ and $c \in F$. In the first case, we obtain that $g^2(x)$ divides T; and since T has no multiple factors, this implies that $g(x) \in F$. But this means that

T divides x^n, which contradicts the hypothesis that T is not divisible by t. In the second case, we must have that $f(x) \in F$. After cancelling a suitable power of x, we obtain that either $Tx = dg_1^2(x)$ or $T = dg_1^2(x)$ for some $d \in F$ and $g_1(x) \in F[x]$. In the former case, we see that x divides T, which implies that t divides T. In the latter case, T has a multiple factor. So both cases lead to a contradiction. $\square$

Now suppose that u is any p-high element of $L(\vartheta)$. Then we can express u uniquely in the form $u = a + b\vartheta$, where a, $b \in L$ and $a \neq 0$. For the sake of contradiction, suppose that $b \neq 0$. Clearly

$$\lambda = u\overline{u} = a^2 - b^2 T$$

is a p-high element of L. Hence $\lambda = ct(\ell)^m$ for some $\ell \in \mathbb{N}$, $m \in \mathbb{Z}$ and $c \in F$. Let $m = 2n + k$, where $k = 0, 1$. If $k = 0$, then $\sqrt{\lambda} \in L$. So after dividing a and b by $\sqrt{\lambda}$ if necessary, we can suppose that $u\overline{u} = 1$. Similarly if $k = 1$, then we can suppose that $u\overline{u} = t(\ell)$. Since u is p-high, it follows that

$$v = -u\overline{u}t(\ell)^{-1} = u^2 t(\ell)^{-1} = (a^2 + b^2 T)t(\ell)^{-1} + (2abt(\ell)^{-1})\vartheta$$

is also p-high in $L(\vartheta)$. Furthermore,

$$v\overline{v} = (u\overline{u})^2 t(\ell)^{-2} = 1.$$

If we can show that $2abt(\ell)^{-1} = 0$, then it follows that $b = 0$. Hence it is enough to consider the case when $u\overline{u} = 1$. Let a, $b \in F(t(\ell))$ and let $x = t(\ell)$. Then there exist polynomials f, g, $h \in F[x]$ such that

$$u = a + b\vartheta = (f/h) + (g/h)\vartheta.$$

Since $u\overline{u} = 1$, we have that $f^2 - g^2 T = h^2$. Since T has no multiple factors, we can suppose that f, g and h are pairwise relatively prime; and we can also suppose that h has been chosen so that its leading coefficient is 1. We shall eventually show that $f + g\vartheta$ is already p-high in $F(x, \vartheta)$. But then Lemma 4.4.4 implies that $g = 0$, which contradicts the assumption that $b = g/h \neq 0$.

Fix some integer $n \geq 1$ and let $m = p^n$. Let $v \in L(\vartheta)$ satisfy $v^m = u$. Let $i \geq 0$ be such that $v \in F(t(\ell + i))$ and let $z = t(\ell + i)$. Then there exist polynomials c, d, $e \in F[z]$ such that $v = (c/e) + (d/e)\vartheta$. Since $(v\overline{v})^m = u\overline{u} = 1$ and F is algebraically closed, we can suppose that $v\overline{v} = 1$. This gives that $c^2 - d^2 T = e^2$; and so we can suppose that c, d and e are pairwise relatively prime and that the leading coefficient of e is 1. Since

$$[(c/e) + (d/e)\vartheta]^m = (f/h) + (g/h)\vartheta,$$

it follows that

$$fe^m = ch\left[c^{m-1} + \binom{m}{2}c^{m-3}d^2T + \cdots + \binom{m}{m-1}d^{m-1}T^{(m-1)/2}\right].$$

Using the facts that $d^2T = c^2 - e^2$ and that

$$1 + \binom{m}{2} + \binom{m}{4} + \cdots + \binom{m}{m-1} = 2^{m-1},$$

we obtain the equality

$$fe^m = ch[2^{m-1}c^{m-1} + e^2\Phi(c, e)],$$

where Φ is a polynomial in two variables over $\mathbb{Q}$. Let $0 \leq j \leq m$ be the greatest integer such that e^j divides h and let $h_1 = h/e^j$. Then

$$fe^{m-j} = ch_1[2^{m-1}c^{m-1} + e^2\Phi(c, e)].$$

Since c and e are relatively prime, it follows that e does not divide the right-hand side of this last equation. Hence $m = j$ and so h_1 divides f. Since f, h are relatively prime and the leading coefficients of h and e are 1, it follows that $h_1 = 1$. Thus $h = e^m$ and $(c + d\vartheta)^m = f + g\vartheta$.

At this point, we only know that c, $d \in F[z]$, where $z = t(\ell + i)$ for some $i \geq 0$. Now we shall use Kummer theory to show that c, $d \in F[x]$. Clearly we can suppose that i was chosen so that $p^i \geq m = p^n$. Thus

$$(c + d\vartheta)^{p^i} = (f + g\vartheta)^{p^i/m} \in F(x, \vartheta).$$

By Theorem 4.4.8, there exist rational functions φ, $\psi \in F(x)$ and an integer $k \in \mathbb{Z}$ such that

$$c(z) + d(z)\vartheta = [\varphi(x) + \psi(x)\vartheta]z^k = [\varphi(z^{p^i}) + \psi(z^{p^i})\vartheta]z^k.$$

Since c, $d \in F[z]$, we can choose φ, ψ, k such that φ, $\psi \in F[x]$ and $k \geq 0$. Consider the equality

$$f(x) + g(x)\vartheta = [\varphi(x) + \psi(x)\vartheta]^m z^{km}.$$

As f and g are relatively prime polynomials, we must have that $k = 0$; and so $c = \varphi$ and $d = \psi$. Since $m = p^n$ was an arbitrary power of p, we have now shown that $f + g\vartheta$ is already p-high in $F(x, \vartheta)$, which is the desired contradiction.

4.5. Notes

The normaliser tower technique was first introduced in Thomas [**49**], where it was used to prove Theorem 4.1.13. The proof of Theorem 4.1.8 is closely based upon Hodges [**16**, Section 5.5]. The proof of Theorem 4.1.7 is an expanded version of the original proof of Fried-Kollar [**10**]. (There is no mention of valuations or Kummer theory in the purely computational proof of Fried-Kollar. While the introduction of these notions neither shortens nor simplifies the proof, I think that it helps to explain why Lemma 4.3.6 is true.)

CHAPTER 5

Hamkins' Theorem

Recently Joel Hamkins has observed that it is possible to define the automorphism tower of an arbitrary (not necessarily centreless) group. The only complication is that, instead of defining the automorphism tower to be an ascending chain of groups, it is now necessary to define the automorphism tower to be a suitable direct system of groups and homomorphisms. In this chapter, we shall prove Hamkins' theorem which says that if G is an arbitrary group, then the automorphism tower of G eventually terminates. Perhaps the most interesting feature of Hamkins' proof is that it gives virtually no information concerning the height $\tau(G)$ of the automorphism tower of an arbitrary group G, apart from the fact that $\tau(G)$ is less than the least inaccessible cardinal κ such that $\kappa > |G|$. Even in the case when G is a finite group, no better upper bound for $\tau(G)$ is currently known than the least inaccessible cardinal.

5.1. Automorphism towers of arbitrary groups

In this section, we shall define the automorphism tower of an arbitrary (not necessarily centreless) group; and we shall prove that the automorphism tower of an arbitrary group eventually terminates. But before we can define the automorphism tower of an arbitrary group, we first need to introduce the notion of the direct limit of a direct system of groups and homomorphisms.

Definition 5.1.1. A partial ordering $\langle \Lambda, \leq \rangle$ is said to be a *directed set* if for all σ, $\nu \in \Lambda$, there exists an element $\lambda \in \Lambda$ such that $\sigma \leq \lambda$ and $\nu \leq \lambda$.

In this book, we shall only need to consider the case when Λ is an ordinal, equipped with its usual well-ordering.

Definition 5.1.2. Let Λ be a directed set. Then a *direct system* of groups and homomorphisms indexed by Λ consists of

(a) a set $\{G_\lambda \mid \lambda \in \Lambda\}$ of groups and

(b) a set $\{\varphi_{\lambda,\mu} : G_\lambda \to G_\mu \mid \lambda \leq \mu \in \Lambda\}$ of homomorphisms

such that the following conditions hold.

(i) $\varphi_{\lambda,\lambda} : G_\lambda \to G_\lambda$ is the identity map for all $\lambda \in \Lambda$.

(ii) $\varphi_{\mu,\nu} \circ \varphi_{\lambda,\mu} = \varphi_{\lambda,\nu}$ for all $\lambda \leq \mu \leq \nu$.

Definition 5.1.3. Suppose that $\{G_\lambda \mid \lambda \in \Lambda\}$ and $\{\varphi_{\lambda,\mu} \mid \lambda \leq \mu \in \Lambda\}$ is a direct system of groups and homomorphisms indexed by Λ. Let G be a group and let $\varphi_\lambda : G_\lambda \to G$ be a homomorphism for each $\lambda \in \Lambda$. Then $(G, \{\varphi_\lambda \mid \lambda \in \Lambda\})$ is the *direct limit* of the direct system if the following conditions hold.

(a) $\varphi_\lambda = \varphi_\mu \circ \varphi_{\lambda,\mu}$ for all $\lambda \leq \mu$.

(b) Suppose that H is a group and that for each $\lambda \in \Lambda$, $\psi_\lambda : G_\lambda \to H$ is a homomorphism such that $\psi_\lambda = \psi_\mu \circ \varphi_{\lambda,\mu}$ for all $\lambda \leq \mu \in \Lambda$. Then there exists a unique homomorphism $\psi : G \to H$ such that $\psi_\lambda = \psi \circ \varphi_\lambda$ for all $\lambda \in \Lambda$.

In this case, we write $G = \varinjlim G_\lambda$.

It is clear that each direct system has at most one direct limit up to isomorphism.

Theorem 5.1.4. *Any direct system of groups and homomorphisms has a direct limit.*

Proof. Suppose that $\{G_\lambda \mid \lambda \in \Lambda\}$ and $\{\varphi_{\lambda,\mu} \mid \lambda \leq \mu \in \Lambda\}$ is a direct system of groups and homomorphisms indexed by Λ. To simplify notation, assume that the groups $\{G_\lambda \mid \lambda \in \Lambda\}$ are pairwise disjoint. Define an equivalence relation E on $X = \bigsqcup_{\lambda\in\Lambda} G_\lambda$ as follows.

- Suppose that $g \in G_\sigma$ and $h \in G_\nu$. Then $g \mathrel{E} h$ iff there exists $\lambda \geq \sigma$, ν such that $\varphi_{\sigma,\lambda}(g) = \varphi_{\nu,\lambda}(h)$.

For each $g \in X$, let $[g] = \{h \in X \mid g \mathrel{E} h\}$ be the corresponding E-equivalence class. Then we can define a group operation on the set $G = \{[g] \mid g \in X\}$ of E-equivalence classes as follows.

- Suppose that $[g]$, $[h] \in G$, where $g \in G_\sigma$ and $h \in G_\nu$. Then

$$[g] \cdot [h] = [\varphi_{\sigma,\lambda}(g) \cdot \varphi_{\nu,\lambda}(h)],$$

where $\lambda \in \Lambda$ is any element such that $\lambda \geq \sigma$, ν.

Finally for each $\lambda \in \Lambda$, define the homomorphism $\varphi_\lambda : G_\lambda \to G$ by $\varphi_\lambda(g) = [g]$. Then it is easily checked that $G = \varinjlim G_\lambda$. □

We are now ready to define the automorphism tower of an arbitrary (not necessarily centreless) group.

DEFINITION 5.1.5. Let G be an arbitrary group. Then the automorphism tower of G consists of the class of groups $\{G_\alpha \mid \alpha \in On\}$, together with the class of canonical homomorphisms $\{\pi_{\beta,\alpha} : G_\beta \to G_\alpha \mid \beta \leq \alpha \in On\}$, defined inductively as follows.

(a) $G_0 = G$ and $\pi_{0,0} = id_{G_0}$.
(b) If $\alpha = \beta + 1$, then $G_\alpha = \mathrm{Aut}(G_\beta)$ and $\pi_{\beta,\alpha} : G_\beta \to G_\alpha$ is the canonical homomorphism which sends each $g \in G_\beta$ to the corresponding inner automorphism $i_g \in \mathrm{Aut}(G_\beta) = G_\alpha$. Furthermore, $\pi_{\alpha,\alpha} = id_{G_\alpha}$ and if $\gamma < \beta$, then $\pi_{\gamma,\alpha} = \pi_{\beta,\alpha} \circ \pi_{\gamma,\beta}$.
(c) If α is a limit ordinal, then we define $(G_\alpha, \{\pi_{\beta,\alpha} \mid \beta < \alpha\})$ to be the direct limit of the direct system of groups $\{G_\beta \mid \beta < \alpha\}$ and canonical homomorphisms $\{\pi_{\gamma,\beta} \mid \gamma \leq \beta < \alpha\}$; and we set $\pi_{\alpha,\alpha} = id_{G_\alpha}$.

The automorphism tower of G is said to terminate if there exists an ordinal α such that the canonical homomorphism $\pi_{\alpha,\alpha+1} : G_\alpha \to G_{\alpha+1}$ is an isomorphism. Once again, this occurs if and only if there exists an ordinal α such that G_α is a complete group. In order for the automorphism tower of G to terminate, it is actually enough that there should exist an ordinal γ such that G_γ is centreless; since then Theorems 2.1.1 and 3.1.13 imply that there exists an ordinal α such that G_α is complete.

THEOREM 5.1.6 (Hamkins [**14**]). *If G is an arbitrary group, then the automorphism tower of G eventually terminates.*

PROOF. Let G be an arbitrary group and let $\langle G_\alpha \mid \alpha \in On\rangle$ be the automorphism tower of G. For each $\alpha \leq \beta$, let $\pi_{\alpha,\beta} : G_\alpha \to G_\beta$ be the corresponding canonical homomorphism. As we explained above, it suffices to show that there exists an ordinal γ such that G_γ is a centreless group. For each ordinal α, let

$$H_\alpha = \{g \in G_\alpha \mid \text{ There exists } \beta \geq \alpha \text{ such that } \pi_{\alpha,\beta}(g) = 1\}.$$

For each $g \in H_\alpha$, let β_g be the least ordinal such that $\pi_{\alpha,\beta_g}(g) = 1$ and let

$$f(\alpha) = \sup\{\beta_g \mid g \in H_\alpha\}.$$

It is easily checked that if $\alpha < \beta$, then $\alpha \leq f(\alpha) \leq f(\beta)$. Now define a strictly increasing ω-sequence of ordinals by

(a) $\gamma_0 = 0$,
(b) $\gamma_{n+1} = f(\gamma_n + 1)$,

and let $\gamma = \sup\{\gamma_n \mid n \in \omega\}$. Then γ is a limit ordinal such that $f(\alpha) < \gamma$ for every $\alpha < \gamma$. We shall prove that G_γ is a centreless group.

Suppose that $g \in Z(G_\gamma)$, so that $\pi_{\gamma,\gamma+1}(g) = 1$. Since γ is a limit ordinal, there exists $\alpha < \gamma$ such that $g = \pi_{\alpha,\gamma}(h)$ for some $h \in G_\alpha$. Notice that

$$\pi_{\alpha,\gamma+1}(h) = \pi_{\gamma,\gamma+1}(\pi_{\alpha,\gamma}(h)) = \pi_{\gamma,\gamma+1}(g) = 1.$$

Hence $\pi_{\alpha,f(\alpha)}(h) = 1$. But this means that

$$g = \pi_{\alpha,\gamma}(h) = \pi_{f(\alpha),\gamma}\left(\pi_{\alpha,f(\alpha)}(h)\right) = \pi_{f(\alpha),\gamma}(1) = 1,$$

as desired. □

Definition 5.1.7. Let G be an arbitrary group.

(a) $\tau(G)$ is the height of the automorphism tower of G; i.e. the least ordinal α such that G_α is a complete group.
(b) $\tau_c(G)$ is the least ordinal γ such that G_γ is a centreless group.

By Theorem 2.1.1 and Corollary 3.3.2, if $\gamma = \tau_c(G)$, then

- $\tau(G) < \gamma + \omega$ if G_γ is finite, and
- $\tau(G) < \left(2^{|G_\gamma|}\right)^+$ if G_γ is infinite.

Unfortunately theproof of Theorem 5.1.6 gives absolutely no upper bound for $\tau_c(G)$. However, with a little more effort, it is possible to obtain an upper bound for $\tau_c(G)$, albeit a seemingly terrible one.

THEOREM 5.1.8 (Hamkins [**14**]). *Let G be an arbitrary group. If κ is a regular uncountable cardinal such that $|G_\alpha| < \kappa$ for all $\alpha < \kappa$, then*

(a) $\tau_c(G) < \kappa$*; and*

(b) $\tau(G) \leq \kappa$.

PROOF. (a) This follows from a slight variation of the proof of Theorem 5.1.6. For each ordinal $\alpha < \kappa$, let

$$K_\alpha = \{g \in G_\alpha \mid \pi_{\alpha,\kappa}(g) = 1\};$$

and for each $g \in K_\alpha$, let β_g be the least ordinal β such that $\pi_{\alpha,\beta}(g) = 1$. Clearly $\beta_g < \kappa$ each $g \in K_\alpha$. Let

$$f(\alpha) = \sup\{\beta_g \mid g \in K_\alpha\}.$$

Since κ is regular and $|K_\alpha| < \kappa$, it follows that $f(\alpha) < \kappa$. Arguing as in the proof of Theorem 5.1.6, we can now find a limit ordinal $\gamma < \kappa$ such that $f(\alpha) < \gamma$ for all $\alpha < \gamma$; and then the proof of Theorem 5.1.6 shows that G_γ is a centreless group.
(b) Let $\gamma = \tau_c(G) < \kappa$. Then the groups from step γ to step κ

$$G_\gamma \to \cdots \to G_\alpha \to \cdots \to G_\kappa$$

in the automorphism tower of G can be identified with the corresponding groups in the classical automorphism tower of the centreless group G_γ; and we can apply the argument in the proof of Theorem 3.1.13 to the increasing chain

$$G_\gamma \trianglelefteq G_{\gamma+1} \trianglelefteq \cdots \trianglelefteq G_\alpha \trianglelefteq \cdots G_\kappa = \bigcup_{\gamma \leq \alpha < \kappa} G_\alpha.$$

Specifically, suppose that $\pi \in \mathrm{Aut}(G_\kappa)$. By Theorem 3.1.10, there exists a club C of κ such that $C \subseteq \kappa \smallsetminus \gamma$ and $\pi[G_\alpha] = G_\alpha$ for all $\alpha \in C$. For each $\alpha \in C$, there exists an element $g_\alpha \in G_{\alpha+1} = \mathrm{Aut}(G_\alpha)$ such that $g_\alpha x g_\alpha^{-1} = \pi(x)$ for all $x \in G_\alpha$. If α, $\beta \in C$ and $\alpha < \beta$, then $g_\alpha^{-1} g_\beta \in C_{G_\kappa}(G_\alpha) \leqslant C_{G_\kappa}(G) = 1$ and so $g_\alpha = g_\beta$. Hence there is a fixed element $g \in G_\kappa$ such that $\pi \restriction G_\alpha = i_g \restriction G_\alpha$ for each $\alpha \in C$ and so $\pi = i_g \in \mathrm{Inn}(G_\kappa)$. Thus $\mathrm{Aut}(G_\kappa) = \mathrm{Inn}(G_\kappa)$ and so $\tau(G) \leq \kappa$. □

Of course, in order to apply Theorem 5.1.8, we must first find a regular uncountable cardinal κ such that $|G_\alpha| < \kappa$ for all $\alpha < \kappa$. In the special case when

G is already centreless, Theorem 3.3.1 implies that we can take $\kappa = \left(2^{|G|}\right)^+$. Unfortunately this is not true in general. For example, in the next section, we shall present an example of a countable abelian group G such that $|\operatorname{Aut}(G)| = 2^\omega$ and $|\operatorname{Aut}(\operatorname{Aut}(G))| = 2^{2^\omega}$. At present, we can find no smaller value for κ than the least inaccessible cardinal greater than $|G|$. (In this case, we can obtain a slightly better bound for $\tau(G)$ than that given by Theorem 5.1.8(b).)

DEFINITION 5.1.9. A regular uncountable cardinal κ is *inaccessible* if $2^\theta < \kappa$ for all cardinals $\theta < \kappa$.

COROLLARY 5.1.10. *Let G be an arbitrary group. If κ is an inaccessible cardinal greater than $|G|$, then $\tau(G) < \kappa$.*

PROOF. An easy induction shows that $|G_\alpha| < \kappa$ for all $\alpha < \kappa$. Hence Theorem 5.1.8 yields that $\tau_c(G) < \kappa$. Let $\gamma = \tau_c(G)$. Then $|G_\gamma| < \kappa$ and it follows that $\tau(G) < \max\{\gamma + \omega, \left(2^{|G_\gamma|}\right)^+\} < \kappa$. □

It has to be admitted that the bound given by Corollary 5.1.10 is very unsatisfactory, particularly when it is remembered that the existence of inaccessible cardinals cannot be proved in ZFC. (For example, see Section IV.10 of Kunen [**26**].) It is especially embarrassing that even when G is a finite group, no better upper bound for $\tau(G)$ is currently known than the first inaccessible cardinal.

5.2. Two examples and many questions

The most interesting open problem on automorphism towers is that of finding a satisfactory upper bound for $\tau(G)$ when G is an arbitrary group. As we saw in the last section, this problem is essentially the same as that of finding a satisfactory upper for $\tau_c(G)$. These problems remain open even in the case when G is finite. We shall begin this section by presenting a natural example of a finite group G such that the automorphism tower of G does not terminate after finitely many steps. Thus Wielandt's theorem does not extend to arbitrary finite groups.

THEOREM 5.2.1. *The automorphism tower of the dihedral group D_8 terminates after exactly $\omega + 1$ steps.*

PROOF. Let $\langle G_\alpha \mid \alpha \in On\rangle$ be the automorphism tower of the dihedral group

$$G = D_8 = \langle a, b \mid a^2 = b^2 = (ab)^4 = 1\rangle;$$

and for each $\alpha \leq \beta$, let $\pi_{\alpha,\beta} : G_\alpha \to G_\beta$ be the corresponding canonical homomorphism. Clearly D_8 has an outer automorphism φ of order 2 which interchanges the involutions a and b; and it is a straightforward exercise to show that $\mathrm{Aut}(D_8) = \langle i_a, \varphi\rangle$. Furthermore, an easy calculation shows that $i_a\varphi$ has order 4 and so $\mathrm{Aut}(D_8) \simeq D_8$. Notice that $Z(D_8) = \langle (ab)^2\rangle$ has order 2. Thus the canonical homomorphism $D_8 \to \mathrm{Aut}(D_8)$ is not an isomorphism; and it follows that the automorphism tower of D_8 does not terminate after finitely many steps.

Another easy calculation shows that $(i_a\varphi)^2 = i_{ab}$. Hence $\pi_{0,1}(ab) = i_{ab}$ is the central involution of $\mathrm{Aut}(D_8)$. It follows that $\pi_{0,2}(ab) = 1$ and so $\pi_{0,2}(a) = \pi_{0,2}(b)$. Thus $G_\omega = \langle \pi_{0,\omega}(a)\rangle$ is cyclic of order 2. It follows that the automorphism tower of $G = D_8$ terminates after exactly $\omega + 1$ steps with the group $G_{\omega+1} = 1$. □

No examples are known of finite groups G such that $\tau(G) > \omega + 1$. However, this is probably because of the serious difficulty of computing the automorphism towers of finite groups. It seems that if G is a finite group, then after a small number of steps either

(a) the sequence $\langle G_n \mid n \in \omega\rangle$ becomes constant up to isomorphism; or
(b) the groups G_n become too complicated to compute.

Of course, the problem is even worse if G is an infinite group.

CONJECTURE 5.2.2. The automorphism tower of an arbitrary finite group terminates after countably many steps.

Theorem 5.1.8 implies that if G is a counterexample to Conjecture 5.2.2, then there exists an ordinal $\alpha < \omega_1$ such that $|G_\alpha| \geq \omega$.

QUESTION 5.2.3. Does there exist a finite group G such that $|G_\alpha| \geq \omega$ for some $\alpha < \omega_1$?

Initially it seems reasonable to concentrate on the special case of whether there exists a finite group G such that G_ω is infinite. (It is not obvious to me whether this special case is equivalent to Question 5.2.3.) If G is such a group, then clearly $|G_n| \to \infty$ as $n \to \infty$. Thus we are led to consider the question of whether

$$\langle |G_n| \mid n \in \omega \rangle$$

is a bounded sequence of integers for every finite group G. This question has already been raised by Scott [**42**] in the following equivalent formulation.

Question 5.2.4. Let G be an arbitrary finite group. Is the sequence of groups $\langle G_n \mid n \in \omega \rangle$ eventually periodic up to isomorphism?

With this wording, Question 5.2.4 almost seems to imply a prior knowledge of the existence of finite groups G such that

(i) $G \not\simeq \mathrm{Aut}(G)$ and
(ii) $G \simeq G_n$ for some $n > 1$.

However, despite raising this question with a number of leading group theorists, I have yet to find anyone who knows of an example of such a group.

Question 5.2.5. Does there exist a finite group G such that the sequence of groups $\langle G_n \mid n \in \omega \rangle$ is eventually periodic up to isomorphism, but *not* eventually constant?

Virtually nothing is known concerning the problem of finding upper bounds for the heights of the automorphism towers of arbitrary infinite groups. As we have already seen, this problem is closely related to the problems of bounding $\tau_c(G)$ and $|G_\alpha|$ in terms of $|G|$. The current state of ignorance concerning the first of these problems is particularly embarrassing. While the only known upper bound for $\tau_c(G)$ is the next inaccesssible cardinal after $|G|$, no example is known of a group such that $\tau_c(G) > \omega + 1$.

Conjecture 5.2.6. For each ordinal α, there exists a group G such that $\tau_c(G) = \alpha$.

Conjecture 5.2.6 is most naturally approached via the notion of the (transfinite) upper central series of a group.

DEFINITION 5.2.7. The *upper central series* of a group G is defined to be the series

$$1 = Z_0(G) \leqslant Z_1(G) \leqslant Z_2(G) \leqslant \cdots \leqslant Z_\alpha(G) \leqslant \cdots$$

such that

(a) $Z_{\alpha+1}(G)/Z_\alpha(G)$ is the centre of $G/Z_\alpha(G)$; and
(b) if α is a limit ordinal, then $Z_\alpha(G) = \bigcup_{\beta<\alpha} Z_\beta(G)$.

If α is the least ordinal such that $Z_\alpha(G) = Z_{\alpha+1}(G)$, then we say that the upper central series of G has length α and that $\zeta(G) = Z_\alpha(G)$ is the *hypercentre* of G. If $\zeta(G) = G$, then we say that G is a *hypercentral* group.

It seems reasonable to expect that for each α, there exists a group G whose upper central series has length exactly exactly α and which satisfies $\tau_c(G) = \alpha$. For example, it might be worth considering the hypercentral groups constructed by McLain [**31**]. Of course, even if this approach to Conjecture 5.2.6 is successful, each of the resulting groups G will satisfy $|G| \geq |\alpha|$ and so $\tau_c(G) < |G|^+$. However, the real question is whether $\tau_c(G)$ can be substantially greater than $|G|$.

QUESTION 5.2.8. Does there exist an infinite group G such that $\tau_c(G) \geq |G|^+$?

Very little is also known concerning the problem of bounding $|G_\alpha|$ in terms of $|G|$. We have seen that if G is centreless, then $|G_\alpha| < \left(2^{|G|}\right)^+$ for all ordinals α. However, this result does not extend to arbitrary groups.

THEOREM 5.2.9. *There exists a countable group G such that $|\operatorname{Aut}(G)| = 2^\omega$ and $|\operatorname{Aut}(\operatorname{Aut}(G))| = 2^{2^\omega}$.*

PROOF. For each prime p, let $\mathbb{Z}[1/p] = \{m/p^n \mid m \in \mathbb{Z}, n \in \mathbb{N}\}$ be the additive group of p-adic rationals. Let $G = \bigoplus_p \mathbb{Z}[1/p]$. Clearly an element $g \in G$ is divisible by p^n for all $n \in \mathbb{N}$ iff $g \in \mathbb{Z}[1/p]$. Hence if $\pi \in \operatorname{Aut}(G)$, then $\pi[\mathbb{Z}[1/p]] = \mathbb{Z}[1/p]$ for each prime p. It is easy to see that any automorphism of $\mathbb{Z}[1/p]$ is just multiplication by some element $u \in U_p = \{\pm p^n \mid n \in \mathbb{Z}\}$, the group of multiplicative units of the ring of p-adic rationals. Thus $\operatorname{Aut}(G) \simeq \prod_p U_p$ and so $|\operatorname{Aut}(G)| = 2^\omega$.

Next note that $U_p \simeq \mathbb{Z} \times C_2$ for each prime p. Hence $\operatorname{Aut}(G) \simeq P \times V$, where P is the direct product of countably many copies of $\mathbb{Z}$ and V is the direct product of countably many copies of C_2. Since each nonzero element of V has order 2, it follows that V is isomorphic to a direct sum of $|V|$ copies of C_2. So we can identify V with the vector space of dimension 2^ω over the field $\mathbb{F}_2$ of two elements. Hence

$$\operatorname{Aut}(\operatorname{Aut}(G)) \geqslant \operatorname{Aut}(P) \times \operatorname{Aut}(V) = \operatorname{Aut}(P) \times GL(V),$$

where $GL(V)$ is the general linear group on the vector space V. It follows that

$$2^{2^\omega} = |GL(V)| \leq |\operatorname{Aut}(\operatorname{Aut}(G))| \leq 2^{2^\omega}.$$

□

DEFINITION 5.2.10. If κ is an infinite cardinal, then the cardinal $\beth_\alpha(\kappa)$ is defined inductively as follows.

(a) $\beth_0(\kappa) = \kappa$.
(b) $\beth_{\alpha+1}(\kappa) = 2^{\beth_\alpha(\kappa)}$.
(c) If α is a limit ordinal, then $\beth_\alpha(\kappa) = \sup_{\beta<\alpha} \beth_\beta(\kappa)$

QUESTION 5.2.11. Does there exist a fixed ordinal β such that if G is an arbitrary infinite group, then $|G_\alpha| \leq \beth_\beta(|G|)$ for all α?

In Chapter 7, we shall prove a result which shows that in the classical case of centreless groups, the upper bound of $\tau(G) < \left(2^{|G|}\right)^+$ cannot be improved in ZFC. It is conceivable that the analogous result is true for the case of arbitrary groups. We shall say a little more about this possibility at the end of Section 7.1.

5.3. Notes

In retrospect, it is surprising that the notion of the automorphism tower of an arbitrary (perhaps centreless) group was not defined before Hamkins' elegant paper [**14**]. Theorems 5.1.6 and 5.1.8 are due to Hamkins [**14**]. Theorem 5.2.1 is an easy consequence of Robinson [**39**, Exercise 1.5.6].

Joel Hamkins has pointed out that it is slightly inaccurate to say that no better bound for $\tau(G)$ is known than the least inaccessible cardinal κ such that $\kappa > |G|$. For example, let λ be the least cardinal such that $V_\lambda \vDash ZFC$ and $\lambda > |G|$. Then an easy Lowenheim-Skolem argument shows that λ is strictly less than κ. Clearly we can assume that $G \in V_\lambda$. Since ZFC proves that the automorphism tower of G

eventually terminates, it follows that

$$V_\lambda \models \text{ The automorphism tower of } G \text{ terminates.}$$

Clearly the notion of the automorphism tower of G is absolute between V_λ and V. Hence $\tau(G) < \lambda$.

CHAPTER 6

Set-theoretic Forcing

The first half of this book presented those results on automorphism towers which can currently be proved using the classical ZFC axioms of set theory. The final three chapters will explain why a number of natural problems concerning automorphism towers are independent of ZFC. For example, we shall prove that if κ is a regular uncountable cardinal, then ZFC cannot decide the question of whether there exists a centreless group G of cardinality κ such that $\tau(G) \geq 2^{\kappa}$. More precisely, in Chapter 7, we shall prove that if κ is a (possibly singular) uncountable cardinal, then it is consistent with ZFC that there exists a centreless group G of cardinality κ such that $\tau(G) > 2^{\kappa}$; and in Chapter 9, we shall prove that it is also consistent with ZFC that $\tau(G) < 2^{\kappa}$ for every centreless group G of regular uncountable cardinality κ. In both cases, we shall use the technique of set-theoretic forcing to construct a suitable model of ZFC which satisfies the relevant group-theoretic statement.

The first five sections of this chapter contain a short introduction to set-theoretic forcing, which is aimed primarily at nonlogicians. In these sections, we shall present a detailed discussion of the fundamental concepts and basic results of set-theoretic forcing. Proofs of results will be provided only when they are both easy and also illustrate important ideas which will be needed in the later chapters of this book. (For a thorough introduction to set-theoretic forcing, the reader should consult the excellent textbook of Kunen [**26**].) In Section 6.8, we shall discuss some of the basic ideas of iterated forcing, including the notion of a reverse Easton iterated forcing.

Those readers who are already familiar with the basic theory of set-theoretic forcing need only read Sections 6.6 and 6.7. In Section 6.6, we shall prove that the height $\tau(G)$ of the automorphism tower of an infinite centreless group G is not necessarily an absolute concept. In Section 6.7, we shall prove some partial absoluteness results for the heights of automorphism towers, which will be used repeatedly in the later chapters of this book.

In the first four sections of this chapter, V will always denote the actual set-theoretic universe.

6.1. Countable transitive models of ZFC

Suppose that σ is a sentence in the first-order language of set theory. By Gödel's Completeness Theorem, σ is consistent with ZFC, written $\mathrm{Con}(ZFC+\sigma)$, iff there exists a countable model $\langle M, E\rangle$ of $ZFC + \sigma$. (Here the relation E is the interpretation of the membership symbol $\in$ in M.) This chapter presents an account of set-theoretic forcing: a technique for constructing countable models of theories such as $ZFC + 2^\omega = \omega_2$, etc. Unfortunately, these constructions cannot be carried out using just the axioms of ZFC, since Gödel's Incompleteness Theorem implies that $\mathrm{Con}(ZFC)$ is not a theorem of ZFC. Throughout this book, our basic assumption is that ZFC is consistent and hence there exists a countable model $\langle M, E\rangle$ of ZFC. In fact, we shall make the stronger assumption that the axioms of ZFC are true in the actual set-theoretic universe V. This implies that there exists a countable model $\langle M, E\rangle$ of ZFC which satisfies the following further hypotheses. (For more details, see Section IV.7 of Kunen [**26**]. As we indicated above, while it is true platonistically that there exists a countable transitive model of ZFC, the proof of its existence cannot be formalised within ZFC. A clear discussion of this point, together with an account of other less platonistic approaches to forcing, can be found in Section VII.9 of Kunen [**26**].)

6.1.1(a) The relation E is the genuine membership relation on M; i.e.

$$E = \{\langle x, y\rangle \in M \times M \mid x \in y\}.$$

Thus our model has the form $\langle M, \in\rangle$.

6.1.1(b) M is a *transitive set*; i.e. if $x \in M$ and $y \in x$, then $y \in M$.

Definition 6.1.1. If the countable model M of ZFC satisfies hypotheses 6.1.1(a) and 6.1.1(b), then M is said to be a *countable transitive model (c.t.m.)*

If M is a c.t.m., then M contains a canonically defined transitive submodel L^M such that $L^M \vDash ZFC + GCH$. (For example, see Chapter VI of Kunen [**26**].)

Hence, whenever it is convenient, we can also assume that the c.t.m. M satisfies GCH.

In the next chapter, we shall prove that there exists a c.t.m. M such that

$$M \vDash \text{ There exists an infinite centreless group } G \text{ such that } \tau(G) > 2^{|G|}.$$

To accomplish this, we must not only construct the c.t.m. M, but also calculate the automorphism tower of G within M. More generally, we shall need to understand the set-theoretic and algebraic properties of various sets and structures within the c.t.m. M. Fortunately, this is not too difficult, since Hypotheses 6.1.1(a) and 6.1.1(b) imply that many basic set-theoretic and algebraic properties $P(x_1, \ldots, x_n)$ are absolute for M; i.e. for all $a_1, \cdots, a_n \in M$,

$$M \vDash P(a_1, \ldots, a_n) \quad \text{iff} \quad V \vDash P(a_1, \ldots, a_n).$$

For example, suppose that $a, b \in M$. It is clear that if $V \vDash a \subseteq b$, then $M \vDash a \subseteq b$. Conversely, suppose that $M \vDash a \subseteq b$. Working within V, the transitivity of M implies that if $c \in a$, then $c \in M$ and hence $c \in b$. Thus $V \vDash a \subseteq b$. Consequently, for all $a, b \in M$,

$$M \vDash a \subseteq b \quad \text{iff} \quad V \vDash a \subseteq b.$$

DEFINITION 6.1.2. A *transitive class model* M of ZFC is a transitive class such that $\langle M, \in\rangle \vDash ZFC$. In particular, we allow the possibility that M is a proper class. In this book, we shall only be interested in the cases when M is either a c.t.m. or the actual universe V.

DEFINITION 6.1.3. Let $\varphi(x_1, \ldots, x_n)$ be a formula with free variables $x_1, \ldots, x_n$ and let M be a transitive class model of ZFC. Then φ is *absolute* for M if for all $a_1, \cdots, a_n \in M$,

$$M \vDash \varphi(a_1, \ldots, a_n) \quad \text{iff} \quad V \vDash \varphi(a_1, \ldots, a_n).$$

Notice that if φ is absolute for every transitive class model of ZFC, then whenever M, N are transitive class models of ZFC such that $M \subseteq N$,

$$M \vDash \varphi(a_1, \ldots, a_n) \quad \text{iff} \quad N \vDash \varphi(a_1, \ldots, a_n)$$

for all $a_1, \cdots, a_n \in M$.

We have already seen that the formula "$x \subseteq y$" is absolute for every transitive class model M of ZFC. With a little more effort, it can be shown that the formulas "x is an ordinal" and "$x = \omega$" are also absolute for every transitive class model M. On the other hand, if M is a c.t.m., then the formula "x is uncountable" is not absolute for M. To see this, let M be a c.t.m. and let $s \in M$ be any element such that $M \vDash s$ is uncountable. Since M is transitive and $s \in M$, it follows that $s \subseteq M$. Hence, working within the actual universe V, we see that s is really only a countably infinite set. Of course, this means that if $f : \omega \to s$ is any of the bijections which witness the countability of s, then necessarily $f \in V \smallsetminus M$.

In order to determine whether a formula φ is absolute, it is usually sufficient to consider whether or not φ is equivalent to a formula which only involves bounded quantifiers.

Definition 6.1.4. Let φ be a formula and let x, y be variables. Then we write $(\exists x \in y)\varphi$ as an abbreviation for $\exists x(x \in y \wedge \varphi)$. Similarly, we write $(\forall x \in y)\varphi$ as an abbreviation for $\forall x(x \in y \to \varphi)$. We say that $(\exists x \in y)$ and $(\forall x \in y)$ are *bounded quantifiers*.

Definition 6.1.5. The set of Δ_0*-formulas* are defined inductively as follows.

(i) "$x \in y$" and "$x = y$" are Δ_0-formulas.
(ii) If φ and ψ are Δ_0-formulas, then $\neg\varphi$, $\varphi \wedge \psi$, $\varphi \vee \psi$, $\varphi \to \psi$ and $\varphi \leftrightarrow \psi$ are also Δ_0-formulas.
(iii) If φ is a Δ_0-formula and x, y are variables, then $(\exists x \in y)\varphi$ and $(\forall x \in y)\varphi$ are also Δ_0-formulas.

Lemma 6.1.6. *Let $\varphi(x_1, \dots, x_n)$, $\psi(x_1, \dots, x_n)$ be formulas with free variables $x_1, \dots, x_n$. If ψ is a Δ_0-formula and $ZFC \vdash \forall x_1 \cdots \forall x_n(\varphi \leftrightarrow \psi)$, then φ is absolute for every transitive class model M of ZFC.*

Proof. Since M is transitive, $s \subseteq M$ for each $s \in M$. Hence each bounded quantifier $(\exists x \in s)$, $(\forall x \in s)$ has the same interpretation in M and V. □

The next result mentions just a few of the many set-theoretic notions which are easily seen to be definable by Δ_0-formulas.

THEOREM 6.1.7. *The following set-theoretic notions are definable by Δ_0-formulas and hence are absolute for every transitive class model M of ZFC.*

(a) *z is an ordered pair.*
(b) *R is a binary relation on the set A.*
(c) *f is a bijection from A onto B.*
(d) *α is an ordinal.*
(e) *n is a natural number.*
(f) *z is a finite sequence.*

$\square$

Many algebraic notions are also easily seen to be definable by Δ_0-formulas. For example, suppose that M is a transitive class model of ZFC and that $\langle G, f, e\rangle \in M$, where $e \in G$ and f is a binary operation on G. (Of course, since M is transitive, it follows that $e,\ f \in M$.) The axioms of group theory are naturally expressed as Δ_0-formulas; e.g.

$$(\forall x \in G)(\exists y \in G) f(x, y) = e.$$

It follows that the notion "$\langle G, f, e\rangle$ is a group" is definable by a Δ_0-formula and hence is absolute for every transitive class model M of ZFC. Similar remarks apply to the notions of a graph, a field, etc.

THEOREM 6.1.8. *The following algebraic notions are definable by Δ_0-formulas and hence are absolute for every transitive class model M of ZFC.*

(a) *$\langle \Gamma, E\rangle$ is a connected graph.*
(b) *$\langle G, \times, 1\rangle$ is a centreless group.*
(c) *$\langle F, +, \times, 0, 1\rangle$ is an algebraically closed field.*

$\square$

We have already seen that if M is a c.t.m., then the notion "x is uncountable" is not absolute for M. Of course, this means that the notion "x is countable" is also not absolute for M. However, suppose that $s \in M$ and that $M \vDash s$ is countable. Then there exists a function $f \in M$ such that $M \vDash f : \omega \to s$ is a surjection. Since the notion of a surjection is definable by a Δ_0-formula, it follows that f is also a surjection in the actual universe V and hence $V \vDash s$ is countable. Consequently, for all $s \in M$, if $M \vDash s$ is countable, then $V \vDash s$ is countable.

Definition 6.1.9. Let $\varphi(x_1, \dots, x_n)$ be a formula with free variables $x_1, \dots, x_n$ and let M be a transitive class model of ZFC. Then φ is *upwards absolute* for M if for every transitive class model N of ZFC such that $M \subseteq N$, for all $a_1, \cdots, a_n \in M$,

$$M \vDash \varphi(a_1, \dots, a_n) \quad \text{implies} \quad N \vDash \varphi(a_1, \dots, a_n).$$

As the above example suggests, there is a simple syntactical criterion which usually suffices to determine whether a notion is upwards absolute.

Definition 6.1.10. The formula φ is a Σ_1*-formula* iff φ has the form

$$\exists y_1 \cdots \exists y_m \psi,$$

where ψ is a Δ_0-formula.

Lemma 6.1.11. *Let $\varphi(x_1, \dots, x_n)$, $\psi(x_1, \dots, x_n)$ be formulas with free variables $x_1, \dots, x_n$. If ψ is a Σ_1-formula and $ZFC \vdash \forall x_1 \cdots \forall x_n(\varphi \leftrightarrow \psi)$, then φ is upwards absolute for every transitive class model M of ZFC.*

□

The following theorem gives two more examples of notions which are easily seen to be definable by Σ_1-formulas. As we shall see later, neither of these notions is absolute.

Theorem 6.1.12. *The following notions are definable by Σ_1-formulas and hence are upwards absolute for every transitive class model M of ZFC.*

(a) *$\langle \Gamma_1, E_1 \rangle$ and $\langle \Gamma_2, E_2 \rangle$ are isomorphic graphs.*
(b) *$\langle G, \times, 1 \rangle$ is a group such that $\tau(G) > 0$.*

□

It occasionally requires a little more thought to determine whether a notion is absolute. For example, consider the notion "G is a simple group." The most obvious definition of this notion is:

$$\neg(\exists H)(H \text{ is a nontrivial proper normal subgroup of } G,$$

which is certainly not a Δ_0-definition; and, at first glance, it seems conceivable that there might exist a transitive class model M of ZFC and a group $G \in M$ such that $M \vDash G$ is simple, while $V \vDash G$ is not simple. Of course, this would mean that G has nontrivial proper normal subgroups in V, but that each such subgroup H satisfies $H \in V \smallsetminus M$. However, notice that if G has a nontrivial proper normal subgroup in V, then it has such a subgroup of the form $H = \langle\, g^G \,\rangle$ for some $1 \neq g \in G$. But, using Theorem 6.1.7(f), we see that $\langle\, g^G \,\rangle$ has a Δ_0-definition in the language of set theory using the parameters g, $G \in M$. It follows that $\langle\, g^G \,\rangle \in M$ and hence $M \vDash G$ is not simple, which is a contradiction. Hence the notion "G is a simple group" is absolute for every transitive class model M of ZFC. An examination of the above argument shows that the notion "G is a simple group" also has the following slightly less obvious Δ_0-defintion:

$$(\forall x \in G)(\forall y \in G)(x \neq 1 \text{ implies } y \in \langle x^G \rangle).$$

We have already noted that if M is a c.t.m., then every set $s \in M$ is really countable within V. However, since $M \vDash ZFC$, it follows that there exists an ordinal $\alpha \in M$ such that

$$M \vDash \alpha \text{ is the least uncountable cardinal.}$$

We shall denote this ordinal by ω_1^M and use a similar notation for the M-versions of other definable objects. For example, $\mathrm{Sym}^M(\omega)$ will denote the group $G \in M$ such that

$$M \vDash G \text{ is the group of permutations of } \omega.$$

Since the notion of a permutation of ω is absolute, we have that

$$\mathrm{Sym}^M(\omega) = \mathrm{Sym}(\omega) \cap M.$$

Here $\mathrm{Sym}(\omega)$ denotes the actual symmetric group on ω. We shall occasionally use notation such as $\mathrm{Sym}^V(\omega)$ when we want to emphasise that we are referring to an object within V rather the M-version within some c.t.m. M.

In the remainder of this section, we shall say a few words concerning the basic strategy of set-theoretic forcing. Suppose that we wish to prove the consistency of $\neg CH$. Let M be any c.t.m. Working within the actual universe V, we see that $\mathcal{P}^M(\omega)$ is really only a countable subset of $\mathcal{P}(\omega)$ and that ω_2^M is a countable ordinal. In particular, there exist uncountably many sets $S \in \mathcal{P}(\omega) \smallsetminus \mathcal{P}^M(\omega)$ and it seems reasonable to hope that we can obtain of model of $ZFC + \neg CH$ by extending M to a suitable c.t.m. N such that $\mathcal{P}^M(\omega) \subsetneq \mathcal{P}^N(\omega)$. For example, let $f \in V \smallsetminus M$

be an injection $f : \omega_2^M \to \mathcal{P}(\omega)$ such that $f(\alpha) = S_\alpha \in \mathcal{P}(\omega) \smallsetminus \mathcal{P}^M(\omega)$ for each $\alpha < \omega_2^M$ and suppose that $N \supsetneq M$ is a c.t.m. such that $f \in N$. Then N contains the ω_2^M-sequence $\langle S_\alpha \mid \alpha < \omega_2^M \rangle$ of distinct "new" subsets $S_\alpha \subseteq \omega$. At first glance, this seems to imply that that CH is false in N. Unfortunately, it only shows that

$$N \vDash 2^\omega \geq |\omega_2^M|,$$

and there remains the possibility that the cardinal ω_2^M of M has been "collapsed" within N; i.e. that N also contains a bijection $g : \omega_1^M \to \omega_2^M$ or, even worse, a bijection $h : \omega \to \omega_2^M$. If this occurs, then

$$N \vDash |\omega_2^M| \leq \omega_1$$

and we have accomplished nothing. To see that this is a genuine problem, suppose that we inadvertently choose the sets $\langle S_\alpha \mid \alpha < \omega_2^M \rangle$ so that $\min S_\alpha \neq \min S_\beta$ for all $\alpha < \beta < \omega_2^M$. Then N also contains the injection $h : \omega_2^M \to \omega$, defined by

$$h(\alpha) = \min S_\alpha,$$

and so ω_2^M is indeed a countable ordinal in N. However, it seems reasonable to hope that this kind of problem will not arise if the sequence $\langle S_\alpha \mid \alpha < \omega_2^M \rangle$ is sufficiently "generic". In the next section, we shall begin our study of set-theoretic forcing: a technique for adjoining "generic" objects to countable transitive models of ZFC.

Convention 6.1.13. From now on, we shall often use phrases such as:

- "let $\kappa \in M$ be a regular cardinal such that $\kappa^{<\kappa} = \kappa$",

instead of the more accurate:

- "let $M \vDash \kappa$ is a regular cardinal such that $\kappa^{<\kappa} = \kappa$".

6.2. Set-theoretic forcing

In this section, V will continue to denote the actual set-theoretic universe.

Definition 6.2.1. A *notion of forcing* is an ordered triple $\langle \mathbb{P}, \leq, 1 \rangle$ such that $\leq$ partially orders $\mathbb{P}$ and 1 is the largest element of $\mathbb{P}$. The elements of $\mathbb{P}$ are called *conditions*.

Slightly abusing notation, we shall usually say that $\mathbb{P}$ is a notion of forcing. On the other hand, when more than one notion of forcing is under discussion, we shall occasionally need to use the more explicit notation $\langle \mathbb{P}, \leq_{\mathbb{P}}, 1_{\mathbb{P}} \rangle$.

A notion of forcing typically arises as a collection of approximations to some object that we wish to "generically adjoin". For example, suppose that we wish to adjoin an ω_2-sequence $\langle g_\alpha \mid \alpha < \omega_2 \rangle$ of distinct functions $g_\alpha \in {}^\omega 2$. Equivalently, we wish to adjoin the associated function $g : \omega_2 \times \omega \to 2$, defined by

$$g(\alpha, n) = g_\alpha(n).$$

Then a suitable notion of forcing would be the set $\mathrm{Fn}(\omega_2 \times \omega, 2)$ of finite approximations to such a function.

Definition 6.2.2. If I, J are any sets, then $\mathrm{Fn}(I, J)$ is the notion of forcing consisting of all functions p such that

(a) $\operatorname{dom} p \subseteq I$,

(b) $\operatorname{ran} p \subseteq J$, and

(c) $|p| < \omega$,

ordered by $q \leq p$ iff $q \supseteq p$.

As this example suggests, if $\mathbb{P}$ is an arbitrary notion of forcing and p, $q \in \mathbb{P}$ are conditions, then the intuitive meaning of $q < p$ is that "q contains more information than p".

Definition 6.2.3. Let $\mathbb{P}$ be a notion of forcing.

(a) The conditions p, $q \in \mathbb{P}$ are *compatible* iff there exists $r \in \mathbb{P}$ such that $r \leq p$, q.

(b) Otherwise, p and q are *incompatible*, written $p \perp q$.

Definition 6.2.4. Let $\mathbb{P}$ be a notion of forcing.

(a) A subset $D \subseteq \mathbb{P}$ is *dense* iff for all $p \in \mathbb{P}$, there exists $q \in D$ such that $q \leq p$.

(b) A subset $G \subseteq \mathbb{P}$ is a *filter* iff the following conditions are satisfied.

  (i) For all p, $q \in G$, there exists $r \in G$ such that $r \leq p$, q.

  (ii) For all p, $q \in \mathbb{P}$, if $p \in G$ and $p \leq q$, then $q \in G$.

For example, consider the notion of forcing $\mathbb{P} = \mathrm{Fn}(\omega_2 \times \omega, 2)$. For each pair of ordinals $\alpha < \beta < \omega_2$, let $C_{\alpha,\beta}$ consist of the conditions $p \in \mathbb{P}$ for which there exists an integer $n \in \omega$ such that

(1) $\langle \alpha, n\rangle$, $\langle \beta, n\rangle \in \operatorname{dom} p$; and
(2) $p(\alpha, n) \neq p(\beta, n)$.

Then $C_{\alpha,\beta}$ is dense in $\mathbb{P}$ for $\alpha < \beta < \omega_2$. Similarly, it is clear that for each pair $\langle \alpha, n\rangle \in \omega_2 \times \omega$, the subset

$$D_{\alpha,n} = \{p \in \mathbb{P} \mid \langle \alpha, n\rangle \in \operatorname{dom} p\}$$

is dense in $\mathbb{P}$. Now suppose that $G \subseteq \mathbb{P}$ is a filter. Then any two of the partial functions p, $q \in G$ are compatible and so $g = \bigcup G$ is a partial function from $\omega_2 \times \omega$ into 2. Clearly g will be a total function iff $G \cap D_{\alpha,n} \neq \emptyset$ for all $\langle \alpha, n\rangle \in \omega_2 \times \omega$. In this case, the corresponding sequence $\langle g_\alpha \mid \alpha < \omega_2\rangle$ will consist of *distinct* functions $g_\alpha \in {}^{\omega}2$ iff $G \cap C_{\alpha,\beta} \neq \emptyset$ for all $\alpha < \beta < \omega_2$. Of course, such a filter G exists iff CH is false in the actual universe V.

Up to this point, our discussion of notions of forcing, dense sets, etc. has taken place in V. However, if we wish to generically adjoin a sequence $\langle g_\alpha \mid \alpha < \omega_2^M\rangle$ of functions $g_\alpha \in {}^{\omega}2$ to a c.t.m. M, then it is necessary to consider the corresponding relativized versions within M. Fortunately, all of the basic ingredients of forcing are absolute, including "$\leq$ partially orders $\mathbb{P}$", "$p \leq q$", "$p \perp q$", "D is a dense subset of $\mathbb{P}$", etc.

Definition 6.2.5. Suppose that M is a c.t.m. and that $\mathbb{P} \in M$ is a notion of forcing. Then the filter $G \subseteq \mathbb{P}$ is *$\mathbb{P}$-generic over M* iff $G \cap D \neq \emptyset$ for all dense subsets D of $\mathbb{P}$ such that $D \in M$.

But how can we tell whether a dense subset D of $\mathbb{P}$ is an element of the c.t.m. M? In practice, this is not a source of difficulty, since each dense set D that is actually used in our arguments will be explicitly definable; i.e. there will always exist a formula $\varphi(y, x_1, \cdots, x_n)$ in the language of set theory with free variables $y, x_1, \cdots, x_n$ and elements $a_1, \cdots, a_n \in M$ such that

$$D = \{p \in \mathbb{P} \mid M \vDash \varphi(p, a_1, \cdots, a_n)\}.$$

Since $M \vDash ZFC$, it follows that $D \in M$. In fact, the formula $\varphi(y, x_1, \cdots, x_n)$ will usually be absolute for M, so that

$$D = \{p \in \mathbb{P} \mid V \vDash \varphi(p, a_1, \cdots, a_n)\}.$$

The countability of the c.t.m. M is used only once in the development of set-theoretic forcing; namely, in the following proof that there always exists a $\mathbb{P}$-generic filter over M.

LEMMA 6.2.6. *Suppose that M is a c.t.m. and that $\mathbb{P} \in M$ is a notion of forcing. For each $p \in \mathbb{P}$, there exists a filter $G \subseteq \mathbb{P}$ such that $p \in G$ and G is $\mathbb{P}$-generic over M.*

PROOF. Let $\{D_n \mid n < \omega\}$ be an enumeration of the countably many dense subsets D of $\mathbb{P}$ such that $D \in M$. (Of course, this enumeration is usually not an element of M. But here we are working within the actual universe V rather than within the countable subuniverse M.) Then we can inductively define a sequence

$$p = p_0 \geq p_1 \geq \cdots \geq p_n \geq \cdots$$

of elements of $\mathbb{P}$ such that $p_{n+1} \in D_n$ for all $n < \omega$. Let

$$G = \{q \in \mathbb{P} \mid \text{ There exists } n < \omega \text{ such that } p_n \leq q\}.$$

Then G is a $\mathbb{P}$-generic filter over M such that $p \in G$. □

For example, let M be a c.t.m. and consider the notion of forcing

$$\mathbb{Q} = \mathrm{Fn}(\omega_2 \times \omega, 2)^M = \mathrm{Fn}(\omega_2^M \times \omega, 2) \in M.$$

Let G be a $\mathbb{Q}$-generic filter over M. For each $\alpha < \beta < \omega_2^M$ and $\langle \alpha, n\rangle \in \omega_2^M \times \omega$, let $C_{\alpha,\beta}$ and $D_{\alpha,n}$ be the corresponding dense subsets of $\mathbb{Q}$, defined in the discussion following Definition 6.2.4. Then $C_{\alpha,\beta}$, $D_{\alpha,n} \in M$ and hence $G \cap C_{\alpha,\beta} \neq \emptyset$ and $G \cap D_{\alpha,n} \neq \emptyset$ for each $\alpha < \beta < \omega_2^M$ and each $\langle \alpha, n\rangle \in \omega_2^M \times \omega$. It follows that $g = \bigcup G$ is a total function from $\omega_2^M \times \omega$ to 2 and that the corresponding sequence $\langle g_\alpha \mid \alpha < \omega_2^M\rangle$ consists of distinct functions $g_\alpha \in {}^\omega 2$.

As we shall see next, for most notions of forcing $\mathbb{P} \in M$, the c.t.m. M does not contain any $\mathbb{P}$-generic filters. (The proof of Lemma 6.2.8 also shows that if $\mathbb{P}$ is an atomless notion of forcing, then there does not exist a filter G which is $\mathbb{P}$-generic over V. Thus it is essential that we force over a countable transitive model M of ZFC rather than over the actual set-theoretic universe V.)

DEFINITION 6.2.7. Let $\mathbb{P}$ be a notion of forcing. An element $p \in \mathbb{P}$ is said to be an *atom* iff there do not exist elements $q_1, q_2 \leq p$ such that $q_1 \perp q_2$. The notion of forcing $\mathbb{P}$ is *nonatomic* iff $\mathbb{P}$ does not contain any atoms.

Notice that if $\mathbb{P} \in M$ and $p \in \mathbb{P}$ is an atom, then $G = \{q \in \mathbb{P} \mid q \leq p \text{ or } p \leq q\}$ is a $\mathbb{P}$-generic filter over M such that $G \in M$.

LEMMA 6.2.8. *Suppose that M is a c.t.m. and that $\mathbb{P} \in M$ is a atomless notion of forcing. If G is a $\mathbb{P}$-generic filter over M, then $G \notin M$.*

PROOF. Suppose that G is a $\mathbb{P}$-generic filter over M such that $G \in M$. Then $D = \mathbb{P} \smallsetminus G \in M$. We claim that D is a dense subset of $\mathbb{P}$. To see this, let $p \in \mathbb{P}$ be any element. Since p is not an atom, there exist $q_1, q_2 \leq p$ such that $q_1 \perp q_2$; and since G is a filter, there exists $i \in \{1, 2\}$ such that $q_i \in \mathbb{P} \smallsetminus G = D$. As G is $\mathbb{P}$-generic over M, we must have that $G \cap D = G \cap (\mathbb{P} \smallsetminus G) \neq \emptyset$, which is a contradiction. □

Suppose that M is a c.t.m., $\mathbb{P} \in M$ is a notion of forcing and that G is a $\mathbb{P}$-generic filter over M. We shall next describe how to construct the corresponding generic extension $M[G]$; i.e. the smallest c.t.m. N of ZFC such that $M \subseteq N$ and $G \in N$. The basic idea is that since $M[G]$ is the smallest c.t.m. containing G, for each element $a \in M[G]$, there should be a corresponding $\mathbb{P}$-name $\tau \in M$ which describes how a may be constructed from G.

DEFINITION 6.2.9. Let $\mathbb{P}$ be a notion of forcing. Then τ is a $\mathbb{P}$-*name* iff

(a) τ is a set of ordered pairs; and
(b) if $\langle \sigma, p \rangle \in \tau$, then σ is a $\mathbb{P}$-name and $p \in \mathbb{P}$.

DEFINITION 6.2.10. Suppose that M is a c.t.m. and that $\mathbb{P} \in M$ is a notion of forcing. Let G be a $\mathbb{P}$-generic filter over M. Then for each $\mathbb{P}$-name $\tau \in M$, we define the corresponding *interpretation* by

$$\tau_G = \{\sigma_G \mid \text{ There exists } p \in G \text{ such that } \langle \sigma, p \rangle \in \tau\}.$$

For example, the empty set $\emptyset$ is trivially a $\mathbb{P}$-name and clearly $\emptyset_G = \emptyset$ for every $\mathbb{P}$-generic filter G. Hence if $p \in \mathbb{P}$, then $\{\langle \emptyset, p \rangle\}$ also a $\mathbb{P}$-name and

$$\{\langle \emptyset, p \rangle\}_G = \begin{cases} \{\emptyset\}, & \text{if } p \in G; \\ \emptyset, & \text{if } p \notin G. \end{cases}$$

Slightly less trivially, let 1 be the largest element of $\mathbb{P}$. Then for each $x \in M$, we can recursively define the *canonical* $\mathbb{P}$*-name* $\check{x} \in M$ by

$$\check{x} = \{\langle \check{y}, 1 \rangle \mid y \in x\}.$$

If G is any $\mathbb{P}$-generic filter, then $1 \in G$ and so an easy induction shows that $\check{x}_G = x$ for all $x \in M$.

Definition 6.2.11. Suppose that M is a c.t.m. and that $\mathbb{P} \in M$ is a notion of forcing. If G is a $\mathbb{P}$-generic filter over M, then the corresponding *generic extension* is defined to be

$$M[G] = \{\tau_G \mid \tau \in M \text{ is a } \mathbb{P}\text{-name }\}.$$

Theorem 6.2.12. *Suppose that M is a c.t.m. and that $\mathbb{P} \in M$ is a notion of forcing. If G is a $\mathbb{P}$-generic filter over M, then*

(a) *$M[G]$ is a countable transitive model of ZFC;*
(b) *$M \subseteq M[G]$ and $G \in M[G]$;*
(c) *M and $M[G]$ contain the same ordinals.*

Furthermore, if N is a transitive class model of ZFC such that $M \subseteq N$ and $G \in N$, then $M[G] \subseteq N$.

A complete proof of Theorem 6.2.12 can be found in Chapter VII of Kunen [**26**]. In this section, we shall only check some of the easier parts of Theorem 6.2.12, in order to illustrate the basic ideas. First we shall prove that the pairing axiom

$$(\forall x)(\forall y)(\exists z)(x \in z \wedge y \in z)$$

holds in $M[G]$. Suppose that a, $b \in M[G]$. Then there exist $\mathbb{P}$-names τ, $\sigma \in M$ such that $\tau_G = a$ and $\sigma_G = b$. Let $\rho = \{\langle \tau, 1 \rangle, \langle \sigma, 1 \rangle\}$. Then $\rho \in M$ is a $\mathbb{P}$-name and clearly $\rho_G = \{a, b\}$.

Next to see that $M \subseteq M[G]$, recall that for each $x \in M$, we have already defined the canonical $\mathbb{P}$-name $\check{x} \in M$ by $\check{x} = \{\langle \check{y}, 1 \rangle \mid y \in x\}$ and that $\check{x}_G = x$ for all $x \in M$.

Finally we shall check that $G \in M[G]$. To see this, let $\Gamma \in M$ be the $\mathbb{P}$-name defined by $\Gamma = \{\langle \check{p}, p \rangle \mid p \in \mathbb{P}\}$. Then we clearly have that $\Gamma_G = G$.

Returning to our earlier example, let M be a c.t.m.,

$$\mathbb{Q} = \mathrm{Fn}(\omega_2 \times \omega, 2)^M = \mathrm{Fn}(\omega_2^M \times \omega, 2) \in M,$$

and suppose that G is a $\mathbb{Q}$-generic filter over M. Then the c.t.m. $M[G]$ contains the ω_2^M-sequence $\langle g_\alpha \mid \alpha < \omega_2^M \rangle \subseteq {}^\omega 2$ of distinct functions $g_\alpha \in {}^\omega 2$. Unfortunately, as we pointed out in Section 6.1, this only shows that

$$M[G] \vDash 2^\omega \geq |\omega_2^M|;$$

and if we wish to prove that $M \vDash \neg CH$, then it remains to be shown that $\omega_2^M = \omega_2^{M[G]}$. In other words, we must show that the cardinal ω_2^M of M has not been collapsed. We shall deal with this problem in the next section. Once again, this is a genuine problem. For example, consider the notion of forcing

$$\mathbb{R} = \mathrm{Fn}(\omega, \omega_2)^M = \mathrm{Fn}(\omega, \omega_2^M) \in M$$

and suppose that H is an $\mathbb{R}$-generic filter over M. For each $\alpha \in \omega_2^M$,

$$D_\alpha = \{p \in \mathbb{R} \mid \alpha \in \operatorname{ran} p\} \in M$$

is a dense subset of $\mathbb{R}$ and so $H \cap D_\alpha \neq \emptyset$. It follows easily that $h = \bigcup H \in M[H]$ is a surjective function from ω onto ω_2^M and hence ω_2^M is a countable ordinal in $M[H]$.

We shall end this section with a short discussion of the forcing relation $\Vdash$. It will soon become clear that this is the key to understanding the relationship between the combinatorial properties of the notion of forcing $\mathbb{P}$ and the statements which are true in the corresponding generic extensions.

Definition 6.2.13. Suppose that $\varphi(x_1, \dots, x_n)$ is a formula with free variables $x_1, \dots, x_n$ and that $\tau_1, \dots, \tau_n \in M$ are $\mathbb{P}$-names. If $p \in \mathbb{P}$, then we say that *p forces* $\varphi(\tau_1, \dots, \tau_n)$, written $p \Vdash \varphi(\tau_1, \dots, \tau_n)$, iff whenever G is a $\mathbb{P}$-generic filter over M such that $p \in G$, then $M[G] \vDash \varphi((\tau_1)_G, \dots, (\tau_n)_G)$.

Occasionally, when more than one forcing notion is under discussion, we shall need to use the more explicit notation $\Vdash_{\mathbb{P}}$.

To test his understanding of Definition 6.2.13, the reader should verify the following basic property of the forcing relation $\Vdash$ on $\mathbb{P}$.

OBSERVATION 6.2.14. *If $p \Vdash \varphi(\tau_1, \dots, \tau_n)$ and $q \le p$, then $q \Vdash \varphi(\tau_1, \dots, \tau_n)$.*

A complete proof of the following theorem can be found in Chapter VII of Kunen [**26**].

THEOREM 6.2.15. *Let $\varphi(x_1, \dots, x_n)$ be a formula with free variables $x_1, \dots, x_n$ and let $\tau_1, \dots, \tau_n \in M$ be $\mathbb{P}$-names.*

(a) *If G is a $\mathbb{P}$-generic filter over M, then $M[G] \vDash \varphi((\tau_1)_G, \dots, (\tau_n)_G)$ iff there exists $p \in G$ such that $p \Vdash \varphi(\tau_1, \dots, \tau_n)$.*

(b) *It may be decided within M whether or not $p \Vdash \varphi(\tau_1, \dots, \tau_n)$; i.e. there exists an alternative relation $\Vdash^*$ definable within M such that whenever $\tau_1, \dots, \tau_n \in M$ are $\mathbb{P}$-names, then*

$$p \Vdash \varphi(\tau_1, \dots, \tau_n) \quad \textit{iff} \quad M \vDash p \Vdash^* \varphi(\tau_1, \dots, \tau_n)$$

In order to get some feeling for why statement 6.2.15(a) is true, let M be a c.t.m. and let

$$\mathbb{Q} = \mathrm{Fn}(\omega_2 \times \omega, 2)^M = \mathrm{Fn}(\omega_2^M \times \omega, 2) \in M.$$

Let G be a $\mathbb{Q}$-generic filter over M. Then we have already seen that $g = \bigcup G$ is a function from $\omega_2^M \times \omega$ to 2. Since $G \in M[G]$, it follows that

$$S = \{n \in \omega \mid g(0, n) = 1\} \in M[G].$$

In fact, if we define

$$\sigma = \{\langle \check{n}, p\rangle \mid p \in \mathbb{P}, \langle 0, n\rangle \in \operatorname{dom} p \text{ and } p(0, n) = 1\},$$

then $\sigma_G = S$. Notice that for each $n \in \omega$,

$$\begin{aligned} M[G] \vDash \check{n}_G \in \sigma_G \quad & \text{iff} \quad \text{there exists } p \in G \text{ such that } p(0, n) = 1 \\ & \text{iff} \quad \text{there exists } p \in G \text{ such that } p \Vdash \check{n} \in \sigma. \end{aligned}$$

As for statement 6.2.15(b), first note that since the definition of $\Vdash$ involves the collection of all $\mathbb{P}$-generic filters, it is far from clear that the forcing relation can be decided within M. However, there is an alternative approach $\Vdash^*$ to forcing which avoids any reference to objects which lie outside M. (This approach $\Vdash^*$ is essentially just the result of a tedious inductive analysis of the forcing relation $\Vdash$, which can be carried out within M. It is briefly described in Section 9.1, where it is needed for another purpose.) But why should we care whether or not the forcing relation $\Vdash$ can be decided within M? Mainly because it allows us to see that various sets defined in terms of $\Vdash$ are elements of the c.t.m. M. For a typical application, let $\varphi(x_1, \dots, x_n)$ be a formula with free variables $x_1, \dots, x_n$ and let $\tau_1, \dots, \tau_n \in M$ be $\mathbb{P}$-names. Since the relation $\Vdash^*$ is definable in M, it follows that the set

$$D = \{p \in \mathbb{P} \mid p \Vdash \varphi(\tau_1, \dots, \tau_n) \text{ or } p \Vdash \neg\varphi(\tau_1, \dots, \tau_n)\}$$

is an element of M. In fact, $D \in M$ is a dense subset of $\mathbb{P}$. To see this, let $p \in \mathbb{P}$ be any element. By Lemma 6.2.6, there exists a $\mathbb{P}$-generic filter G over M such that $p \in G$. Suppose, for example, that $M[G] \vDash \varphi((\tau_1)_G, \dots, (\tau_n)_G)$. By Theorem 6.2.15(a), there exists $q \in G$ such that $q \Vdash \varphi(\tau_1, \dots, \tau_n)$. Since G is a filter, there exists $r \in G$ such that $r \leq p$, q and Observation 6.2.14 implies that $r \in D$. A similar argument proves the following basic property of the forcing relation $\Vdash$ on $\mathbb{P}$.

OBSERVATION 6.2.16. *Suppose that $B \in M$ and that $\tau_1, \dots, \tau_n \in M$ are $\mathbb{P}$-names. If $p \Vdash (\exists x \in \check{B})\varphi(x, \tau_1, \dots, \tau_n)$, then the set*

$$\{q \in \mathbb{P} \mid \textit{There exists } b \in B \textit{ such that } q \Vdash \varphi(\check{b}, \tau_1, \dots, \tau_n)\} \in M$$

is dense below p.

Here a subset $E \subseteq \mathbb{P}$ is said to be *dense below* p iff for all $q \leq p$, there exists $r \in E$ such that $r \leq q$.

6.3. Preserving cardinals

Throughout this section, M will denote a c.t.m. of ZFC and V will denote the actual set-theoretic universe.

Suppose that $\mathbb{P} \in M$ is a notion of forcing. In this section, we shall isolate certain combinatorial properties of $\mathbb{P}$ which are sufficient to ensure that $\mathbb{P}$ preserves various cardinals $\theta \in M$. We shall make frequent use of the following easy observation.

LEMMA 6.3.1. *Suppose that G is a $\mathbb{P}$-generic filter over M and that $p \in G$. If $E \in M$ is dense below p, then $G \cap E \neq \emptyset$.*

PROOF. Let $D = \{q \in \mathbb{P} \mid q \in E \text{ or } q \perp p\}$. Then $D \in M$ and it is easily checked that D is a dense subset of $\mathbb{P}$. Thus $G \cap D \neq \emptyset$. Since the elements of the filter G are pairwise compatible and $p \in G$, it follows that $G \cap E \neq \emptyset$. □

DEFINITION 6.3.2. Suppose that $\mathbb{P} \in M$ is a notion of forcing and that $\kappa \in M$ is an infinite cardinal.

(a) $\mathbb{P}$ *preserves cardinals greater or equal to* κ if whenever G is a $\mathbb{P}$-generic filter over M and $\theta \in M$ is a cardinal such that $\theta \geq \kappa$, then θ remains a cardinal in $M[G]$.

(b) $\mathbb{P}$ *preserves cofinalities greater or equal to* κ if whenever G is a $\mathbb{P}$-generic filter over M and $\gamma \in M$ is a limit ordinal with $\mathrm{cf}^M(\gamma) \geq \kappa$, then $\mathrm{cf}^{M[G]}(\gamma) = \mathrm{cf}^M(\gamma)$.

If the cardinal $\theta \in M$ is collapsed in the generic extension $M[G]$, then there exists an ordinal $\alpha < \theta$ such that $M[G]$ contains a surjective function $f : \alpha \to \theta$. Clearly this is impossible unless θ is an uncountable cardinal in M. Hence if $\mathbb{P}$ preserves cardinals greater or equal to ω_1, then $\mathbb{P}$ preserves every cardinal $\theta \in M$. In this case, we shall simply say that $\mathbb{P}$ preserves cardinals. A similar remark applies to cofinalities.

LEMMA 6.3.3. *Suppose that $\mathbb{P} \in M$ is a notion of forcing and that $\kappa \in M$ is an infinite cardinal.*

(a) *If $\mathbb{P}$ preserves cofinalities greater or equal to κ, then $\mathbb{P}$ also preserves cardinals greater or equal to κ.*

(b) *If $\mathbb{P}$ does not preserve cofinalities greater or equal to κ, then there exists a regular cardinal $\theta \in M$ with $\theta \geq \kappa$ and a $\mathbb{P}$-generic filter G over M such that $M[G] \vDash \theta$ is not a regular cardinal.*

Proof. (a) Suppose $\mathbb{P}$ preserves cofinalities greater or equal to κ and let G be a $\mathbb{P}$-generic filter over M. Let $\lambda \in M$ be a cardinal such that $\lambda \geq \kappa$. First suppose that λ is a regular cardinal in M. Since $\mathrm{cf}^{M[G]}(\lambda) = \mathrm{cf}^{M}(\lambda) = \lambda$, it follows that λ is also a regular cardinal in $M[G]$. On the other hand, if λ is a singular cardinal, then there exists a sequence $\langle \lambda_i \mid i < \mathrm{cf}^M(\lambda)\rangle$ of regular cardinals in M such that $\lambda = \sup \lambda_i$ and each $\lambda_i \geq \kappa$. Since each λ_i remains a cardinal in $M[G]$, it follows that λ is also a limit cardinal in $M[G]$.

(b) Suppose that whenever $\theta \in M$ is a regular cardinal such that $\theta \geq \kappa$ and G is a $\mathbb{P}$-generic filter over M, then θ remains a regular cardinal in $M[G]$. Let $\gamma \in M$ be a limit ordinal such that $\lambda = \mathrm{cf}^M(\gamma) \geq \kappa$. By Lemma 3.1.3, there exists a function $f \in M$ such that $f : \lambda \to \gamma$ is strictly increasing and cofinal. Applying Lemma 3.1.4 in $M[G]$, together with the fact that λ remains a regular cardinal in $M[G]$, we obtain that $\mathrm{cf}^{M[G]}(\gamma) = \mathrm{cf}^{M[G]}(\lambda) = \lambda$. □

Definition 6.3.4. Let $\mathbb{P}$ be a notion of forcing. A subset $\mathcal{A} \subseteq \mathbb{P}$ is an *antichain* iff the elements of $\mathcal{A}$ are pairwise incompatible.

Definition 6.3.5. Let $\mathbb{P}$ be a notion of forcing and let κ be an infinite cardinal. Then $\mathbb{P}$ has the *κ-chain condition* (*κ-c.c.*) iff every antichain of $\mathbb{P}$ has cardinality less than κ. If $\kappa = \omega_1$, then we say that $\mathbb{P}$ has the *countable chain condition* (*c.c.c.*)

Theorem 6.3.6. *Suppose that $\mathbb{P} \in M$ is a notion of forcing and that $\kappa \in M$ is a regular cardinal. If $M \vDash \mathbb{P}$ has the κ-*c.c.*, then $\mathbb{P}$ preserves cofinalities and cardinals greater or equal to κ.*

Of course, if $\mathbb{P} \in M$ is a notion of forcing, then $\mathbb{P}$ is really countable and hence has the *c.c.c.* in the actual set-theoretic universe V. However, as the statement of Theorem 6.3.6 indicates, in order to determine whether $\mathbb{P}$ preserves cardinals or cofinalities, we need to understand the combinatorial properties of $\mathbb{P}$ within the c.t.m. M. For example, consider the notion of forcing $\mathbb{P} = \mathrm{Fn}(\omega, \omega_1)^M = \mathrm{Fn}(\omega, \omega_1^M) \in M$. For each $\alpha < \omega_1^M$, let $p_\alpha \in \mathbb{P}$ be a condition such that $p_\alpha(0) = \alpha$. Then $\mathcal{A} = \{p_\alpha \mid \alpha < \omega_1^M\} \in M$ and $M \vDash \mathcal{A}$ is an uncountable antichain in $\mathbb{P}$. It is easily checked that if G is a $\mathbb{P}$-generic filter over M, then $g = \bigcup G \in M[G]$ is a surjective function from ω onto ω_1^M. Thus $\mathbb{P}$ does not preserve ω_1^M.

As we shall soon see, Theorem 6.3.6 is an easy consequence of the following important result.

LEMMA 6.3.7. *Suppose that $\mathbb{P} \in M$ is a notion of forcing, $\kappa \in M$ is a regular cardinal and that $M \vDash \mathbb{P}$ has the κ-c.c.. Let G be a $\mathbb{P}$-generic filter over M. Suppose that A, $B \in M$ and that $f \in M[G]$ with $f : A \to B$. Then there exists a function $F \in M$ such that*

(a) *$F : A \to \mathcal{P}(B)$;*
(b) *for all $a \in A$, $f(a) \in F(a)$; and*
(c) *for all $a \in A$, $M \vDash |F(a)| < \kappa$.*

PROOF. Let $\tau \in M$ be a $\mathbb{P}$-name such that $\tau_G = f$. Then there exists $p \in G$ such that

$$p \Vdash \tau \text{ is a function from } \check{A} \text{ into } \check{B}.$$

Define the function $F : A \to \mathcal{P}(B)$ by

$$F(a) = \{b \in B \mid \text{ There exists } q \leq p \text{ such that } q \Vdash \tau(\check{a}) = \check{b}\}.$$

Since the relation $\Vdash$ is decidable in M, it follows that $F \in M$. Let $a \in A$ and suppose that $f(a) = b$. Then there exists $r \in G$ such that $r \Vdash \tau(\check{a}) = \check{b}$. Let $q \in G$ satisfy $q \leq p, r$. Then $q \Vdash \tau(\check{a}) = \check{b}$ and so $b \in F(a)$.

Finally we shall show that $M \vDash |F(a)| < \kappa$. Applying the Axiom of Choice within M, there exists a function $Q \in M$ with $Q : F(a) \to \mathbb{P}$ such that for all $b \in F(a)$, $Q(b) \leq p$ and $Q(b) \Vdash \tau(\check{a}) = \check{b}$. We claim that if b_1, $b_2 \in F(a)$ are distinct, then $Q(b_1) \perp Q(b_2)$. If not, then there exists a $\mathbb{P}$-generic filter H over M such that $Q(b_1)$, $Q(b_2) \in H$. Clearly we also have that $p \in H$ and so

$$M[H] \vDash \text{ The function } \tau_H : A \to B \text{ satisfies } \tau_H(a) = b_1 \text{ and } \tau_H(a) = b_2,$$

which is a contradiction. Thus $Q : F(a) \to \mathbb{P}$ is an injection and $\{Q(b) \mid b \in F(a)\}$ is an antichain in $\mathbb{P}$. Since $M \vDash \mathbb{P}$ has the κ-*c.c.*, it follows that $M \vDash |F(a)| < \kappa$. □

PROOF OF THEOREM 6.3.6. Suppose not. Then by Lemma 6.3.3, there exists a regular cardinal $\theta \in M$ with $\theta \geq \kappa$ and a $\mathbb{P}$-generic filter G over M such that $M[G] \vDash \theta$ is not a regular cardinal. Thus there exists an ordinal $\alpha < \theta$ and a function $f \in M[G]$ such that f maps α cofinally into θ. By Lemma 6.3.7, there exists a function $F \in M$ such that

(a) $F : \alpha \to \mathcal{P}(\theta)$;

(b) for all $\xi < \alpha$, $f(\xi) \in F(\xi)$; and

(c) for all $\xi < \alpha$, $M \vDash |F(\xi)| < \kappa$.

Let $S = \bigcup_{\xi<\alpha} F(\xi)$. Then $S \in M$ and S is a cofinal subset of θ. But computing the cardinality of S in M, we see that $M \vDash |S| < \theta$ and this contradicts the assumption that $M \vDash \theta$ is a regular cardinal. □

Let $\mathbb{Q} = \mathrm{Fn}(\omega_2 \times \omega, 2)^M = \mathrm{Fn}(\omega_2^M \times \omega, 2) \in M$ and let G be a $\mathbb{Q}$-generic filter over M. In Section 6.2, we saw that $M[G] \vDash 2^\omega \geq |\omega_2^M|$. Our next target will be to prove that $M \vDash \mathbb{Q}$ has the *c.c.c.* This implies that $\mathbb{Q}$ preserves cofinalities and cardinals. In particular, $\omega_2^M = \omega_2^{M[G]}$ and so $M[G] \vDash 2^\omega \geq \omega_2$.

Since $M \vDash ZFC$, in order to prove that $M \vDash \mathbb{Q}$ has the *c.c.c.*, it is enough to prove in ZFC that $\mathbb{P} = \mathrm{Fn}(\omega_2 \times \omega, 2)$ has the *c.c.c.* in the actual set-theoretic universe V.

Definition 6.3.8. A family $\mathcal{D}$ of sets is said to be a Δ-*system* iff there exists a fixed set R such that $A \cap B = R$ whenever A, B are distinct elements of $\mathcal{D}$. In this case, R is said to be the *root* of the Δ-system.

Theorem 6.3.9. *If $\mathcal{B}$ is any uncountable family of finite sets, then there is an uncountable family $\mathcal{D} \subseteq \mathcal{B}$ which forms a Δ-system.*

Proof. After passing to a suitable subset of $\mathcal{B}$ if necessary, we can assume that $|\mathcal{B}| = \omega_1$. Since $\mathrm{cf}(\omega_1) = \omega_1$, we can also assume that there exists an integer $n \geq 1$ such that $|B| = n$ for all $B \in \mathcal{B}$. Suppose inductively that the theorem holds for every family $\mathcal{F}$ of finite sets such that $|\mathcal{F}| = \omega_1$ for which there exists an integer $m < n$ with $|F| = m$ for all $F \in \mathcal{F}$. Let $S = \bigcup\{B \mid B \in \mathcal{B}\}$. There are two cases to consider.

First suppose that there exists an element $s \in S$ such that $\mathcal{C} = \{B \in \mathcal{B} \mid s \in B\}$ has cardinality ω_1. Applying the inductive hypothesis, there exists an uncountable Δ-system $\mathcal{D}^* \subseteq \mathcal{C}^* = \{B \smallsetminus \{s\} \mid B \in \mathcal{C}\}$ with root R^*. It follows that

$$\mathcal{D} = \{D \cup \{s\} \mid D \in \mathcal{D}^*\} \subseteq \mathcal{B}$$

is an uncountable Δ-system with root $R = R^* \cup \{s\}$.

Otherwise, for each $s \in S$, there exist only countably many sets $B \in \mathcal{B}$ such that $s \in B$. In this case, we shall show that $\mathcal{B}$ contains an uncountable subfamily $\mathcal{D} = \{B_\beta \mid \beta < \omega_1\}$ of pairwise disjoint sets. Of course, this means that $\mathcal{D}$ is an uncountable Δ-system with root $R = \emptyset$. Suppose that $\alpha < \omega_1$ and that we have defined the set $B_\beta \in \mathcal{B}$ for each $\beta < \alpha$. Let $S_\alpha = \bigcup\{B_\beta \mid \beta < \alpha\}$. Then S_α is a countable set and hence there exist only countably many $B \in \mathcal{B}$ such that $B \cap S_\alpha \neq \emptyset$. Consequently, there exists $B_\alpha \in \mathcal{B}$ such that $B_\alpha \cap S_\alpha = \emptyset$. □

Theorem 6.3.10. *If J is a countable set, then* $\mathrm{Fn}(I, J)$ *has the c.c.c.*

Proof. Suppose that $\mathcal{A} = \{p_\alpha \mid \alpha < \omega_1\}$ is an uncountable antichain in $\mathrm{Fn}(I, J)$. Since J is countable, if $F \subset I$ is a finite subset, then there exist only countably many functions $f : F \to J$. Hence, after passing to a suitable subset of $\mathcal{A}$ if necessary, we can suppose that if $\alpha < \beta < \omega_1$, then $\operatorname{dom} p_\alpha \neq \operatorname{dom} p_\beta$. Applying Theorem 6.3.9 to the family $\{\operatorname{dom} p_\alpha \mid \alpha < \omega_1\}$, there exists a subset $X \subseteq \omega_1$ with $|X| = \omega_1$ and a finite subset $R \subset I$ such that $\{\operatorname{dom} p_\alpha \mid \alpha \in X\}$ forms a Δ-system with root R. Since there are only countably many possibilities for $p_\alpha \restriction R$, there exists a subset $Y \subseteq X$ with $|Y| = \omega_1$ and a fixed function $f : R \to J$ such that $p_\alpha \restriction R = f$ for all $\alpha \in Y$. But this means that the conditions $\{p_\alpha \mid \alpha \in Y\}$ are pairwise compatible, which contradicts the hypothesis that $\mathcal{A}$ is an antichain. □

Hence if $\mathbb{Q} = \mathrm{Fn}(\omega_2 \times \omega, 2)^M \in M$, then $M \vDash \mathbb{Q}$ has the *c.c.c.* Hence if G is a $\mathbb{P}$-generic filter over M, then $\omega_2^M = \omega_2^{M[G]}$ and so $M[G] \vDash 2^\omega \geq \omega_2$. This completes the proof of the consistency of $ZFC + \neg CH$.

Unfortunately, many useful notions of forcing fail to have the *c.c.c.* For example, consider the following variant of $\mathrm{Fn}(I, J)$.

Definition 6.3.11. If κ is an infinite cardinal and I, J are any sets, then $\mathrm{Fn}(I, J, \kappa)$ is the notion of forcing consisting of all functions p such that

(a) $\operatorname{dom} p \subseteq I$,
(b) $\operatorname{ran} p \subseteq J$, and
(c) $|p| < \kappa$,

ordered by $q \leq p$ iff $q \supseteq p$.

If I is infinite, $|J| \geq 2$ and $\kappa > \omega$, then $\mathrm{Fn}(I, J, \kappa)$ contains uncountable antichains. However, we still have the following analogues of Theorems 6.3.9 and 6.3.10.

THEOREM 6.3.12 (The Δ-System Lemma). *Let κ be an infinite cardinal. Let $\theta > \kappa$ be a regular cardinal such that $\lambda^{<\kappa} < \theta$ for all cardinals $\lambda < \theta$. If $\mathcal{B}$ is a family of sets such that $|\mathcal{B}| \geq \theta$ and $|B| < \kappa$ for all $B \in \mathcal{B}$, then there exists a Δ-system $\mathcal{D} \subseteq \mathcal{B}$ such that $|\mathcal{D}| = \theta$.*

PROOF. This is Theorem II.1.6 of Kunen [**26**]. $\square$

In practice, we shall usually use the following special case of the Δ-System Lemma.

COROLLARY 6.3.13. *Let κ be an infinite cardinal such that $\kappa^{<\kappa} = \kappa$. If $\mathcal{B}$ is a family of sets such that $|\mathcal{B}| = \kappa^+$ and $|B| < \kappa$ for all $B \in \mathcal{B}$, then there exists a Δ-system $\mathcal{D} \subseteq \mathcal{B}$ such that $|\mathcal{D}| = \kappa^+$.*

$\square$

THEOREM 6.3.14. *If κ is an infinite cardinal and I, J are any sets, then $\mathrm{Fn}(I, J, \kappa)$ has the $(|J|^{<\kappa})^+$-c.c.*

PROOF. This is Lemma VII.6.10 of Kunen [**26**]. $\square$

In particular, suppose that $\kappa \in M$ is a regular uncountable cardinal such that $M \vDash 2^{<\kappa} = \kappa$. Let $\mathbb{P} = \mathrm{Fn}(\kappa^{++} \times \kappa, 2, \kappa)^M \in M$. By Theorem 6.3.14, $M \vDash \mathbb{P}$ has the κ^+-c.c. Hence, by Theorem 6.3.6, $\mathbb{P}$ preserves cofinalities and cardinals greater or equal to κ^+. As we shall soon see, $\mathbb{P}$ also preserves cardinals less than or equal to κ. In the remainder of this section, we shall discuss the relevant combinatorial property of $\mathbb{P}$.

DEFINITION 6.3.15. Suppose that $\mathbb{P} \in M$ is a notion of forcing and that $\kappa \in M$ is an infinite cardinal.

(a) $\mathbb{P}$ *preserves cardinals less than or equal to κ* if whenever G is a $\mathbb{P}$-generic filter over M and $\theta \in M$ is a cardinal such that $\theta \leq \kappa$, then θ remains a cardinal in $M[G]$.

(b) $\mathbb{P}$ *preserves cofinalities less than or equal to* κ if whenever G is a $\mathbb{P}$-generic filter over M and $\gamma \in M$ is a limit ordinal with $\mathrm{cf}^M(\gamma) \leq \kappa$, then $\mathrm{cf}^{M[G]}(\gamma) = \mathrm{cf}^M(\gamma)$.

The proof of the following lemma is essentially identical to that of Lemma 6.3.3.

LEMMA 6.3.16. *Suppose that* $\mathbb{P} \in M$ *is a notion of forcing and that* $\kappa \in M$ *is an infinite cardinal.*

(a) *If* $\mathbb{P}$ *preserves cofinalities less than or equal to* κ*, then* $\mathbb{P}$ *also preserves cardinals less than or equal to* κ*.*
(b) *If* $\mathbb{P}$ *does not preserve cofinalities less than or equal to* κ*, then there exists a regular cardinal* $\theta \in M$ *with* $\theta \leq \kappa$ *and a* $\mathbb{P}$*-generic filter* G *over* M *such that* $M[G] \vDash \theta$ *is not a regular cardinal.*

□

DEFINITION 6.3.17. Let $\mathbb{P}$ be a notion of forcing and let κ be an infinite cardinal. Then $\mathbb{P}$ is *κ-closed* iff whenever $\gamma < \kappa$ and $\langle p_\xi \mid \xi < \gamma \rangle$ is a decreasing sequence of elements of $\mathbb{P}$, then there exists an element $q \in \mathbb{P}$ such that $q \leq p_\xi$ for all $\xi < \gamma$.

LEMMA 6.3.18. *If* κ *is a regular cardinal and* I*,* J *are any sets, then* $\mathrm{Fn}(I, J, \kappa)$ *is* κ*-closed.*

PROOF. Suppose that $\gamma < \kappa$ and that $\langle p_\xi \mid \xi < \gamma \rangle$ is a decreasing sequence of elements of $\mathrm{Fn}(I, J, \kappa)$. Then $q = \bigcup_{\xi<\gamma} p_\xi \in \mathrm{Fn}(I, J, \kappa)$ and $q \leq p_\xi$ for all $\xi < \gamma$. □

THEOREM 6.3.19. *Suppose that* $\mathbb{P} \in M$ *is a notion of forcing and that* $\kappa \in M$ *is an infinite cardinal. If* $M \vDash \mathbb{P}$ *is* κ*-closed, then* $\mathbb{P}$ *preserves cofinalities and cardinals less than or equal to* κ*.*

As we shall soon see, Theorem 6.3.19 is an easy consequence of the following important result.

LEMMA 6.3.20. *Suppose that* $\mathbb{P} \in M$ *is a notion of forcing,* $\kappa \in M$ *is an infinite cardinal and that* $M \vDash \mathbb{P}$ *is* κ*-closed. Let* G *be a* $\mathbb{P}$*-generic filter over* M*. Suppose* A*,* $B \in M$ *with* $|A| < \kappa$ *and that* $f \in M[G]$ *with* $f : A \to B$*. Then* $f \in M$*.*

Proof. Let $\tau \in M$ be a $\mathbb{P}$-name such that $\tau_G = f$. Then there exists $p \in G$ such that

$$p \Vdash \tau \text{ is a function from } \check{A} \text{ into } \check{B}.$$

Let $E \in M$ be the set consisting of the conditions $q \in \mathbb{P}$ such that there exists a function $g \in {}^AB \cap M$ with $q \Vdash \tau = \check{g}$. By Lemma 6.3.1, it is enough to prove that E is dense below p.

Let $q \in \mathbb{P}$ with $q \leq p$. For the rest of this proof, we shall work within M. Let $|A| = \lambda < \kappa$ and let $\{a_\alpha \mid \alpha < \lambda\}$ be an enumeration of the elements of A. We shall define sequences $\langle q_\alpha \mid \alpha \leq \lambda \rangle$ and $\langle b_\alpha \mid \alpha < \lambda \rangle$ of elements of $\mathbb{P}$, B respectively by transfinite recursion such that the following conditions are satisfied.

(a) $q_0 = q$.

(b) If $\beta < \alpha \leq \lambda$, then $q_\alpha \leq q_\beta$.

(c) If $\alpha < \lambda$, then $q_{\alpha+1} \Vdash \tau(\check{a}_\alpha) = \check{b}_\alpha$.

First suppose that $\alpha \leq \lambda$ is a limit ordinal. Since $\mathbb{P}$ is κ-closed, there exists $q_\alpha \in \mathbb{P}$ such that $q_\alpha \leq q_\beta$ for all $\beta < \alpha$. Now suppose that $\alpha = \beta + 1$. Since $q_\beta \leq p$, it follows that

$$q_\beta \Vdash (\exists x \in \check{B})\tau(\check{a}_\beta) = x.$$

Hence there exist $q_\alpha \leq q_\beta$ and $b_\beta \in B$ such that $q_\alpha \Vdash \tau(\check{a}_\beta) = \check{b}_\beta$. Finally let $g \in {}^AB$ be the function defined by $g(a_\alpha) = b_\alpha$ for all $\alpha < \lambda$. Then $q_\lambda \Vdash \tau = \check{g}$. □

Of course, if $\mathbb{P} \in M$ is an arbitrary notion of forcing, then $\mathbb{P}$ is trivially ω-closed. Hence if A, $B \in M$ and A is finite, then $({}^AB)^M = ({}^AB)^{M[G]}$ for every $\mathbb{P}$-generic filter G over M.

Proof of Theorem 6.3.19. Suppose not. Then by Lemma 6.3.16, there exists a regular cardinal $\theta \in M$ with $\theta \leq \kappa$ and a $\mathbb{P}$-generic filter G over M such that $M[G] \vDash \theta$ is not a regular cardinal. Thus there exists an ordinal $\alpha < \theta$ and a function $f \in M[G]$ such that f maps α cofinally into θ. By Lemma 6.3.20, $f \in M$, contradicting the assumption that θ is a regular cardinal in M. □

Combining Theorems 6.3.6 and 6.3.19, we obtain the following important result.

Theorem 6.3.21. *Suppose that $\mathbb{P} \in M$ is a notion of forcing and that $\kappa \in M$ is an infinite cardinal. If $M \vDash \mathbb{P}$ is κ-closed and has the κ^+-c.c., then $\mathbb{P}$ preserves cofinalities and cardinals.*

$\square$

6.4. Nice $\mathbb{P}$-names

In this section, M will denote a c.t.m. of ZFC.

Let $\mathbb{Q} = \mathrm{Fn}(\omega_2 \times \omega, 2)^M \in M$ and let G be a $\mathbb{Q}$-generic filter over M. Then we have already seen that $M[G] \vDash 2^\omega \geq \omega_2$. In this section, we shall study the problem of computing the exact value of 2^ω within $M[G]$. There is an obvious approach to finding an upper bound for 2^ω in $M[G]$; namely, we should calculate the number of $\mathbb{Q}$-names $\tau \in M$ such that $\tau_G \subseteq \omega$. Unfortunately, this initial approach does not give any useful information, since the collection of such $\mathbb{Q}$-names is a proper subclass of M. To see this, let $q \in \mathbb{Q}$ be any condition such that $q \notin G$. Then for every $x \in M$, $\tau_x = \{\langle \check{x}, q \rangle\} \in M$ is a $\mathbb{Q}$-name such that $(\tau_x)_G = \emptyset$. Of course, if $x \notin \omega$, then τ_x seems a perverse choice for a $\mathbb{Q}$-name for a subset of ω. The next three results show that the above stategy works if we restrict our attention to more natural $\mathbb{Q}$-names for subsets of ω.

DEFINITION 6.4.1. Let $\mathbb{P}$ be a notion of forcing and let B be any set. Then τ is a *nice $\mathbb{P}$-name for a subset of B* iff τ has the form

$$\bigcup \{\{\check{b}\} \times A_b \mid b \in B\},$$

where each A_b is an antichain in $\mathbb{P}$.

LEMMA 6.4.2. *Suppose that $\mathbb{P} \in M$ is a notion of forcing and that G is a $\mathbb{P}$-generic filter over M.*

(a) *If $B \in M$ and $C \in M[G]$ with $C \subseteq B$, then there exists a nice $\mathbb{P}$-name $\tau \in M$ for a subset of B such that $\tau_G = C$.*

(b) *Suppose that λ, $\theta \in M$ are cardinals and that*

$$M \vDash \theta \text{ is the number of nice } \mathbb{P}\text{-names for subsets of } \lambda.$$

Then $M[G] \vDash 2^\lambda \leq |\theta|$.

PROOF. (a) This is a special case of Lemma VII.5.12 of Kunen [**26**].

(b) If $\langle \tau_\alpha \mid \alpha < \theta \rangle \in M$ is an enumeration of the nice $\mathbb{P}$-names for subsets of λ, then $\langle (\tau_\alpha)_G \mid \alpha < \theta \rangle \in M[G]$ is an enumeration with repetitions of $\mathcal{P}(\lambda)^{M[G]}$. $\square$

Lemma 6.4.3. *If $\mathbb{P}$ has the κ-c.c., then for every set B, there exist at most $(|\mathbb{P}|^{<\kappa})^{|B|}$ nice $\mathbb{P}$-names for subsets of B.*

Proof. Let $\mathbb{A}$ be the set of antichains of $\mathbb{P}$. Since $\mathbb{P}$ has the κ-c.c., it follows that $|\mathbb{A}| \leq |\mathbb{P}|^{<\kappa}$. Each nice $\mathbb{P}$-names for a subset of B is essentially just a function from B to $\mathbb{A}$ and there are at most $(|\mathbb{P}|^{<\kappa})^{|B|}$ such functions. □

Theorem 6.4.4. *Suppose that κ, λ, $\theta \in M$ are infinite cardinals such that $\kappa^{<\kappa} = \kappa \leq \lambda < \theta$ and $\theta^\lambda = \lambda$. Let $\mathbb{P} = \mathrm{Fn}(\theta \times \lambda, 2, \kappa)^M \in M$ and let G be a $\mathbb{P}$-generic filter over M.*

(a) *$M \vDash \mathbb{P}$ is κ-closed and has the κ^+-c.c.*
(b) *$\mathbb{P}$ preserves cofinalities and cardinals.*
(c) *$M[G] \vDash 2^\lambda = \theta$.*

Proof. Throughout this proof, we shall work within the c.t.m. M. First note that since $\kappa^{<\kappa} = \kappa$, it follows that κ is a regular cardinal and hence $\mathbb{P}$ is κ-closed. It also follows that $2^{<\kappa} = \kappa$ and so Theorem 6.3.14 implies that $\mathbb{P}$ has the κ^+-c.c. Consequently, by Theorem 6.3.21, $\mathbb{P}$ preserves cofinalities and cardinals. To see that $M[G] \vDash 2^\lambda \geq \theta$, let $g = \bigcup G \in M[G]$ and for each $\alpha < \theta$, let $g_\alpha : \lambda \to 2$ be the function defined by $g_\alpha(\xi) = g(\alpha, \xi)$. By considering the appropriate dense subsets of $\mathbb{P}$, it is easily checked that $g_\alpha \neq g_\beta$ for all $\alpha < \beta < \theta$. Finally to see that $M[G] \vDash 2^\lambda \leq \theta$, it is enough to show that there are at most θ nice $\mathbb{P}$-names for subsets of λ. Since there are only $\theta^{<\kappa}$ possibilities for the domain D of an element p of $\mathbb{P}$ and at most $2^{<\kappa}$ possibilities for p once the domain D is fixed, it follows that $|\mathbb{P}| = \theta^{<\kappa} \cdot 2^{<\kappa} = \theta$. Hence, by Lemma 6.4.3, the number of nice $\mathbb{P}$-names for subsets of λ is at most $(|\mathbb{P}|^{<\kappa})^\lambda = \theta$. □

By König's Theorem, if λ is an infinite cardinal, then $\mathrm{cf}(2^\lambda) > \lambda$. Theorem 6.4.4 implies that this is the only constraint on 2^λ that can be proved in ZFC. To see this, let $M \vDash GCH$. If $\lambda, \theta \in M$ are infinite cardinals, then $\theta^\lambda = \theta$ iff $\mathrm{cf}(\theta) > \lambda$. Hence, in this case, if

$$\mathbb{P} = \mathrm{Fn}(\theta \times \lambda, 2, \omega)^M = \mathrm{Fn}(\theta \times \lambda, 2)^M \in M$$

and if G be a $\mathbb{P}$-generic filter over M, then $M[G] \vDash 2^\lambda = \theta$. It is easily checked that if $\kappa \in M$ is any cardinal such that $\omega \leq \kappa < \lambda$, then we also have that $M[G] \vDash 2^\kappa = \theta$.

If $\lambda \in M$ is a regular cardinal, then the following result shows that we can force simultaneously that 2^λ is arbitrarily large and also that λ is the first cardinal which violates GCH.

COROLLARY 6.4.5. *Suppose that $M \vDash GCH$ and that λ, $\theta \in M$ are infinite cardinals such that λ is regular and* $\mathrm{cf}(\theta) > \lambda$. *Let $\mathbb{P} = \mathrm{Fn}(\theta \times \lambda, 2, \lambda)^M \in M$ and let G be a $\mathbb{P}$-generic filter over M.*

(a) *$\mathbb{P}$ preserves cofinalities and cardinals.*
(b) *$M[G] \vDash 2^\lambda = \theta$.*
(c) *$M[G] \vDash 2^\mu = \mu^+$ for every cardinal μ such that $\omega \leq \mu < \lambda$.*

PROOF. Once again, we shall work within the c.t.m. M. Using GCH, since λ is regular and $\mathrm{cf}(\theta) > \lambda$, it follows that $\lambda^{<\lambda} = \lambda$ and $\theta^\lambda = \lambda$. Hence (a) and (b) are immediate consequences of Theorem 6.4.4. Finally suppose that μ is a cardinal such that $\omega \leq \mu < \lambda$. By Lemma 6.3.20, since $\mathbb{P}$ is λ-closed, it follows that $({}^\mu 2)^{M[G]} = ({}^\mu 2)^M$. Hence $M[G] \vDash 2^\mu = \mu^+$. □

The situation is much more interesting when λ is a singular cardinal. In 1974, confounding all expectations, Silver [**46**] proved that GCH cannot first fail at a singular cardinal of uncountable cofinality. Assuming the consistency of suitable large cardinals, Magidor [**29**] had already shown that it was consistent that GCH first fails at $\aleph_\omega$. Later, again assuming the consistency of suitable large cardinals, Gitik and Magidor [**12**] proved that if α is any countably infinite ordinal such that $\alpha > \omega$, then it is consistent that GCH first fails at $\aleph_\omega$ and that $2^{\aleph_\omega} = \aleph_{\alpha+1}$. On the other hand, Shelah [**44**] has used PCF theory to prove the remarkable result that if $2^{\aleph_n} < \aleph_\omega$ for all $n < \omega$, then $2^{\aleph_\omega} < \aleph_{\omega_4}$. A well-written survey of this fascinating area can be found in Jech [**21**].

6.5. Some observations and conventions

In this section, we shall make a number of simple but useful observations; and we shall discuss some of the conventions that will be used in the remaining sections of this book.

Suppose that we wish to establish the consistency of $ZFC + \sigma$, where σ is a sentence in the first-order language of set theory. Then it is enough to find a c.t.m.

M and a notion of forcing $\mathbb{P} \in M$ for which there exists a $\mathbb{P}$-generic filter G over M such that $M[G] \vDash \sigma$. In practice, we usually find a notion of forcing $\mathbb{P} \in M$ such that $M[G] \vDash \sigma$ for every $\mathbb{P}$-generic filter G over M. For this reason, we are usually not concerned about which particular $\mathbb{P}$-generic filter is chosen.

Convention 6.5.1. If $\mathbb{P} \in M$ is a notion of forcing, then we often denote the corresponding generic extension by $M^{\mathbb{P}}$ if we do not wish to specify a particular $\mathbb{P}$-generic filter over M.

We have already defined what it means for a condition $p \in \mathbb{P}$ to force a statement to be true in the generic extension $M^{\mathbb{P}}$. In the later chapters of this book, it will be useful to have the notion of a condition $p \in \mathbb{P}$ deciding a statement in $M^{\mathbb{P}}$.

Definition 6.5.2. Suppose that $\varphi(x_1, \dots, x_n)$ is a formula with free variables $x_1, \dots, x_n$ and that $\tau_1, \dots, \tau_n \in M$ are $\mathbb{P}$-names. If $p \in \mathbb{P}$, then we say that p *decides* $\varphi(\tau_1, \dots, \tau_n)$ iff $p \Vdash \varphi(\tau_1, \dots, \tau_n)$ or $p \Vdash \neg\varphi(\tau_1, \dots, \tau_n)$.

In Section 6.2, we observed that

$$D = \{p \in \mathbb{P} \mid p \text{ decides } \varphi(\tau_1, \dots, \tau_n)\} \in M$$

is a dense subset of $\mathbb{P}$. It will also be useful have the notion of a condition $p \in \mathbb{P}$ deciding a function or relation in $M^{\mathbb{P}}$.

Definition 6.5.3. Suppose that A, $B \in M$ and that $\tilde{f}$, $\tilde{R} \in M$ are $\mathbb{P}$-names for which there exist conditions q, $r \in \mathbb{P}$ such that

$$q \Vdash \tilde{f} \text{ is a function from } \check{A} \text{ to } \check{B}$$

and

$$r \Vdash \tilde{R} \text{ is an } n\text{-ary relation on } \check{A}$$

Then we say that p *decides* $\tilde{f}$ iff $p \leq q$ and there exists a function $g \in M$ with $g : A \to B$ such that $p \Vdash \tilde{f} = \check{g}$. Similarly, we say that p *decides* $\tilde{R}$ iff $p \leq r$ and there exists an n-ary relation $S \in M$ on A such that $p \Vdash \tilde{R} = \check{S}$.

For example, suppose also that $\mathbb{P} \in M$ is κ-closed and that $A \in M$ satisfies $|A| < \kappa$. Then the proof of Lemma 6.3.20 shows that

$$D = \{p \in \mathbb{P} \mid p \text{ decides } \tilde{f}\}$$

is dense below q; and a similar argument shows that

$$E = \{p \in \mathbb{P} \mid p \text{ decides } \tilde{R}\}$$

is dense below r.

Next suppose that $M \vDash GCH$ and that $\lambda, \theta \in M$ are infinite cardinals such that λ is regular and $\mathrm{cf}(\theta) > \lambda$. In Section 6.4, we saw that if $\mathbb{P} = \mathrm{Fn}(\theta \times \lambda, 2, \lambda)^M \in M$, then the following statements are true in $M^{\mathbb{P}}$:

(i) $2^\lambda = \theta$; and

(ii) $2^\mu = \mu^+$ for every cardinal μ such that $\omega \leq \mu < \lambda$.

In Chapter 9, it will be useful to notice that we could just as well have used the slightly simpler notion of forcing $\mathbb{Q} = \mathrm{Fn}(\theta, 2, \lambda)^M \in M$. To see this, let $\varphi \in M$ be a bijection between θ and $\theta \times \lambda$. Then φ induces a corresponding isomorphism $f : \mathbb{P} \to \mathbb{Q}$, defined by

$$f(p) = p \circ \varphi.$$

As the reader would expect, isomorphic notions of forcing give rise to identical generic extensions.

THEOREM 6.5.4. *Suppose that* $\mathbb{P}$, $\mathbb{Q} \in M$ *are notions of forcing and that* $f \in M$ *is an isomorphism from* $\mathbb{P}$ *onto* $\mathbb{Q}$. *If* G *is a* $\mathbb{P}$*-generic filter over* M, *then* $H = f[G] \in M[G]$ *is a* $\mathbb{Q}$*-generic filter over* M *and* $M[G] = M[H]$.

PROOF. It is easily checked that $H = f[G]$ is a $\mathbb{Q}$-generic filter over M. Since $H = f[G] \in M[G]$, it follows that $M[H] \subseteq M[G]$. Similarly, since $G = f^{-1}[H] \in M[H]$, we have that $M[G] \subseteq M[H]$ and so $M[G] = M[H]$. □

The analogous result is also true when $\mathbb{Q}$ is a dense sub-order of $\mathbb{P}$.

THEOREM 6.5.5. *Suppose that* $\mathbb{P}$, $\mathbb{Q} \in M$ *are notions of forcing and that* $\mathbb{P}$ *is a dense sub-order of* $\mathbb{Q}$.

(a) *If* G *is a* $\mathbb{P}$*-generic filter over* M, *then*

$$H = \{q \in \mathbb{Q} \mid (\exists p \in G)(p \leq q)\}$$

is a $\mathbb{Q}$*-generic filter over* M.

(b) *If* H *is a* $\mathbb{Q}$*-generic filter over* M, *then* $G = H \cap \mathbb{P}$ *is a* $\mathbb{P}$*-generic filter over* M.

Furthermore, in both cases, $M[G] = M[H]$.

PROOF. (a) This is completely routine.

(b) It is easily checked that G satisfies the following properties:

(i) For all p, $q \in G$, there exists $r \in \mathbb{P}$ such that $r \leq p$, q.

(ii) For all p, $q \in \mathbb{P}$, if $p \in G$ and $p \leq q$, then $q \in G$.

(iii) If $D \in M$ is a dense subset of $\mathbb{P}$, then $G \cap D \neq \emptyset$.

Thus it only remains to prove that for all p, $q \in G$, there exists $r \in G$ such that $r \leq p$, q. To see this, let p, $q \in G$ and consider

$$D_{p,q} = \{r \in \mathbb{P} \mid r \perp p \text{ or } r \perp q \text{ or } r \leq p, q\}.$$

It is easily checked that $D_{p,q} \in M$ is a dense subset of $\mathbb{P}$ and hence there exists an element $r \in G \cap D_{p,q}$. Since the elements of G are pairwise compatible, it follows that $r \leq p$, q. □

In the early stages of an account of set-theoretic forcing, it is necessary to carefully distinguish between a notion of forcing such as $\mathrm{Fn}(\omega_2 \times \omega, 2)$ and its relativised version $\mathrm{Fn}(\omega_2 \times \omega, 2)^M = \mathrm{Fn}(\omega_2^M \times \omega, 2)$ within the c.t.m. M. For this reason, in the first four sections of this chapter, the symbol V was reserved to stand for the actual set-theoretic universe, and we kept track of which arguments took place within V and which within M. However, after the basic ideas of forcing are understood, this extremely careful approach begins to feel increasingly tedious and clumsy. For example, in order to prove that $\mathrm{Fn}(\omega_2 \times \omega, 2)^M$ preserves cofinalities and cardinalities, we first proved that $\mathrm{Fn}(\omega_2 \times \omega, 2)$ has the *c.c.c.* within V and then deduced that

$$M \vDash \mathrm{Fn}(\omega_2 \times \omega, 2)^M \text{ has the } c.c.c.$$

However, notice that it was not really necessary to conduct any of our argument within V; since $M \vDash ZFC$, our original argument relativised to M shows directly that $\mathrm{Fn}(\omega_2 \times \omega, 2)^M$ has the *c.c.c.* within M.

From now on, V will denote some fixed c.t.m., usually referred to as the *ground model*; and, unless otherwise specified, all of our arguments will take place in V. Furthermore, we shall usually write $\mathrm{Fn}(\omega_2 \times \omega, 2)$, $\mathrm{Sym}(\omega)$, ω_2, etc. for the relativised versions of these objects within V, instead of the more accurate

$\mathrm{Fn}(\omega_2 \times \omega, 2)^V$, $\mathrm{Sym}^V(\omega)$, ω_2^V, etc. (Of course, whenever it seems genuinely helpful, we shall revert to using notation such as $\mathrm{Sym}^V(\omega)$. For example, this occurs in the next section when we compute the automorphism tower of $\mathrm{Sym}^V(\omega)$ within a generic extension M of the ground model V.) In a similar vein, if $\mathbb{P} \in V$ is a notion of forcing, then we shall usually use phrases such as

- "$\mathbb{P}$ has the *c.c.c.*"

instead of the more accurate

- "$V \vDash \mathbb{P}$ has the *c.c.c.*"

In summary, we shall write as though V still continued to denote the actual set-theoretic universe. (As we mentioned earlier, there is only one point in the development of forcing where it is essential that V is a c.t.m. rather than the set-theoretic universe; namely, the proof of the existence of a $\mathbb{P}$-generic filter over V.)

6.6. $\tau(G)$ is not absolute

Throughout this section, V will denote the ground model.

In this section, we shall prove that the height $\tau(G)$ of the automorphism tower of an infinite centreless group G is not necessarily an absolute concept. First we shall present an example of a centreless group $G \in V$ for which there exists a notion of forcing $\mathbb{P} \in V$ such that $\tau^{V^{\mathbb{P}}}(G) > \tau^V(G)$. The idea is very simple: we shall let G be a suitably chosen complete group and let $\mathbb{P}$ be a notion of forcing which adjoins a "new automorphism" $\pi \in \mathrm{Aut}^{V^{\mathbb{P}}}(G) \smallsetminus \mathrm{Aut}^V(G)$. Then clearly π must be an outer automorphism and so $\tau^{V^{\mathbb{P}}}(G) > \tau^V(G) = 0$. Of course, G cannot be an arbitrary complete group. For example, it is impossible to adjoin a new automorphism to a finitely generated group H. (To see this, suppose that H is generated by the finite subset F and let $\mathbb{P}$ be any notion of forcing. If $\pi \in \mathrm{Aut}^{V^{\mathbb{P}}}(G)$, then π is uniquely determined by its restriction $\pi \upharpoonright F$. By the remark following Lemma 6.3.20, $\pi \upharpoonright F \in V$ and hence $\pi \in V$.) For a more interesting example of such a group, consider $G = \mathrm{Sym}(\omega)$. In Section 1.3, we proved that $\mathrm{Sym}(\omega)$ is a complete group. Now let $\mathbb{P} \in V$ be any notion of forcing which adjoins a new permutation to ω; for example, we could take

$$\mathbb{P} = \{p \in \mathrm{Fn}(\omega, \omega) \mid p \text{ is an injection } \}.$$

Then

$$V^{\mathbb{P}} \vDash G = \mathrm{Sym}^V(\omega) \text{ is a proper subgroup of } \mathrm{Sym}(\omega),$$

and so the results of Section 1.3 can no longer be applied to G within $V^{\mathbb{P}}$. However, the following result shows that G remains complete within $V^{\mathbb{P}}$.

THEOREM 6.6.1. *Let $G = Sym^V(\omega)$. If $\mathbb{P} \in V$ be any notion of forcing, then G remains a complete group in the generic extension $V^{\mathbb{P}}$.*

PROOF. It is more convenient to consider the isomorphic group $H = \mathrm{Sym}^V(\mathbb{Z})$. We shall work within the generic extension $M = V^{\mathbb{P}}$. Let $\pi \in \mathrm{Aut}^M(H)$ be an arbitrary automorphism. Since

$$\mathrm{Alt}^M(\mathbb{Z}) = \mathrm{Alt}^V(\mathbb{Z}) \leqslant H \leqslant \mathrm{Sym}^M(\mathbb{Z}),$$

Corollary 4.1.5 implies that there exists a permutation $\varphi \in \mathrm{Sym}^M(\mathbb{Z})$ such that $\pi(h) = \varphi h \varphi^{-1}$ for all $h \in H$. In particular, this is true when $\sigma \in H$ is the permutation such that $\sigma(z) = z + 1$ for all $z \in \mathbb{Z}$. Let $g = \varphi\sigma\varphi^{-1}$. Then $g \in H$ and $g(\varphi(z)) = \varphi(z + 1)$ for all $z \in \mathbb{Z}$. Hence if $n \geq 0$, then $\varphi(n) = g^n(\varphi(0))$ and $\varphi(-n) = (g^{-1})^n(\varphi(0))$. Since $g \in H = \mathrm{Sym}^V(\mathbb{Z})$, it follows that $g \in V$ and hence $\varphi \in V$. Thus π is an inner automorphism of H. □

The situation becomes more interesting if we consider the setwise stabiliser $S_{\mathcal{U}}$ of a nonprincipal ultrafilter over over the set ω of natural numbers.

DEFINITION 6.6.2. A *nonprincipal ultrafilter* over the set ω is a collection $\mathcal{U}$ of subsets of ω satisfying the following conditions:

(i) If $A, B \in \mathcal{U}$, then $A \cap B \in \mathcal{U}$.
(ii) If $A \in \mathcal{U}$ and $A \subseteq B \subseteq \omega$, then $B \in \mathcal{U}$.
(iii) For all $A \subseteq \omega$, either $A \in \mathcal{U}$ or $\omega \smallsetminus A \in \mathcal{U}$.
(iv) If F is a finite subset of ω, then $F \notin \mathcal{U}$.

Equivalently, if $\mu : \mathcal{P}(\omega) \to \{0, 1\}$ is the function such that $\mu(A) = 1$ if and only if $A \in \mathcal{U}$, then μ is a finitely additive probability measure on ω such that $\mu(F) = 0$ for all finite subsets F of ω. It is well-known that there are exactly 2^{2^ω} distinct nonprincipal ultrafilters over ω. For example, see Theorem 56 of Jech [**19**].

DEFINITION 6.6.3. If $\mathcal{U}$ is a nonprincipal ultrafilter filter over ω, then its *setwise stabiliser* is the group

$$S_{\mathcal{U}} = \{\pi \in \mathrm{Sym}(\omega) \mid \text{ For all } X \subseteq \omega, X \in \mathcal{U} \text{ iff } \pi[X] \in \mathcal{U}\}.$$

The next two results collect together some of the basic algebraic properties of the group $S_{\mathcal{U}}$. Recall that if $\pi \in \mathrm{Sym}(\omega)$, then

$$\mathrm{fix}(\pi) = \{n \in \omega \mid \pi(n) = n\}$$

and

$$\mathrm{supp}(\pi) = \{n \in \omega \mid \pi(n) \neq n\}.$$

THEOREM 6.6.4. *Let $\mathcal{U}$ be a nonprincipal ultrafilter on ω.*

(a) $S_{\mathcal{U}} = \{\pi \in \mathrm{Sym}(\omega) \mid \mathrm{fix}(\pi) \in \mathcal{U}\}$.
(b) $S_{\mathcal{U}}$ *is a maximal proper subgroup of* $\mathrm{Sym}(\omega)$.
(c) $S_{\mathcal{U}}$ *is a complete group.*

PROOF. Clauses (a) and (b) are a restatement of Theorem 6.4 of Macpherson and Neumann [**28**]. Clearly $S_{\mathcal{U}}$ contains the subgroup $\mathrm{Fin}(\omega)$ of finite permutations of ω. Hence Corollary 4.1.5 implies that $S_{\mathcal{U}}$ is centreless and that the automorphism tower of $S_{\mathcal{U}}$ coincides with the normaliser tower of $S_{\mathcal{U}}$ in $\mathrm{Sym}(\omega)$. Since $S_{\mathcal{U}}$ is a maximal proper subgroup of $\mathrm{Sym}(\omega)$, it follows that $N_{\mathrm{Sym}(\omega)}(S_{\mathcal{U}}) = S_{\mathcal{U}}$ and hence $S_{\mathcal{U}}$ is a complete group. □

THEOREM 6.6.5. *If $\mathcal{A}$, $\mathcal{B}$ are nonprincipal ultrafilters over ω, then the following are equivalent.*

(a) $S_{\mathcal{A}} \simeq S_{\mathcal{B}}$.
(b) *There exists* $\pi \in \mathrm{Sym}(\omega)$ *such that* $\pi[\mathcal{A}] = \mathcal{B}$.
(c) *There exists* $\pi \in \mathrm{Sym}(\omega)$ *such that* $\pi S_{\mathcal{A}} \pi^{-1} = S_{\mathcal{B}}$.

Hence there exist 2^{2^ω} pairwise nonisomorphic groups of the form $S_{\mathcal{U}}$ for some nonprincipal ultrafilter $\mathcal{U}$ over ω.

PROOF. Clearly that (b) and (c) are equivalent. It is also clear that (c) implies (a). Finally if $\theta : S_{\mathcal{A}} \to S_{\mathcal{B}}$ is an isomorphism, Theorem 4.1.3 implies that there exists $\pi \in \mathrm{Sym}(\omega)$ such that $\theta(g) = \pi g \pi^{-1}$ for all $g \in S_{\mathcal{A}}$. Hence $\pi S_{\mathcal{A}} \pi^{-1} = S_{\mathcal{B}}$. □

Now suppose that $\mathcal{U} \in V$ is an ultrafilter and let

$$G = S^V_{\mathcal{U}} = \{\pi \in \mathrm{Sym}^V(\omega) \mid \text{ For all } X \in \mathcal{P}^V(\omega), X \in \mathcal{U} \text{ iff } \pi[X] \in \mathcal{U}\}.$$

By Theorem 6.6.4(c), $V \vDash G$ is a complete group. Our next target is to show that there exists a *c.c.c.* notion of forcing $\mathbb{P} \in V$ which adjoins an outer automorphism of $S^V_{\mathcal{U}}$. Our argument is based upon the following simple observations.

Definition 6.6.6. If A, $B \subseteq \omega$, then $A \subseteq^* B$ iff $|A \smallsetminus B| < \omega$. In this case, we say that A is *almost contained in* B.

Lemma 6.6.7. *If $\mathcal{U}$ is a nonprincipal ultrafilter filter over ω, then there does not exist an infinite subset $T \subseteq \omega$ such that $T \subseteq^* A$ for all $A \in \mathcal{U}$.*

Proof. Suppose that such a set T exists. If $T \notin \mathcal{U}$, then $\omega \smallsetminus T \in \mathcal{U}$ and hence $T \subseteq^* \omega \smallsetminus T$, which is a contradiction. Hence $T \in \mathcal{U}$. Express $T = X \sqcup Y$ as the disjoint union of two infinite subsets. Since $T \in \mathcal{U}$, it follows that either $X \in \mathcal{U}$ or $Y \in \mathcal{U}$. But this implies that either $T \subseteq^* X$ or $T \subseteq^* Y$, which is also a contradiction. □

Lemma 6.6.8. *Let $\mathcal{U} \in V$ be a nonprincipal ultrafilter filter over ω. Suppose that the notion of forcing $\mathbb{P} \in V$ adjoins an infinite subset $T \subseteq \omega$ such that $T \subseteq^* A$ for all $A \in \mathcal{U}$. Then $\mathbb{P}$ adjoins an outer automorphism of $S^V_{\mathcal{U}}$.*

Proof. Let $G = S^V_{\mathcal{U}}$. Throughout this proof, we shall work within the generic extension $M = V^{\mathbb{P}}$. (Of course, $\mathcal{U}$ will no longer be an ultrafilter in M.) Let $T \subseteq \omega$ be an infinite subset such that $T \subseteq^* A$ for all $A \in \mathcal{U}$ and let $\varphi \in \mathrm{Sym}^M(\omega)$ be any permutation such that $\mathrm{supp}(\varphi) = T$. By Lemma 6.6.7, $\varphi \notin V$ and so $\varphi \notin G$. We shall show that φ normalises G in $\mathrm{Sym}^M(\omega)$. To see this, let $\pi \in G$ be an arbitrary element. By Theorem 6.6.4(a), $\mathrm{fix}(\pi) \in \mathcal{U}$. Since $T \subseteq^* \mathrm{fix}(\pi)$ and

$$\mathrm{supp}(\varphi) \cap \mathrm{supp}(\pi) = \mathrm{supp}(\varphi) \smallsetminus \mathrm{fix}(\pi) = T \smallsetminus \mathrm{fix}(\pi),$$

it follows that $|\,\mathrm{supp}(\varphi) \cap \mathrm{supp}(\pi)| < \omega$. Hence $\varphi\pi\varphi^{-1}\pi^{-1} \in \mathrm{Fin}(\omega)$ and so

$$\varphi\pi\varphi^{-1} \in \mathrm{Fin}(\omega)\pi \subseteq G.$$

Since φ normalises G and $\varphi \notin G$, it follows that φ induces an outer automorphism of G via conjugation. □

DEFINITION 6.6.9. Let $\mathcal{U}$ be a nonprincipal ultrafilter on ω. Then the associated *Mathias notion of forcing* $\mathbb{P}_{\mathcal{U}}$ consists of the conditions $p = \langle s, A \rangle$, where $s : n \to \omega$ is a strictly increasing function for some $n < \omega$ and $A \in \mathcal{U}$. The ordering on $\mathbb{P}_{\mathcal{U}}$ is given by $\langle t, B \rangle \leq \langle s, A \rangle$ iff

(i) $t \supseteq s$ and $B \subseteq A$; and
(ii) $t(\ell) \in A$ for all $\ell \in \operatorname{dom} t \smallsetminus \operatorname{dom} s$.

The Mathias notion of forcing $\mathbb{P}_{\mathcal{U}}$ is designed to generically adjoin an infinite subset $T \subseteq \omega$ such that $T \subseteq^* A$ for all $A \in \mathcal{U}$. Intuitively, the condition $p = \langle s, A \rangle \in \mathbb{P}_{\mathcal{U}}$ consists of

(i) a finite approximation $s : n \to \omega$ to the increasing enumeration of the desired set $T = \{m_\ell \mid \ell < \omega\}$; together with
(ii) a "promise" that $m_\ell \in A$ for all $\ell \geq n$.

LEMMA 6.6.10. *$\mathbb{P}_{\mathcal{U}}$ has the c.c.c.*

PROOF. Suppose that $\{p_\alpha \mid \alpha < \omega_1\}$ is an uncountable subset of $\mathbb{P}_{\mathcal{U}}$; say, $p_\alpha = \langle s_\alpha, A_\alpha \rangle$ for each $\alpha < \omega_1$. Since there are only countably many possibilities for s_α, there exists a fixed function s and an uncountable subset $I \subseteq \omega_1$ such that $s_\alpha = s$ for all $\alpha \in I$. If α, $\beta \in I$, then $\langle s, A_\alpha \cap A_\beta \rangle \leq p_\alpha, p_\beta$. Hence $\mathbb{P}_{\mathcal{U}}$ does not contain any uncountable antichains. □

LEMMA 6.6.11. *For each $A \in \mathcal{U}$,*

$$D_A = \{\langle s, B \rangle \mid B \subseteq A\}$$

is a dense subset of $\mathbb{P}_{\mathcal{U}}$.

PROOF. If $p = \langle s, C \rangle$ is any element of $\mathbb{P}_{\mathcal{U}}$, then $q = \langle s, A \cap C \rangle \in D_A$ and $q \leq p$. □

LEMMA 6.6.12. *If $n < \omega$, then*

$$D_n = \{\langle s, A \rangle \mid n \subseteq \operatorname{dom} s\}$$

is a dense subset of $\mathbb{P}_{\mathcal{U}}$.

PROOF. Let $p = \langle s, A\rangle$ be any element of $\mathbb{P}_{\mathcal{U}}$ and let $m = \max\{n, \operatorname{dom} s\}$. Since A is an infinite subset of ω, we can extend s to a strictly increasing function $t : m \to \omega$ such that $t(\ell) \in A$ for all $\ell \in \operatorname{dom} t \smallsetminus \operatorname{dom} s$. Then $q = \langle t, A\rangle \in D_n$ and $q \leq p$. □

THEOREM 6.6.13. *Let $\mathcal{U} \in V$ be a nonprincipal ultrafilter over ω and let $\mathbb{P}_{\mathcal{U}}$ be the associated Mathias notion of forcing. Then*

(a) *$\mathbb{P}_{\mathcal{U}}$ has the c.c.c.; and*
(b) *in $V^{\mathbb{P}_{\mathcal{U}}}$, there exists an infinite subset $T \subseteq \omega$ such that $T \subseteq^* A$ for all $A \in \mathcal{U}$.*

PROOF. Let $V^{\mathbb{P}_{\mathcal{U}}} = V[G]$, where G is a $\mathbb{P}_{\mathcal{U}}$-generic filter over V, and define $f \in V[G]$ by

$$f = \bigcup \{s \mid \text{ There exists } A \in \mathcal{U} \text{ such that } \langle s, A\rangle \in G\}.$$

Then Lemma 6.6.12 implies that $f \in {}^{\omega}\omega$ and clearly f is strictly increasing. Let $T = f[\omega] \in V[G]$. We shall show that $T \subseteq^* A$ for all $A \in \mathcal{U}$. To see this, fix some $A \in \mathcal{U}$. Applying Lemma 6.6.11, let $p = \langle s, B\rangle \in G$ be a condition such that $B \subseteq A$. Let $\operatorname{dom} s = n$ and let m be any integer such that $m \geq n$. Then there exists a condition $q = \langle t, C\rangle \in G$ such that $q \leq p$ and $m \in \operatorname{dom} t$. This implies that $f(m) = t(m) \in B$. Thus

$$\{f(m) \mid n \leq m < \omega\} \subseteq B \subseteq A$$

and so $T \subseteq^* A$. □

This completes the proof of the following result.

THEOREM 6.6.14. *If $\mathcal{U} \in V$ is a nonprincipal ultrafilter over ω and $G = S_{\mathcal{U}}^V$, then*

(a) *G is a complete group; and*
(b) *there exists a c.c.c. notion of forcing $\mathbb{P}$ which adjoins an outer automorphism of G.*

In particular, $\tau^{V^{\mathbb{P}}}(G) > \tau^V(G)$.

□

Perhaps more surprisingly, in the remainder of this section, we shall present an example of a centreless group $G \in V$ for which there exists a notion of forcing $\mathbb{P}$ such that $\tau^{V^{\mathbb{P}}}(G) < \tau^V(G)$.

DEFINITION 6.6.15. The notion of forcing $\mathbb{P}$ is said to *adjoin a new real* if there exists an element $R \in \mathcal{P}^{V^{\mathbb{P}}}(\omega) \smallsetminus \mathcal{P}^V(\omega)$.

Of course, because of the natural bijections between $\mathcal{P}(\omega)$, $\mathrm{Sym}(\omega)$, the field $\mathbb{R}$ of real numbers, etc., it follows that if $\mathbb{P}$ is a notion of forcing, then $\mathcal{P}^{V^{\mathbb{P}}}(\omega) \neq \mathcal{P}^V(\omega)$ iff $\mathrm{Sym}^{V^{\mathbb{P}}}(\omega) \neq \mathrm{Sym}^V(\omega)$ iff $\mathbb{R}^{V^{\mathbb{P}}} \neq \mathbb{R}^V$, etc.

THEOREM 6.6.16. *There exists a centreless group $G \in V$ of cardinality 2^ω which satisfies the following conditions.*

(a) $\tau^V(G) = 2$.
(b) *If $\mathbb{P}$ is any notion of forcing which adjoins a new real, then $\tau^{V^{\mathbb{P}}}(G) = 1$.*

The proof of Theorem 6.6.16 makes use of the following two results.

LEMMA 6.6.17. *Let $\mathbb{P}$ be a notion of forcing which adjoins a new real and let $M = V^{\mathbb{P}}$ be the corresponding generic extension.*

(a) $\mathrm{Sym}^V(\omega)$ *is self-normalising in* $\mathrm{Sym}^M(\omega)$.
(b) $\mathrm{Sym}^V(\omega) \not\simeq \mathrm{Sym}^M(\omega)$ *in the actual set theoretic universe.*

PROOF. (a) If $\varphi \in \mathrm{Sym}^M(\omega)$ normalises $\mathrm{Sym}^V(\omega)$, then φ induces an automorphism π of $\mathrm{Sym}^V(\omega)$ via conjugation. By Theorem 6.6.1, π is an inner automorphism and hence $\varphi \in \mathrm{Sym}^V(\omega)$.

(b) In this part of the proof, we shall work within the actual set-theoretic universe. Suppose that $\theta : \mathrm{Sym}^V(\omega) \to \mathrm{Sym}^M(\omega)$ is an isomorphism. By Theorem 4.1.3, there exists a permutation φ of ω such that $\theta(g) = \varphi g \varphi^{-1}$ for all $g \in \mathrm{Sym}^V(\omega)$. Arguing as in the proof of Theorem 6.6.1, we see that $\varphi \in M$. But then we have that $\varphi \in \mathrm{Sym}^M(\omega)$ and that $\varphi\, \mathrm{Sym}^V(\omega)\varphi^{-1} = \mathrm{Sym}^M(\omega)$, which is impossible since $\mathrm{Sym}^V(\omega)$ is a proper subgroup of $\mathrm{Sym}^M(\omega)$. □

Of course, Lemma 6.6.17(b) implies that $M \vDash \mathrm{Sym}^V(\omega) \not\simeq \mathrm{Sym}^M(\omega)$.

The following result is simply the observation that the construction of Fried and Kollár in the proof of Theorem 4.1.7 is upwards absolute.

THEOREM 6.6.18. *Let $\Gamma = \langle X, E\rangle \in V$ be any graph. Then there exists a field $K_\Gamma \in V$ of cardinality* $\max\{|X|, \omega\}$ *such that whenever $\mathbb{P}$ is a (possibly trivial) notion of forcing and $M = V^{\mathbb{P}}$, then the following conditions are satisfied.*

(a) *X is an $\mathrm{Aut}^M(K_\Gamma)$-invariant subset of K_Γ.*
(b) *The restriction mapping, $\pi \mapsto \pi \restriction X$, is an isomorphism from $\mathrm{Aut}^M(K_\Gamma)$ onto $\mathrm{Aut}^M(\Gamma)$.*

□

Here a notion of forcing $\mathbb{P}$ is said to be *trivial* if $|\mathbb{P}| = 1$. Of course, in this case, we have that $V^{\mathbb{P}} = V$.

PROOF OF THEOREM 6.6.16. Let $\Gamma = \langle X_1 \sqcup X_2, E\rangle \in V$ be the complete bipartite graph such that

(i) $|X_1| = |X_2| = \omega$; and
(ii) if $v \in X_i$ and $w \in X_j$, then v and w are adjacent iff $i \neq j$.

Let $K_\Gamma \in V$ be the field given by Theorem 6.6.18. If $\mathbb{P}$ is a (possibly trivial) notion of forcing, then

$$\mathrm{Aut}^{V^{\mathbb{P}}}(\Gamma) = [\mathrm{Sym}^{V^{\mathbb{P}}}(X_1) \times \mathrm{Sym}^{V^{\mathbb{P}}}(X_2)] \rtimes \langle \sigma \rangle,$$

where $\sigma \in V$ is any involution which interchanges the null subgraphs X_1 and X_2. For rest of this proof, we shall identify $\mathrm{Aut}^{V^{\mathbb{P}}}(K_\Gamma)$ with $\mathrm{Aut}^{V^{\mathbb{P}}}(\Gamma)$. Let $H = \mathrm{Sym}^V(X_1)$ and let $G = PGL(2, K_\Gamma) \rtimes H$. (Notice that if $K \in V$ is any field, then $PGL(2, K)$ is absolute for V. In particular, it is not necessary to write $PGL^V(2, K_\Gamma)$ or $PGL^{V^{\mathbb{P}}}(2, K_\Gamma)$.) By Theorems 4.1.6 and 6.6.18, whenever $\mathbb{P}$ is a (possibly trivial) notion of forcing, then for each ordinal α,

$$G_\alpha^{V^{\mathbb{P}}} = PGL(2, K_\Gamma) \rtimes N_\alpha^{V^{\mathbb{P}}}(H),$$

where $N_\alpha^{V^{\mathbb{P}}}(H)$ is the αth group in the normaliser tower of H in $\mathrm{Aut}^{V^{\mathbb{P}}}(\Gamma)$. Clearly in V, the normaliser tower of H in $\mathrm{Aut}^V(\Gamma)$ is given by

$$H = \mathrm{Sym}^V(X_1) \lhd \mathrm{Sym}^V(X_1) \times \mathrm{Sym}^V(X_2) \lhd \mathrm{Aut}^V(\Gamma).$$

Hence $\tau^V(G) = 2$.

Now let $\mathbb{P}$ be any notion of forcing which adjoins a real and let $M = V^{\mathbb{P}}$. By Lemma 6.6.17(a), $H = \mathrm{Sym}^V(X_1)$ is self-normalising in $\mathrm{Sym}^M(X_1)$. Hence the normaliser of H in $\mathrm{Aut}^M(\Gamma)$ is $N = \mathrm{Sym}^V(X_1) \times \mathrm{Sym}^M(X_2)$. Lemma 6.6.17(b) implies that N is self-normalising in $\mathrm{Aut}^M(\Gamma)$. Hence $\tau^M(G) = 1$. $\square$

In this section, we have seen that the height of the automorphism tower of an infinite centreless group may either increase or decrease in a generic extension. In Chapter 8, we shall prove that it is consistent that for every infinite cardinal $\lambda \in V$ and every ordinal $\alpha < \lambda$, there exists a centreless group $G \in V$ with the following properties.

(a) $\tau^V(G) = \alpha$.
(b) If β is any ordinal such that $1 \leq \beta < \lambda$, then there exists a notion of forcing $\mathbb{P} \in V$, which preserves cofinalities and cardinals, such that $\tau^{V^{\mathbb{P}}}(G) = \beta$.

6.7. An absoluteness theorem for automorphism towers

Throughout this section, V will denote the ground model.

In this section, we shall prove the following absoluteness theorem for automorphism towers, which will be used repeatedly in the later chapters of this book.

THEOREM 6.7.1. *Let $\mathbb{P} \in V$ be a κ^+-closed notion of forcing and let $M = V^{\mathbb{P}}$ be the corresponding generic extension. If $G \in V$ is a centreless group of cardinality κ, then $G^M_\alpha = G^V_\alpha$ for all $\alpha \geq 0$. Hence $\tau^M(G) = \tau^V(G)$.*

We shall prove that $G^M_\alpha = G^V_\alpha$ by induction on $\alpha \geq 0$. Clearly the result holds when $\alpha = 0$, since

$$G^M_0 = G = G^V_0.$$

In order to prove the result when $\alpha = 1$, we must show that

$$G^M_1 = \mathrm{Aut}^M(G) = \mathrm{Aut}^V(G) = G^V_1.$$

It is clear that $\mathrm{Aut}^V(G) \leqslant \mathrm{Aut}^M(G)$. On the other hand, since $\mathbb{P} \in V$ is κ^+-closed and $|G| = \kappa$, Lemma 6.3.20 implies that if $\pi \in \mathrm{Aut}^M(G)$, then $\pi \in V$. Hence $\mathrm{Aut}^M(G) \leqslant \mathrm{Aut}^V(G)$. Unfortunately, this argument is not enough to prove that

$$G^M_2 = \mathrm{Aut}^M(G^V_1) = \mathrm{Aut}^V(G^V_1) = G^V_2,$$

since it is possible that $|G^V_1| \geq \kappa^+$. However, note that Lemma 6.3.20 implies that $\pi \restriction G \in V$ for every $\pi \in \mathrm{Aut}^M(G^V_1)$. Hence the following lemma implies that

$\mathrm{Aut}^M(G_1^V) = \mathrm{Aut}^V(G_1^V)$. Continuing in this fashion, we shall obtain Theorem 6.7.1. (Lemma 6.7.2 will also play a crucial role in Chapter 9.)

LEMMA 6.7.2. *Let $\mathbb{P} \in V$ be a notion of forcing and let $M = V^{\mathbb{P}}$ be the corresponding generic extension. Suppose that $G \in V$ is a centreless group and that $\pi \in \mathrm{Aut}^M(G_\alpha^V)$ for some $\alpha \geq 0$. If $\pi \restriction G \in V$, then $\pi \in V$.*

PROOF. For each $\beta \leq \alpha$, let $\pi_\beta = \pi \restriction G_\beta^V$. Working inside V, we shall prove inductively that $\pi_\beta \in V$ for all $\beta \leq \alpha$. In particular, $\pi = \pi_\alpha \in V$. By assumption, $\pi_0 = \pi \restriction G \in V$. Now suppose that $\beta > 0$ and that $\pi_\gamma \in V$ for all $\gamma < \beta$. First suppose that $\beta = \gamma + 1$ is a successor ordinal.

CLAIM 6.7.3. *If $g \in G_{\gamma+1}^V$, then $\pi(g)$ is the unique element $h \in N_{G_\alpha^V}(\pi_\gamma[G_\gamma^V])$ such that $h\pi_\gamma(a)h^{-1} = \pi_\gamma(gag^{-1})$ for all $a \in G_\gamma^V$.*

PROOF OF CLAIM 6.7.3. Let $g \in G_{\gamma+1}^V$. Then g normalises G_γ^V and for all $a \in G_\gamma^V$,

$$\pi(g)\pi_\gamma(a)\pi(g)^{-1} = \pi(gag^{-1}) = \pi_\gamma(gag^{-1}).$$

Since Theorem 1.1.10 implies that

$$C_{G_\alpha^V}(\pi_\gamma[G_\gamma^V]) = C_{G_\alpha^V}(\pi[G_\gamma^V]) = \pi[C_{G_\alpha^V}(G_\gamma^V)] = 1,$$

it follows that $\pi(g)$ is the unique such element. □

By Claim 6.7.3, $\pi_{\gamma+1} = \pi \restriction G_{\gamma+1}^V$ is explicitly definable from G_α^V and π_γ. Since G_α^V, $\pi_\gamma \in V$, it follows that $\pi_{\gamma+1} \in V$. Finally if β is a limit ordinal, then it is clear that $\pi_\beta = \bigcup_{\gamma<\beta} \pi_\gamma \in V$. □

PROOF OF THEOREM 6.7.1. We shall argue by induction on α that $G_\alpha^M = G_\alpha^V$. This is clearly true when $\alpha = 0$ and no difficulties arise when α is a limit ordinal. Assume inductively that $G_\alpha^M = G_\alpha^V$ for some $\alpha \geq 0$. Then it is enough to show that $\mathrm{Aut}^M(G_\alpha^V) = \mathrm{Aut}^V(G_\alpha^V)$. Clearly $\mathrm{Aut}^V(G_\alpha^V) \leqslant \mathrm{Aut}^M(G_\alpha^V)$. On the other hand, since $\mathbb{P} \in V$ is κ^+-closed and $|G| = \kappa$, Lemma 6.3.20 implies that $\pi \restriction G \in V$ for each $\pi \in \mathrm{Aut}^M(G_\alpha^V)$. Hence Lemma 6.7.2 implies that $\mathrm{Aut}^M(G_\alpha^V) \leqslant \mathrm{Aut}^V(G_\alpha^V)$. □

6.8. Iterated forcing

Throughout this section, V will denote the ground model.

In this section, we shall discuss some of the basic ideas of iterated forcing, including the notion of a reverse Easton iterated forcing. (Clear accounts of the theory of iterated forcing can be found in Baumgartner [**2**], Jech [**20**] and Kunen [**26**].)

To prove the consistency of $\tau_{\omega_1} < \tau_{\omega_2}$, it is enough to find a notion of forcing $\mathbb{P}$ which adjoins a centreless group G of cardinality ω_2 such that $\tau(G) \geq (2^{\omega_1})^+$; and this can be done using the techniques developed in the earlier sections of this chapter. However, these techniques are often not enough to prove the consistency of statements involving the existence of groups of arbitrarily large cardinalities. For example, in Section 7.4, we shall prove that it is consistent that τ_κ is a strictly increasing function of κ. To accomplish this, we shall begin with a ground model V satisfying GCH and we shall then construct a generic extension M in which the following statements are true:

(a) GCH holds.
(b) For each regular uncountable cardinal κ, there exists a centreless group T of cardinality κ such that $\tau(T) = \kappa^+$.

It follows easily that τ_κ is strictly increasing in M. To see this, working within M, suppose that θ, λ are infinite cardinals such that $\theta < \lambda$. Since GCH holds, it follows that $\tau_\theta < \left(2^\theta\right)^+ = \theta^{++}$. Let κ be a regular (necessarily uncountable) cardinal such that $\theta < \kappa \leq \lambda$. Then there exists a centreless group T of cardinality κ such that $\tau(T) = \kappa^+ \geq \theta^{++}$ and so $\tau_\kappa > \tau_\theta$. Applying Corollary 3.2.2, we obtain that $\tau_\theta < \tau_\kappa \leq \tau_\lambda$.

In order to construct M, it seems natural to proceed as follows. First we should force with a notion of forcing $\mathbb{P}_0 \in V_0 = V$ which adjoins a centreless group G_0 of cardinality ω_1 such that $\tau(G_0) = \omega_2$. Next we should force with a notion of forcing $\mathbb{P}_1 \in V_1 = V_0^{\mathbb{P}_0}$ which adjoins a centreless group G_1 of cardinality ω_2 such that $\tau(G_0) = \omega_3$. Continuing in this fashion, for each $n < \omega$, we should construct an n-stage iterated forcing

$$V = V_0 \subset V_1 = V_0^{\mathbb{P}_0} \subset \cdots \subset V_{i+1} = V_i^{\mathbb{P}_i} \subset \cdots \subset V_n = V_{n-1}^{\mathbb{P}_{n-1}},$$

where each $\mathbb{P}_i \in V_i$ adjoins a centreless group G_i of cardinality ω_{i+1} such that $\tau(G_i) = \omega_{i+2}$. However, this approach runs into difficulties at stage ω, as it turns out that $V_\omega = \bigcup_{n<\omega} V_n$ is not a model of ZFC. (Since a new subset of $\aleph_\omega$ is adjoined at every step of the iteration, it follows that

$$V_\omega \nvDash (\exists x)(\forall y)(y \in x \leftrightarrow y \subseteq \aleph_\omega).$$

Hence the Powerset Axiom fails in V_ω.) For this reason, it is necessary to take a completely different approach to iterated forcing.

We shall begin by reconsidering a typical 2-stage iteration

$$V = V_0 \subset V_1 = V_0^{\mathbb{P}_0} \subset V_2 = V_1^{\mathbb{P}_1} = (V_0^{\mathbb{P}_0})^{\mathbb{P}_1},$$

where $\mathbb{P}_0 \in V_0$ and $\mathbb{P}_1 \in V_1$. Our first target is to find a single notion of forcing $\mathbb{Q} \in V_0$ such that

$$(V_0^{\mathbb{P}_0})^{\mathbb{P}_1} = V_0^{\mathbb{Q}}.$$

If $\mathbb{P}_1 \in V_0$, then it is relatively straightforward to show that $\mathbb{Q} = \mathbb{P}_0 \times \mathbb{P}_1$ satisfies our requirements. (For example, see Section VIII.1 of Kunen [**26**].) Unfortunately, in most examples, we have that $\mathbb{P}_1 \in V_1 \smallsetminus V_0$. However, even in this case, V_0 always contains a $\mathbb{P}_0$-name for the notion of forcing $\mathbb{P}_1$. In more detail, let $V_0^{\mathbb{P}_0} = V_0[G_0]$, where G_0 is a $\mathbb{P}_0$-generic filter over V_0. Then there exists a $\mathbb{P}_0$-name τ and a condition $p \in G_0$ such that $\tau_{G_0} = \mathbb{P}_1$ and

$$p \Vdash \tau \text{ is a partially ordered set with greatest element.}$$

In fact, as explained in Section VIII.5 of Kunen [**26**], the $\mathbb{P}_0$-name τ can be chosen so that $p = 1_{\mathbb{P}_0}$.

Definition 6.8.1. Suppose that $\mathbb{P}$ is a notion of forcing. Then a $\mathbb{P}$-*name for a notion of forcing* is a triple of $\mathbb{P}$-names $\langle \tilde{\mathbb{Q}}, \leq_{\tilde{\mathbb{Q}}}, 1_{\tilde{\mathbb{Q}}} \rangle$ such that

$$1_{\mathbb{P}} \Vdash \langle \tilde{\mathbb{Q}}, \leq_{\tilde{\mathbb{Q}}} \rangle \text{ is a partial order with largest element } 1_{\tilde{\mathbb{Q}}}.$$

Slightly abusing notation, we shall usually say that $\tilde{\mathbb{Q}}$ is a $\mathbb{P}$-name for a notion of forcing.

Definition 6.8.2. Suppose that $\mathbb{P}$ is a notion of forcing and that $\tilde{\mathbb{Q}}$ is a $\mathbb{P}$-name for a notion of forcing. Then $\mathbb{P} * \tilde{\mathbb{Q}}$ is the notion of forcing with underlying set

$$\{\langle p, \tilde{q}\rangle \mid p \in \mathbb{P} \text{ and } 1_{\mathbb{P}} \Vdash \tilde{q} \in \tilde{\mathbb{Q}}\},$$

partially ordered by

$$\langle p_1, \tilde{q}_1\rangle \leq \langle p_2, \tilde{q}_2\rangle \quad \text{iff} \quad p_1 \leq_{\mathbb{P}} p_2 \text{ and } p_1 \Vdash \tilde{q}_1 \leq_{\tilde{\mathbb{Q}}} \tilde{q}_2.$$

(Strictly speaking, we should identify those conditions $\langle p_1, \tilde{q}_1\rangle$ and $\langle p_2, \tilde{q}_2\rangle$ such that $\langle p_1, \tilde{q}_1\rangle \leq \langle p_2, \tilde{q}_2\rangle$ and $\langle p_2, \tilde{q}_2\rangle \leq \langle p_1, \tilde{q}_1\rangle$. Of course, this occurs iff $p_1 = p_2$ and $p_1 \Vdash \tilde{q}_1 = \tilde{q}_2$.)

The next result shows that $\mathbb{P} * \tilde{\mathbb{Q}}$ satisfies our requirements; i.e. that forcing with $\mathbb{P} * \tilde{\mathbb{Q}}$ is equivalent to first forcing with $\mathbb{P}$ and then forcing with $\tilde{\mathbb{Q}}_G$ over $V^{\mathbb{P}} = V[G]$. (While the notion of forcing $\mathbb{P} * \tilde{\mathbb{Q}}$ was primarily introduced in order to be able to define iterated forcing constructions of transfinite length, it should be mentioned that $\mathbb{P} * \tilde{\mathbb{Q}}$ often plays a crucial role in understanding the properties of 2-stage iterations. For example, see the proof of Theorem 8.3.1.)

Theorem 6.8.3. *Let $\mathbb{P} \in V$ be a notion of forcing and let $\tilde{\mathbb{Q}} \in V$ be a $\mathbb{P}$-name for a notion of forcing.*

(a) *Suppose that G is a $\mathbb{P}$-generic filter over V and that H is a $\tilde{\mathbb{Q}}_G$-generic filter over $V[G]$. Then*

$$G * H = \{\langle p, \tilde{q}\rangle \in \mathbb{P} * \tilde{\mathbb{Q}} \mid p \in G \text{ and } \tilde{q}_G \in H\}$$

is a $\mathbb{P} * \tilde{\mathbb{Q}}$-generic filter over V and $V[G * H] = V[G][H]$.

(b) *Conversely, suppose that K is a $\mathbb{P} * \tilde{\mathbb{Q}}$-generic filter over V. Then*

$$G = \{p \in \mathbb{P} \mid \text{ There exists } \tilde{q} \text{ such that } \langle p, \tilde{q}\rangle \in K\}$$

is a $\mathbb{P}$-generic filter over V,

$$H = \{\tilde{q}_G \mid \text{ There exists } p \text{ such that } \langle p, \tilde{q}\rangle \in K\}$$

is a $\tilde{\mathbb{Q}}_G$-generic filter over $V[G]$ and $K = G * H$.

Proof. This is Theorem 1.1 of Baumgartner [**2**]. □

The next result confirms that $\mathbb{P} * \tilde{\mathbb{Q}}$ satisfies the expected chain conditions, etc.

THEOREM 6.8.4. *Let $\mathbb{P}$ be a notion of forcing and let $\tilde{\mathbb{Q}}$ be a $\mathbb{P}$-name for a notion of forcing. Let λ be a regular uncountable cardinal and let θ be a cardinal such that $\theta^\mu = \theta$ for all $\mu < \lambda$.*

(a) *If $\mathbb{P}$ has the λ-c.c. and $1_\mathbb{P} \Vdash \tilde{\mathbb{Q}}$ has the λ-c.c. , then $\mathbb{P} * \tilde{\mathbb{Q}}$ has the λ-c.c.*
(b) *If $\mathbb{P}$ is λ-closed and $1_\mathbb{P} \Vdash \tilde{\mathbb{Q}}$ is λ-closed , then $\mathbb{P} * \tilde{\mathbb{Q}}$ is λ-closed.*
(c) *Suppose that $\mathbb{P}$ has the λ-c.c. and that $|\mathbb{P}| \leq \theta$. If $1_\mathbb{P} \Vdash |\tilde{\mathbb{Q}}| \leq \theta$, then $|\mathbb{P} * \tilde{\mathbb{Q}}| \leq \theta$.*

PROOF. This combines Theorem 2.1, Corollary 2.6 and Lemma 3.2 of Baumgartner [**2**]. □

It is now relatively straightforward to define iterated forcing constructions of arbitrary transfinite length.

DEFINITION 6.8.5. Let $\alpha \geq 0$. Then an *α-stage iteration* is a sequence of forcing notions $\langle \mathbb{P}_\beta \mid \beta \leq \alpha \rangle$ which satisfies the following conditions.

(a) If $\beta \leq \alpha$ and $p \in \mathbb{P}_\beta$, then $p = \langle p(\gamma) \mid \gamma < \beta \rangle$ is a β-sequence.
(b) If $\beta = 0$, then $\mathbb{P}_0 = \{\emptyset\}$ is the trivial notion of forcing.
(c) If $\beta = \gamma + 1$, then there exists a $\mathbb{P}_\gamma$-name $\tilde{\mathbb{Q}}_\gamma$ for a notion of forcing such that
  (i) $p \in \mathbb{P}_\beta$ iff $p \restriction \gamma \in \mathbb{P}_\gamma$ and $1_{\mathbb{P}_\gamma} \Vdash_{\mathbb{P}_\gamma} p(\gamma) \in \tilde{\mathbb{Q}}_\gamma$; and
  (ii) $p \leq_{\mathbb{P}_\beta} q$ iff $p \restriction \gamma \leq_{\mathbb{P}_\gamma} q \restriction \gamma$ and $p \restriction \gamma \Vdash_{\mathbb{P}_\gamma} p(\gamma) \leq_{\mathbb{P}_\gamma} q(\gamma)$.
  Thus $\mathbb{P}_\beta \simeq \mathbb{P}_\gamma * \tilde{\mathbb{Q}}_\gamma$.
(c) If β is a limit ordinal, then
  (i) if $p \in \mathbb{P}_\beta$, then $p \restriction \gamma \in \mathbb{P}_\gamma$ for all $\gamma < \beta$; and
  (ii) $p \leq_{\mathbb{P}_\beta} q$ iff $p \restriction \gamma \leq_{\mathbb{P}_\gamma} q \restriction \gamma$ for all $\gamma < \beta$.
  Furthermore, $\mathbb{P}_\beta$ satisfies the following closure properties:
  (iii) $1_{\mathbb{P}_\beta} = \langle 1_{\tilde{\mathbb{Q}}_\gamma} \mid \gamma < \beta \rangle \in \mathbb{P}_\beta$; and
  (iv) if $\gamma < \beta$, $p \in \mathbb{P}_\beta$, $q \in \mathbb{P}_\gamma$ and $q \leq_{\mathbb{P}_\gamma} p \restriction \gamma$, then $r \in \mathbb{P}_\beta$, where $r \restriction \gamma = q$ and $r(\xi) = p(\xi)$ for all $\gamma \leq \xi < \beta$.

If β is a limit ordinal, then clause 6.8.5(c) is not enough to completely determine the notion of forcing $\mathbb{P}_\beta$. While there are other possibilities, in this book, we shall always take $\mathbb{P}_\beta$ to be either the direct limit or the inverse limit of $\langle \mathbb{P}_\gamma \mid \gamma < \beta \rangle$.

DEFINITION 6.8.6. Suppose that β is a limit ordinal and that $\langle \mathbb{P}_\gamma \mid \gamma \leq \beta \rangle$ is a β-stage iteration.

(a) $\mathbb{P}_\beta$ is the *direct limit* of $\langle \mathbb{P}_\gamma \mid \gamma < \beta \rangle$ if $p \in \mathbb{P}_\beta$ iff there exists $\gamma < \beta$ such that $p \upharpoonright \gamma \in \mathbb{P}_\gamma$ and $p(\xi) = 1_{\tilde{\mathbb{Q}}_\xi}$ for all $\gamma \leq \xi < \beta$.

(b) $\mathbb{P}_\beta$ is the *inverse limit* of $\langle \mathbb{P}_\gamma \mid \gamma < \beta \rangle$ if $p \in \mathbb{P}_\beta$ iff $p \upharpoonright \gamma \in \mathbb{P}_\gamma$ for all $\gamma < \beta$.

EXAMPLE 6.8.7. Suppose that the α-stage iteration $\langle \mathbb{P}_\beta \mid \beta \leq \alpha \rangle$ satisfies the following conditions.

(a) If β is a limit ordinal such that $\mathrm{cf}(\beta) > \omega$, then $\mathbb{P}_\beta$ is the direct limit of $\langle \mathbb{P}_\gamma \mid \gamma < \beta \rangle$.

(b) If β is a limit ordinal such that $\mathrm{cf}(\beta) = \omega$, then $\mathbb{P}_\beta$ is the inverse limit of $\langle \mathbb{P}_\gamma \mid \gamma < \beta \rangle$.

Then for each $p \in \mathbb{P}_\alpha$, the set $\{\xi < \alpha \mid p(\xi) \neq 1_{\tilde{\mathbb{Q}}_\xi}\}$ is countable. In this case, we say that $\langle \mathbb{P}_\beta \mid \beta \leq \alpha \rangle$ is a *countable support iteration*.

If $\langle \mathbb{P}_\beta \mid \beta \leq \alpha \rangle$ is an α-stage iteration and $\beta < \alpha$, then we shall often identify each condition $p \in \mathbb{P}_\beta$ with the corresponding condition $p' \in \mathbb{P}_\alpha$ defined by

$$p'(\gamma) = \begin{cases} p(\gamma), & \text{if } \gamma < \beta; \\ 1_{\tilde{\mathbb{Q}}_\gamma}, & \text{if } \beta \leq \gamma < \alpha. \end{cases}$$

With this convention, we have then that

$$\mathbb{P}_0 \subseteq \mathbb{P}_1 \subseteq \cdots \subseteq \mathbb{P}_\beta \subseteq \cdots \subseteq \mathbb{P}_\alpha.$$

Furthermore, if $\beta \leq \alpha$ is a limit ordinal, then $\mathbb{P}_\beta$ is the direct limit of $\langle \mathbb{P}_\gamma \mid \gamma \leq \beta \rangle$ iff $\mathbb{P}_\beta = \bigcup_{\gamma < \beta} \mathbb{P}_\gamma$.

Suppose that $\langle \mathbb{P}_\beta \mid \beta \leq \alpha \rangle$ is an α-stage iteration and that $\beta < \alpha$. Then the basic idea is that forcing with $\mathbb{P}_\alpha$ should be equivalent to first forcing with $\mathbb{P}_\beta$ and then later forcing with the notions of forcing corresponding to $\tilde{\mathbb{Q}}_\gamma$ for $\beta \leq \gamma < \alpha$. To state this result precisely, for each $p \in \mathbb{P}_\alpha$, let

$$p^\beta = \langle p(\gamma) \mid \beta \leq \gamma < \alpha \rangle,$$

so that $p = (p \upharpoonright \beta)^\frown p^\beta$; and let

$$\mathbb{P}_{\beta\alpha} = \{p^\beta \mid p \in \mathbb{P}_\alpha\}.$$

If G_β is a $\mathbb{P}_\beta$-generic filter over V, then we can define a partial ordering $\leq_{\mathbb{P}_{\beta\alpha},G_\beta}$ of $\mathbb{P}_{\beta\alpha}$ in $V[G_\beta]$ by

$$r \leq_{\mathbb{P}_{\beta\alpha},G_\beta} s \quad \text{iff} \quad (\exists p \in G_\beta) p^\frown r \leq_{\mathbb{P}_\alpha} p^\frown s.$$

Let $\tilde{\mathbb{P}}_{\beta\alpha}$ be the canonically chosen $\mathbb{P}_\beta$-name for the notion of forcing

$$\langle \mathbb{P}_{\beta\alpha}, \leq_{\mathbb{P}_{\beta\alpha},G_\beta} \rangle \in V[G_\beta].$$

THEOREM 6.8.8. *With the above notation, $\mathbb{P}_\alpha$ is isomorphic to a dense sub-order of $\mathbb{P}_\beta * \tilde{\mathbb{P}}_{\beta\alpha}$. Furthermore, if G_α is a $\mathbb{P}_\alpha$-generic filter over V, then*

$$G_\beta = \{p \restriction \beta \mid p \in G_\alpha\}$$

is a $\mathbb{P}_\beta$-generic filter over V.

PROOF. This combines Theorems 1.2 and 5.1 of Baumgartner [**2**]. $\square$

Hence, by Theorem 6.5.5, forcing with $\mathbb{P}_\alpha$ is equivalent to first forcing with $\mathbb{P}_\beta$ and then with the notion of forcing

$$\langle \mathbb{P}_{\beta\alpha}, \leq_{\mathbb{P}_{\beta\alpha},G_\beta} \rangle \in V^{\mathbb{P}_\beta} = V[G_\beta].$$

Also notice that if we identify $\mathbb{P}_\beta$ with the corresponding sub-order of $\mathbb{P}_\alpha$, then $G_\beta = G_\alpha \cap \mathbb{P}_\beta$.

As we mentioned earlier, in Section 7.4, beginning with a ground model V satisfying GCH, we shall construct a generic extension M in which the following statements are true:

(a) GCH holds.

(b) For each regular uncountable cardinal κ, there exists a centreless group T of cardinality κ such that $\tau(T) = \kappa^+$.

Here the generic extension M will have the form $V^{\mathbb{P}_\infty}$, where $\mathbb{P}_\infty$ is a notion of forcing which iteratively adjoins a suitable group T for each regular uncountable cardinal κ. In other words, $\mathbb{P}_\infty$ is the limit of an iteration $\langle \mathbb{P}_\beta \mid \beta \in On \rangle$ along the entire class On of ordinals of V. In particular, $\mathbb{P}_\infty$ will be a proper class of V. In the remainder of this section, we shall discuss some of the basic properties of proper class forcing.

Let V be a c.t.m. of ZFC. Recall that $C \subseteq V$ is said to be a *class* of V iff there exists a formula $\varphi(x, y_1, \cdots, y_n)$ with free variables x, $y_1, \cdots, y_n$ and elements $a_1, \cdots, a_n \in V$ such that

$$C = \{c \in V \mid V \vDash \varphi(c, a_1, \cdots, a_n)\}.$$

In particular, every set $a \in V$ is a class of V. The class C is said to be a *proper class* of V iff $C \notin V$. Notice that, from the viewpoint of the actual set-theoretic universe, V has only countably many classes.

Now suppose that $\mathbb{P}$, $\leq$ are classes of V such that $\langle \mathbb{P}, \leq \rangle$ is a partial order with greatest element 1. Then a filter $G \subseteq \mathbb{P}$ is said to be $\mathbb{P}$-generic over V iff $G \cap D \neq \emptyset$ whenever D is a class of V which is dense in $\mathbb{P}$. Since V really has only countably many classes, the proof of Lemma 6.2.6 shows that for each $p \in \mathbb{P}$, there exists a $\mathbb{P}$-generic filter G over V such that $p \in G$. We can now define the notion of a $\mathbb{P}$-name $\tau \in V$ and its interpretation τ_G exactly as in Definitions 6.2.9 and 6.2.10. Finally if G is a $\mathbb{P}$-generic filter G over V, then we once again define the corresponding generic extension by

$$V[G] = \{\tau_G \mid \tau \in V \text{ is a } \mathbb{P}\text{-name }\}.$$

While much of the theory of forcing generalises in a routine fashion to the broader context of class forcing, there are some significant differences. For example, the obvious candidate for a $\mathbb{P}$-name for G is

$$\Gamma = \{\langle \check{p}, p \rangle \mid p \in \mathbb{P}\}.$$

However, if $\mathbb{P}$ is a proper class of V, then Γ is also a proper class of V and so $\Gamma \notin V$. In fact, $G \notin V[G]$ for most proper class forcing notions. A much more serious difference is that $V[G]$ may not be a model of ZFC. For example, the proper class forcing notion $\mathbb{P} = \mathrm{Fn}(On \times \omega, 2)$ adjoins a proper class of subsets of ω and hence the Powerset Axiom fails in the corresponding generic extension $V^{\mathbb{P}}$.

In this book, we shall only need to consider some very well-behaved proper class forcing notions, known as *reverse Easton iterations*. (This name is especially misleading: reverse Easton iterated forcing was introduced by Silver and the iteration proceeds in the usual direction rather than in the reverse direction. Clear introductions to reverse Easton forcing can be found in Baumgartner [**2**] and Menas

[**32**].) For example, in Chapters 7 and 8, we shall make use of proper class iterated forcing notions which fit into the following framework.

HYPOTHESIS 6.8.9. V is a c.t.m. satisfying $ZFC + GCH$ and $\langle \mathbb{P}_\beta \mid \beta \in On \rangle$ is a sequence of forcing notions satisfying the following conditions.

(1) If $\beta = 0$, then $\mathbb{P}_0 = \{\emptyset\}$ is the trivial notion of forcing.
(2) If β is a limit ordinal which is not inaccessible, then $\mathbb{P}_\beta$ is the inverse limit of $\langle \mathbb{P}_\gamma \mid \gamma < \beta \rangle$.
(3) If β is an inaccessible cardinal, then $\mathbb{P}_\beta$ is the direct limit of $\langle \mathbb{P}_\gamma \mid \gamma < \beta \rangle$.
(4) If $\beta = \gamma + 1$ is a successor ordinal, then exists a possibly trivial notion of forcing $\mathbb{Q}_\gamma \in V^{\mathbb{P}_\gamma}$ such that $\mathbb{P}_\beta \simeq \mathbb{P}_\gamma * \widetilde{\mathbb{Q}}_\gamma$. (Here $\widetilde{\mathbb{Q}}_\gamma$ denotes a $\mathbb{P}_\gamma$-name of the notion of forcing $\mathbb{Q}_\gamma \in V^{\mathbb{P}_\gamma}$.) Furthermore, $\mathbb{Q}_\gamma$ is chosen so that the following conditions hold.
    (a) If γ is not a regular uncountable cardinal, then $\mathbb{Q}_\gamma = \mathbb{P}_0$ is the trivial notion of forcing.
    (b) If $\gamma = \kappa$ is a regular uncountable cardinal, then $\mathbb{Q}_\gamma \in V^{\mathbb{P}_\gamma}$ is a notion of forcing of cardinality at most κ^+ such that

$$V^{\mathbb{P}_\kappa} \models \mathbb{Q}_\kappa \text{ is } \kappa\text{-closed and has the } \kappa^+\text{-}c.c.$$

Let $\mathbb{P}_\infty$ be the direct limit of $\langle \mathbb{P}_\beta \mid \beta \in On \rangle$; and for each $\beta \in On$, let $\tilde{\mathbb{P}}_{\beta\infty}$ be the canonically chosen $\mathbb{P}_\beta$-name for a proper class notion of forcing such that $\mathbb{P}_\infty$ is isomorphic to a dense sub-order of $\mathbb{P}_\beta * \tilde{\mathbb{P}}_{\beta\infty}$. As usual, if $\beta \in On$, then we identify $\mathbb{P}_\beta$ with the corresponding sub-order of $\mathbb{P}_\infty$.

Let G be a $\mathbb{P}_\infty$-generic filter over V and let $V[G]$ be the corresponding generic extension. For each $\beta \in On$, let $G_\beta = G \cap \mathbb{P}_\beta$ and let $\mathbb{P}_{\beta\infty} = (\tilde{\mathbb{P}}_{\beta\infty})_{G_\beta} \in V[G_\beta]$.

THEOREM 6.8.10. *With the above hypotheses, the following hold.*

(a) $\mathbb{P}_\infty$ *preserves cofinalities and cardinals.*
(b) *For all* $\beta \in On$, $V[G_\beta]$ *is a model of* $ZFC + GCH$.
(c) *If* κ *is a regular uncountable cardinal, then* $V[G_{\kappa+1}] \models \mathbb{P}_{\kappa+1\infty}$ *is* κ^+*-closed.*
(d) $V[G] = \bigcup_{\beta \in On} V[G_\beta]$.
(e) $V[G]$ *is a model of* $ZFC + GCH$.

PROOF. This follows easily from Menas [**32**]. $\square$

In particular, if the proper class iteration $\langle \mathbb{P}_\beta \mid \beta \in On \rangle$ satisfies Hypothesis 6.8.9, then $V^{\mathbb{P}_\beta} \vDash GCH$ for all $\beta \in On$. Consequently, in the successor stages $\beta = \gamma + 1$ of the construction of such an iteration, we can assume inductively that $V^{\mathbb{P}_\gamma} \vDash GCH$. (Of course, any reader who is already familiar with the theory of iterated forcing can simply perform this induction himself. This remark is only included for those readers who do not yet fit into this category.)

6.9. Notes

The account of set-theoretic forcing in Sections 6.1 to 6.5 is closely based on Kunen's textbook [**26**]. The results in Sections 6.6 and 6.7 first appeared in Thomas [**50**]. The accounts of iterated forcing and reverse Easton forcing in Section 6.8 follow Baumgartner [**2**] and Menas [**32**].

CHAPTER 7

Forcing Long Automorphism Towers

In Chapter 3, we proved that if G is a centreless group of infinite cardinality κ, then the automorphism tower of G terminates after less than $(2^\kappa)^+$ steps; and we also pointed out that $\tau_\kappa < (2^\kappa)^+$, so that $(2^\kappa)^+$ is not the best possible upper bound for $\tau(G)$. The main result of this chapter says that if $\kappa > \omega$, then it is impossible in ZFC to prove a better upper bound for $\tau(G)$ than $(2^\kappa)^+$. (It remains an open question whether a better upper bound can be found in the case when $\kappa = \omega$.) During the proof of our main result, we shall prove two other theorems which are of some interest in their own right.

In Section 7.2, we shall study the question of which groups G can be embedded in the quotient group $\mathrm{Sym}(\kappa)/\mathrm{Sym}_\kappa(\kappa)$. Of course, such a group G must satisfy $|G| \leq 2^\kappa$. But this is not really a restriction on the structure of G, since 2^κ can be made arbitrarily large in generic extensions of the ground model V. In Section 7.2, we shall prove that if κ is a regular uncountable cardinal such that $\kappa^{<\kappa} = \kappa$ and G is an arbitrary group, then there exists a cardinal-preserving notion of forcing $\mathbb{Q}$ such that G is embeddable in $\mathrm{Sym}(\kappa)/\mathrm{Sym}_\kappa(\kappa)$ within the generic extension $V^{\mathbb{Q}}$. On the other hand, the analogous result is false if we consider the question of which groups can be embedded in $\mathrm{Sym}(\kappa)$. For example, in Section 7.1, we shall prove that the alternating group $\mathrm{Alt}(\kappa^+)$ cannot be embedded in $\mathrm{Sym}(\kappa)$.

In Section 7.3, we shall study the question of which groups G can be realised up to isomorphism as the automorphism groups of first-order structures $\mathcal{M}$ of cardinality κ. There is one obvious constraint on the structure of such a group G; namely, since G acts on a structure $\mathcal{M}$ of cardinality κ, it follows that G is embeddable in $\mathrm{Sym}(\kappa)$. (For example, the alternating group $\mathrm{Alt}(\kappa^+)$ cannot be isomorphic to the automorphism group of a structure of cardinalty κ.) It turns out that if κ is a regular uncountable cardinal such that $\kappa^{<\kappa} = \kappa$, then this is the only constraint on the structure of G. In Section 7.3, we shall prove that if κ is such a cardinal and G is an arbitrary subgroup of $\mathrm{Sym}(\kappa)$, then there exists a cardinal-preserving notion of forcing $\mathbb{P}$ such that in the generic extension $V^{\mathbb{P}}$, there exists a first-order structure $\mathcal{M}$ of cardinality κ such that $G \simeq \mathrm{Aut}(\mathcal{M})$. (A theorem of Dudley [**5**] and Solecki [**47**] shows that the analogous result is false for $\kappa = \omega$ and this is the point where the proof of our main result breaks down in the case when $\kappa = \omega$.)

In the final three sections of this chapter, we shall present some further applications of our main theorems and shall discuss some of the many remaining open problems.

7.1. The nonexistence of a better upper bound

In Chapter 3, we proved that if κ is an infinite cardinal, then $\tau_\kappa < (2^\kappa)^+$. In this section, we shall sketch the proof of the following result, which can be interpreted as saying that if $\kappa > \omega$, then it is impossible to prove a better upper bound for τ_κ in ZFC.

THEOREM 7.1.1. *Let $V \vDash GCH$ and let κ, $\lambda \in V$ be uncountable cardinals such that $\kappa < \mathrm{cf}(\lambda)$. Let α be any ordinal such that $\alpha < \lambda^+$. Then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that the following statements are true in the corresponding generic extension $V^{\mathbb{P}}$.*

(a) $2^\kappa = \lambda$.
(b) *There exists a centreless group T of cardinality κ such that $\tau(T) = \alpha$.*

Our methods do not enable us to deal with the case when $\kappa = \omega$ and the following question remains open.

QUESTION 7.1.2. Is it true, or even consistent, that there exists a countable centreless group G such that $\tau(G) \geq \omega_1$?

Initially we shall only consider the case when κ is a regular uncountable cardinal. We do not need to assume the full GCH, but rather only that $\kappa^{<\kappa} = \kappa$. From now on, fix such a regular uncountable cardinal κ. Let λ be any candidate for 2^κ; i.e. λ is a cardinal such that $\lambda^\kappa = \lambda$. And let α be any ordinal such that $\alpha < \lambda^+$. We shall describe how to construct a cardinal-preserving notion of forcing $\mathbb{P}$ such that the following statements are true in the corresponding generic extension $V^{\mathbb{P}}$.

(a) $2^\kappa = \lambda$.

(b) There exists a centreless group T of cardinality κ such that $\tau(T) = \alpha$.

By Theorem 4.1.9, it is enough to produce a structure $\mathcal{M} \in V^{\mathbb{P}}$ for a first-order language L and a subgroup H of $\mathrm{Aut}(\mathcal{M})$ such that the following conditions are satisfied:

(i) $|\mathcal{M}| = |L| = \kappa$;
(ii) $|H| \leq \kappa$; and
(iii) the normaliser tower of H in $\mathrm{Aut}(\mathcal{M})$ terminates after exactly α steps.

Roughly speaking, our strategy will be to

(1) first construct a pair of groups, $H \leqslant G$, such that $|H| \leq \kappa$ and the normaliser tower of H in G terminates after exactly α steps; and
(2) then attempt to find a cardinal-preserving notion of forcing $\mathbb{P}$ which adjoins a structure $\mathcal{M}$ of cardinality κ such that $G \simeq \mathrm{Aut}(\mathcal{M})$.

Of course, there are many groups G for which such a notion of forcing $\mathbb{P}$ cannot possibly exist. For example, de Bruijn [**3**] has shown that the alternating group $\mathrm{Alt}(\kappa^+)$ cannot be embedded in $\mathrm{Sym}(\kappa)$. Since $\mathrm{Alt}(\kappa^+)$ cannot act on a set of cardinality κ, there is certainly no cardinal-preserving notion of forcing $\mathbb{P}$ which adjoins a structure $\mathcal{M}$ of cardinality κ such that $\mathrm{Alt}(\kappa^+) \simeq \mathrm{Aut}(\mathcal{M})$.

Our next definition singles out a combinatorial condition which is satisfied by all those groups which are embeddable in $\mathrm{Sym}(\kappa)$. (See Proposition 7.1.5.) Conversely, in Theorem 7.1.6, we shall show that if a group G satisfies this combinatorial condition, then there exists a cardinal-preserving notion of forcing $\mathbb{P}$ which adjoins a structure $\mathcal{M}$ of cardinality κ such that $G \simeq \mathrm{Aut}(\mathcal{M})$.

Definition 7.1.3. Suppose that κ is a regular uncountable cardinal such that $\kappa^{<\kappa} = \kappa$. Then a group G is said to satisfy the κ^+-*compatibility condition* if it has the following property. Suppose that H is a group such that $|H| < \kappa$. Suppose that $\langle f_i \mid i < \kappa^+ \rangle$ is a sequence of embeddings $f_i : H \to G$ and let $H_i = f_i[H]$ for each $i < \kappa^+$. Then there exist ordinals $i < j < \kappa^+$ and a surjective homomorphism $\varphi : \langle H_i, H_j \rangle \to H_i$ such that

(a) $\varphi \circ f_j = f_i$; and
(b) $\varphi \upharpoonright H_i = id_{H_i}$.

Of course, if $i < j < \kappa^+$, then there can exist at most one homomorphism $\varphi : \langle H_i, H_j \rangle \to H_i$ satisfying conditions 7.1.3(a) and 7.1.3(b). For we must define the restriction $\varphi \restriction H_i \cup H_j$ by

$$\varphi(h) = \begin{cases} h & \text{if } h \in H_i; \\ (f_i \circ f_j^{-1})(h) & \text{if } h \in H_j. \end{cases}$$

By the Δ-System Lemma, we can assume that there exists a fixed subgroup $K \leqslant G$ such that $H_i \cap H_j = K$ for all $i < j < \kappa^+$; and an easy counting argument allows us to also assume that there exists a fixed homomorphism $\theta : K \to H$ such that $f_i^{-1} \restriction K = \theta$ for all $i < \kappa^+$. This implies that the partial mapping $\varphi \restriction H_i \cup H_j$ is well-defined for all $i < j < \kappa^+$. However, there still remains the question of finding a pair of ordinals $i < j < \kappa^+$ such that this partial mapping extends to a well-defined homomorphism $\varphi : \langle H_i, H_j \rangle \to H_i$.

Example 7.1.4. In order to get an better understanding of Definition 7.1.3, it will probably be helpful to see an example of a group which *fails* to satisfy the κ^+-compatibility condition. So we shall show that $\mathrm{Alt}(\kappa^+)$ does not satisfy the κ^+-compatibility condition. Let $H = \mathrm{Alt}(4)$. For each $3 \leq i < \kappa^+$, let $\Delta_i = \{0, 1, 2, i\}$ and let $f_i : \mathrm{Alt}(4) \to \mathrm{Alt}(\Delta_i)$ be an isomorphism. If $3 \leq i < j < \kappa^+$, then

$$\langle \mathrm{Alt}(\Delta_i), \mathrm{Alt}(\Delta_j) \rangle = \mathrm{Alt}(\Delta_i \cup \Delta_j) \simeq \mathrm{Alt}(5).$$

Since $\mathrm{Alt}(5)$ is a simple group, there does not exist a surjective homomorphism from $\langle \mathrm{Alt}(\Delta_i), \mathrm{Alt}(\Delta_j) \rangle$ onto $\mathrm{Alt}(\Delta_i)$.

Proposition 7.1.5. *Let κ be a regular uncountable cardinal such that $\kappa^{<\kappa} = \kappa$ and let $G \leqslant \mathrm{Sym}(\kappa)$. Then G satisfies the κ^+-compatibility condition.*

Proof. Let H be a group such that $|H| < \kappa$ and let $\langle f_i \mid i < \kappa^+ \rangle$ be a sequence of embeddings $f_i : H \to G$. For each $i < \kappa^+$, let $H_i = f_i[H]$ and let Z_i be a subset of κ chosen so that

(a) $|Z_i| < \kappa$;
(b) Z_i is H_i-invariant; and
(c) $g \restriction Z_i \neq id_{Z_i}$ for all $1 \neq g \in H_i$.

Since there are only $\kappa^{<\kappa} = \kappa$ possibilities for Z_i, after passing to a suitable subsequence if necessary, we can suppose that

(1) there exists a fixed subset Z such that $Z_i = Z$ for all $i < \kappa^+$.

Similarly we can suppose that the following condition holds.

(2) For each ordinal $i < \kappa^+$, let $r_i : H_i \to \mathrm{Sym}(Z)$ be the restriction mapping, $g \mapsto g \restriction Z$. Then $r_i \circ f_i = r_j \circ f_j$ for all $i < j < \kappa^+$.

Fix any pair of ordinals i, j such that $i < j < \kappa^+$. Let $\rho : \langle H_i, H_j \rangle \to \mathrm{Sym}(Z)$ be the restriction mapping, $g \mapsto g \restriction Z$. Then $\rho\left[\langle H_i, H_j \rangle\right] = r_i\left[H_i\right]$ and so we can define a surjective homomorphism $\varphi : \langle H_i, H_j \rangle \to H_i$ by $\varphi = r_i^{-1} \circ \rho$. Clearly $\varphi \restriction H_i = id_{H_i}$. By condition (2), if $h \in H$, then

$$(\rho \circ f_j)(h) = (r_j \circ f_j)(h) = (r_i \circ f_i)(h)$$

and so

$$(\varphi \circ f_j)(h) = (r_i^{-1} \circ \rho \circ f_j)(h) = f_i(h).$$

Hence $\varphi \circ f_j = f_i$. □

If the group H in the proof of Proposition 7.1.5 happens to be finite, then the assumption that $\kappa^{<\kappa} = \kappa$ is not needed. So combining Proposition 7.1.5 and Example 7.1.4, we obtain an alternative proof of de Bruijn's theorem that $\mathrm{Alt}(\kappa^+)$ does not embed into $\mathrm{Sym}(\kappa)$.

We shall prove the following theorem in Section 7.3.

THEOREM 7.1.6. *Let κ be a regular uncountable cardinal such that $\kappa^{<\kappa} = \kappa$ and let G be a group which satisfies the κ^+-compatibility condition. Let L be a first-order language consisting of κ binary relation symbols. Then there exists a notion of forcing $\mathbb{P}$ such that*

(a) *$\mathbb{P}$ is κ-closed;*
(b) *$\mathbb{P}$ has the κ^+-c.c.; and*
(c) *$\mathbb{P}$ adjoins an L-structure $\mathcal{M}$ of cardinality κ such that $G \simeq \mathrm{Aut}(\mathcal{M})$.*

Furthermore, if $|G| = \theta$, then $|\mathbb{P}| = \max\{\kappa, \theta^{<\kappa}\}$.

Combining Proposition 7.1.5 and Theorem 7.1.6, we see that if κ is an uncountable cardinal such that $\kappa^{<\kappa} = \kappa$ and G is an arbitrary subgroup of $\mathrm{Sym}(\kappa)$, then there exists a cardinal-preserving notion of forcing $\mathbb{P}$ and a structure $\mathcal{M} \in V^{\mathbb{P}}$ of cardinality κ such that $G \simeq \mathrm{Aut}(\mathcal{M})$. This is the point at which our proof breaks down in the case when $\kappa = \omega$, since the analogue of Theorem 7.1.6 is false for $\kappa = \omega$.

For example, let A be a free abelian group such that $\omega < |A| \leq 2^{\omega}$. Then it is an easy exercise to show that A is embeddable in $\mathrm{Sym}(\omega)$. However, Dudley [**5**] has shown that the automorphism group of a countable first-order structure is never an uncountable free abelian group. (This result was independently rediscovered by Solecki [**47**].)

We have now reduced our problem to that of finding a pair of subgroups

$$H \leqslant G < \mathrm{Sym}(\kappa)$$

such that $|H| \leq \kappa$ and the normaliser tower of H in G terminates after exactly α steps. We cannot expect to always find such a pair in the ground model V. For example, we might have chosen $\alpha > 2^{\kappa}$. However, such a pair always exists in a suitably chosen generic extension of V. In Section 7.2, we shall prove that there is a notion of forcing $\mathbb{Q}$ such that

(1) $\mathbb{Q}$ is κ-closed;
(2) $\mathbb{Q}$ has the κ^{+}- *c.c.*.

and such that the following statements are true in the generic extension $V^{\mathbb{Q}}$.

(a) $2^{\kappa} = \lambda$.
(b) There exists a pair of subgroups $H \leqslant G < \mathrm{Sym}(\kappa)$ such that $|H| = \kappa$ and the normaliser tower of H in G terminates after exactly α steps.

Working within $V^{\mathbb{Q}}$, we have that G satisfies the κ^{+}-compatibility condition; and since $\mathbb{Q}$ is κ-closed, we still have that $\kappa^{<\kappa} = \kappa$. Hence we can use Theorem 7.1.6 to generically adjoin a structure $\mathcal{M} \in V^{\mathbb{Q} * \tilde{\mathbb{P}}}$ of cardinality κ such that $G \simeq \mathrm{Aut}(\mathcal{M})$. Since the normaliser tower of H in G is an upwards absolute notion, we can now use Theorem 4.1.9 in $V^{\mathbb{Q} * \tilde{\mathbb{P}}}$ to obtain a centreless group T of cardinality κ such that $\tau(T) = \alpha$. (Here $\tilde{\mathbb{P}}$ denotes a $\mathbb{Q}$-name of the notion of forcing $\mathbb{P} \in V^{\mathbb{Q}}$, which is given by Theorem 7.1.6.)

Finally suppose that κ is a singular cardinal. Once again, let λ be any cardinal such that $\lambda^{\kappa} = \lambda$ and let α be any ordinal such that $\alpha < \lambda^{+}$. Then we also have that $\lambda^{\omega_1} = \lambda$. By the above argument, there is a generic extension $V^{\mathbb{Q} * \tilde{\mathbb{P}}}$ in which the following statements are true.

(i) $2^{\omega_1} = 2^{\kappa} = \lambda$.
(ii) There exists a centreless group T of cardinality ω_1 such that $\tau(T) = \alpha$.

First suppose that $\alpha \geq 1$. Then $G = T \times \mathrm{Alt}(\kappa)$ is a centreless group of cardinality κ; and by Theorem 3.2.1, $\tau(G) = \tau(T) = \alpha$. On the other hand, if $\alpha = 0$, then we can let G be any complete group of cardinality κ. For example, we could let $G = PGL(2, K)$, where K is a rigid field of cardinality κ. (The existence of such a field follows from Theorems 4.1.8 and 4.1.7.)

Some readers may feel dissatisfied with the proof in the case when κ is singular; and it has to be admitted that the argument does avoid confronting the difficulties in this case. In particular, all problems of the following type remain open.

Conjecture 7.1.7. It is consistent that

(a) $2^{\aleph_n} = \aleph_{n+1}$ for all $n \in \omega$; and
(b) there exists a centreless group G of cardinality $\aleph_\omega$ such that $\tau(G) \geq \aleph_{\omega+1}$.

In the remainder of this section, we shall reconsider the question of whether it is possible to find a better upper bound for $\tau(G)$ when G is an arbitrary (not necessarily centreless) group. Recall that, in Section 5.1, we were only able to prove that if κ is the least inaccessible cardinal such that $\kappa > |G|$, then $\tau(G) < \kappa$. It is natural to ask whether there exists an analogue of Theorem 7.1.1, which would show that it is impossible to prove a better upper bound in ZFC. Of course, in this case, we would also have to rule out such upper bounds as $\beth_{|G|}(|G|)$, etc. For example, the following conjecture would serve our purpose.

Conjecture 7.1.8. Let $V \vDash GCH$ and let $\kappa \in V$ be an infinite cardinal. Let $\lambda \in V$ be the least inaccessible cardinal such that $\lambda > \kappa$ and let α be any ordinal such that $\alpha < \lambda$. Then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that the following statements are true in the corresponding generic extension $V^{\mathbb{P}}$.

(a) GCH holds.
(b) There exists a group G of cardinality κ such that $\tau(G) = \alpha$.

It is even conceivable that the corresponding conjecture might hold for finite groups. In other words, if $V \vDash GCH$ and α is less than the first inaccessible cardinal, then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that the following statements are true in the corresponding generic extension $V^{\mathbb{P}}$.

(a) GCH holds.

(b) There exists a finite group G such that $\tau(G) = \alpha$.

Of course, in this case, the finite group G would already exist in the ground model V. But this is not a contradiction, since there is no reason to expect that the automorphism tower of a finite group G should be an absolute notion. (For example, it could be that G_ω is an infinite group and then the rest of the automorphism tower could perhaps be manipulated via suitable notions of forcing.) However, this would mean that there exists a fixed finite group G such that for cofinally many ordinals α less than the first inaccessible cardinal, it is consistent that $\tau(G) = \alpha$. When I raised this possibility in a recent talk, Leo Harrington asked me whether I had a specific candidate in mind for such a group!

7.2. Realising normaliser towers within infinite symmetric groups

Let κ be a regular uncountable cardinal such that $\kappa^{<\kappa} = \kappa$. In this section, we shall study the problem of realising long normaliser towers within $\mathrm{Sym}(\kappa)$. In particular, we shall prove that if α is any ordinal, then there exists a generic extension $V^{\mathbb{Q}}$ such that a normaliser tower of height α can be realised in $\mathrm{Sym}^{V^{\mathbb{Q}}}(\kappa)$.

We shall make use of the ascending chain of groups $\langle W_\alpha \mid \alpha \in On \rangle$ which we defined in Section 4.1. For the reader's convenience, we repeat the definition here.

Definition 7.2.1. The ascending chain of groups

$$W_0 \leqslant W_1 \leqslant \cdots \leqslant W_\alpha \leqslant W_{\alpha+1} \leqslant \dots$$

is defined inductively as follows.

(a) $W_0 = C_2$, the cyclic group of order 2.

(b) Suppose that $\alpha = \beta + 1$. Then

$$W_\beta = W_\beta \oplus 1 \leqslant \left[W_\beta \oplus W_\beta^* \right] \rtimes \langle \sigma_{\beta+1} \rangle = W_{\beta+1}.$$

Here W_β^* is an isomorphic copy of W_β; and $\sigma_{\beta+1}$ is an element of order 2 which interchanges the factors $W_\beta \oplus 1$ and $1 \oplus W_\beta^*$ of the direct sum $W_\beta \oplus W_\beta^*$ via conjugation. Thus $W_{\beta+1}$ is isomorphic to the wreath product $W_\beta \text{ Wr } C_2$.

(c) If α is a limit ordinal, then $W_\alpha = \bigcup_{\beta<\alpha} W_\beta$.

By Lemmas 4.1.11 and 4.1.12, the groups $\langle W_\alpha \mid \alpha \in On \rangle$ satisfy the following properties.

(i) If $1 \leq n < \omega$, then the normaliser tower of W_0 in W_n terminates after exactly $n+1$ steps.

(ii) If $\alpha \geq \omega$, then the normaliser tower of W_0 in W_α terminates after exactly α steps.

(iii) $|W_\alpha| \leq \max\{|\alpha|, \omega\}$ for all ordinals α.

Unfortunately, the group W_{κ^+} does not satisfy the κ^+-compatibility condition. To see this, let

$$H = C_2 \text{ Wr } C_2 = [\langle a \rangle \oplus \langle b \rangle] \rtimes \langle c \rangle,$$

where the involution c interchanges a and b via conjugation; and for each limit ordinal $i < \kappa^+$, let $f_i : H \to W_{\kappa^+}$ be the embedding such that $f_i(a) = \sigma_{i+1}$ and $f_i(c) = \sigma_{i+2}$. (Here we are using the notation which was introduced in Definition 7.2.1.) Let i, j be any limit ordinals such that $i < j < \kappa^+$. Then

$$\langle H_i, H_j \rangle \simeq ((C_2 \text{ Wr } C_2) \text{ Wr } C_2) \text{ Wr } C_2.$$

Suppose that there exists a surjective homomorphism $\varphi : \langle H_i, H_j \rangle \to H_i$ such that

(a) $\varphi \circ f_j = f_i$ and

(b) $\varphi \restriction H_i = id_{H_i}$.

Then $\varphi(\sigma_{j+1}) = \varphi(\sigma_{i+1}) = \sigma_{i+1}$ and $\varphi(\sigma_{j+2}) = \varphi(\sigma_{i+2}) = \sigma_{i+2}$. Consider the element $x = \sigma_{j+1}\sigma_{j+2}\sigma_{j+1}\sigma_{j+2} \in H_j$. Then it is easily checked that

(1) x lies in the centre of H_j and

(2) $\sigma_{j+1} y \sigma_{j+1}^{-1} = xyx^{-1}$ for all $y \in H_i$.

Thus $z = \varphi(x)$ lies in the centre of H_i. Since $\varphi(\sigma_{j+1}) = \sigma_{i+1}$, we find that

$$\sigma_{i+1} y \sigma_{i+1}^{-1} = zyz^{-1} = y$$

for all $y \in H_i$. But this contradicts the fact that σ_{i+1} is a noncentral element of H_i.

Thus if $\alpha \geq \kappa^+$, then W_α is not embeddable in $\mathrm{Sym}(\kappa)$. However, the above argument does not rule out the possibility that W_α is embeddable in the quotient group $\mathrm{Sym}(\kappa)/\mathrm{Sym}_\kappa(\kappa)$; and this is enough for our purposes. (Recall that $\mathrm{Sym}_\kappa(\kappa)$

denotes the normal subgroup of $\mathrm{Sym}(\kappa)$ consisting of all those permutations φ such that $|\operatorname{supp}(\varphi)| < \kappa$.)

LEMMA 7.2.2. *Let κ be an infinite cardinal such that $\kappa^{<\kappa} = \kappa$. Suppose that γ is an ordinal and that there exists an embedding*

$$f : W_\gamma \to \mathrm{Sym}(\kappa)/\operatorname{Sym}_\kappa(\kappa).$$

Then for each ordinal $\alpha \leq \gamma$, there exist groups $H \leqslant G \leqslant \mathrm{Sym}(\kappa)$ such that $|H| = \kappa$ and the normaliser tower of H in G terminates after exactly α steps.

PROOF. For each ordinal $\alpha \leq \gamma$, let H^α be the subgroup of $\mathrm{Sym}(\kappa)$ such that $f[W_\alpha] = H^\alpha/\operatorname{Sym}_\kappa(\kappa)$. Since $|\operatorname{Sym}_\kappa(\kappa)| = \kappa^{<\kappa} = \kappa$, it follows that $|H^0| = \kappa$.

CLAIM 7.2.3. *For each ordinal β,*

$$f[N_\beta(W_0, W_\alpha)] = N_\beta(H^0, H^\alpha)/\operatorname{Sym}_\kappa(\kappa).$$

PROOF. We will argue by induction on β. The result is clear when $\beta = 0$ and no difficulties arise when β is a limit ordinal. Suppose that $\beta = \xi + 1$ and that the result holds for ξ. Let $R = N_\xi(H^0, H^\alpha)$ and for each subgroup K such that $\operatorname{Sym}_\kappa(\kappa) \leqslant K \leqslant H^\alpha$, let $\overline{K} = K/\operatorname{Sym}_\kappa(\kappa)$. Then

$$f[N_{\xi+1}(W_0, W_\alpha)] = N_{\overline{H^\alpha}}(\overline{R})$$

and so we must show that

$$N_{\overline{H^\alpha}}(\overline{R}) = \overline{N_{H^\alpha}(R)}.$$

But this is an immediate consequence of the Correspondence Theorem for subgroups of quotient groups, together with the observation that the normaliser of any subgroup L is the largest subgroup M such that $L \trianglelefteq M$. □

It is now easy to complete the proof of Lemma 7.2.2. Applying Claim 7.2.3, we see that if $\alpha \geq \omega$, then the normaliser tower of H^0 in H^α terminates after exactly α steps; and that if $2 \leq \alpha = n < \omega$, then the normaliser tower of H^0 in H^{n-1} terminates after exactly n steps. This just leaves the cases when $\alpha = 0, 1$. When $\alpha = 0$, we can take $H = G = \mathrm{Alt}(\kappa)$; and when $\alpha = 1$, we can take $H = \mathrm{Alt}(\kappa)$ and $G = \mathrm{Sym}(\kappa)$. □

The next result implies that if $\omega < \kappa^{<\kappa} = \kappa$ and W is any group, then there exists a cardinal-preserving notion of forcing $\mathbb{Q}$ such that in $V^{\mathbb{Q}}$, the group W is embeddable in $\mathrm{Sym}(\kappa)/\mathrm{Sym}_\kappa(\kappa)$.

LEMMA 7.2.4. *Suppose that $\kappa \in V$ is a regular uncountable cardinal such that $\kappa^{<\kappa} = \kappa$. If W is any group, then there exists a notion of forcing $\mathbb{Q}$ such that*

(1) *$\mathbb{Q}$ is κ-closed;*
(2) *$\mathbb{Q}$ has the κ^+-c.c.; and*
(3) *in the generic extension $V^{\mathbb{Q}}$, there exists an isomorphic embedding*

$$f : W \to (\mathrm{Sym}(\kappa)/\mathrm{Sym}_\kappa(\kappa))^{V^{\mathbb{Q}}}.$$

Furthermore, if $|W| = \theta$, then $|\mathbb{Q}| = \max\{\kappa, \theta^{<\kappa}\}$.

PROOF. Let $\Omega = \bigcup_{\alpha<\kappa}\{\alpha\} \times \alpha$. We will work with the symmetric group $\mathrm{Sym}(\Omega)$ rather than with $\mathrm{Sym}(\kappa)$. Let $\mathbb{Q}$ be the notion of forcing consisting of the conditions

$$p = (\delta_p, H_p, E_p)$$

such that the following hold.

(a) $\omega \leq \delta_p < \kappa$.
(b) H_p is a subgroup of W such that $|H_p| \leq |\delta_p|$.
(c) E_p is a function which assigns a permutation $e^p_{\pi,\xi} \in \mathrm{Sym}(\{\xi\} \times \xi)$ to each pair $(\pi, \xi) \in H_p \times \delta_p$.

We set $q = (\delta_q, H_q, E_q) \leq p = (\delta_p, H_p, E_p)$ iff

(1) $\delta_p \leq \delta_q$;
(2) $H_p \leqslant H_q$;
(3) $E_p \subseteq E_q$; and
(4) if $\delta_p \leq \xi < \delta_q$, then the restriction to H_p of the function, $\pi \mapsto e^q_{\pi,\xi}$, is an isomorphic embedding of H_p into $\mathrm{Sym}(\{\xi\} \times \xi)$.

CLAIM 7.2.5. *$\mathbb{Q}$ is κ-closed.*

PROOF OF CLAIM 7.2.5. This is clear. □

CLAIM 7.2.6. *$\mathbb{Q}$ has the κ^+-c.c..*

Proof of Claim 7.2.6. Suppose that $p_i = (\delta_{p_i}, H_{p_i}, E_{p_i}) \in \mathbb{Q}$ for $i < \kappa^+$. Then, after passing to a suitable subsequence of $\langle p_i \mid i < \kappa^+ \rangle$, we can suppose that the following condition holds.

(i) There exists a fixed ordinal δ such that $\delta_{p_i} = \delta$ for all $i < \kappa^+$.

Now consider the family $\{H_{p_i} \mid i < \kappa^+\}$ of subgroups of W. Since $\kappa^{<\kappa} = \kappa$, the Δ-System Lemma says that there exists a fixed subgroup H and a subset $I \subset \kappa^+$ of cardinality κ^+ such that $H_{p_i} \cap H_{p_j} = H$ for all pairs of distinct elements $i, j \in I$. Hence we can also suppose that the following condition holds.

(ii) There exists a fixed subgroup H such that

$$H_{p_i} \cap H_{p_j} = H$$

for all $i < j < \kappa^+$.

Finally since $|H| < \kappa$ and $\kappa^{<\kappa} = \kappa$, there are at most κ many functions

$$E : H \times \delta \to \bigcup_{\xi < \delta} \mathrm{Sym}(\{\xi\} \times \xi).$$

Hence we can also suppose that the following condition holds.

(iii) There exists a fixed function E such that $E_{p_i} \restriction H \times \delta = E$ for all $i < \kappa^+$.

Now fix any two ordinals $i < j < \kappa^+$. Let $H^+ = \langle H_{p_i}, H_{p_j} \rangle$ be the subgroup generated by $H_{p_i} \cup H_{p_j}$ and let $E^+ : H^+ \times \delta \to \bigcup_{\xi<\delta} \mathrm{Sym}(\{\xi\} \times \xi)$ be any extension of $E_{p_i} \cup E_{p_j}$ which satisfies condition (c). Then $q = (\delta, H^+, E^+)$ is a common lower bound of p_i and p_j. □

Claim 7.2.7. *For each $\alpha < \kappa$, the set $C_\alpha = \{q \in \mathbb{Q} \mid \delta_q \geq \alpha\}$ is dense in $\mathbb{Q}$.*

Proof of Claim 7.2.7. Let $p = (\delta_p, H_p, E_p) \in \mathbb{Q}$. Then we can suppose that $\delta_p < \alpha$. We can define an isomorphic embedding $\varphi : H_p \to \mathrm{Sym}(H_p)$ by setting $\varphi(h)(x) = hx$ for all $x \in H_p$. Since $|H_p| \leq |\delta_p|$, it follows that there exists an isomorphic embedding $\varphi_\xi : H_p \to \mathrm{Sym}(\{\xi\} \times \xi)$ for each $\delta_p \leq \xi < \alpha$. Hence there exists a condition $q = (\delta_q, H_q, E_q) \leq p$ such that $H_q = H_p$ and $\delta_q = \alpha$. □

Claim 7.2.8. *For each $\pi \in W$, the set $D_\pi = \{q \in \mathbb{Q} \mid \pi \in H_q\}$ is dense in $\mathbb{Q}$.*

Proof of Claim 7.2.8. Let $p = (\delta_p, H_p, E_p) \in \mathbb{Q}$. Then we can suppose that $\pi \notin H_p$. Let $H^+ = \langle H_p, \pi \rangle$ be the subgroup generated by $H_p \cup \{\pi\}$. Let $\delta = \delta_p$

and let $E^+ : H^+ \times \delta \to \bigcup_{\xi<\delta} \operatorname{Sym}(\{\xi\} \times \xi)$ be any extension of E_p which satisfies condition (c). Then $q = (\delta, H^+, E^+) \leq p$. $\square$

Let F be a $\mathbb{Q}$-generic filter over V and let $V^{\mathbb{Q}} = V[F]$ be the corresponding generic extension. Working within $V^{\mathbb{Q}}$, for each $\pi \in W$, let

$$e(\pi) = \bigcup \{e^p_{\pi,\xi} \mid \text{ There exists } p \in F \text{ such that } \pi \in H_p \text{ and } \xi < \delta_p\}.$$

Then $e(\pi) \in \operatorname{Sym}(\Omega)$. Let $\operatorname{Sym}_\kappa(\Omega) = \{\psi \in \operatorname{Sym}(\Omega) \mid |\operatorname{supp}(\psi)| < \kappa\}$ and define the function

$$f : W \to \operatorname{Sym}(\Omega)/\operatorname{Sym}_\kappa(\Omega)$$

by $f(\pi) = e(\pi)\operatorname{Sym}_\kappa(\Omega)$. Then it is enough to show that f is an isomorphic embedding.

Claim 7.2.9. *If $1 \neq \pi \in W$, then $f(\pi) \neq 1$.*

Proof of Claim 7.2.9. Choose a condition $p = (\delta_p, H_p, E_p) \in F$ such that $\pi \in H_p$. If ξ is any ordinal such that $\delta_p \leq \xi < \kappa$, then $e(\pi) \upharpoonright \{\xi\} \times \xi \neq id_{\{\xi\}\times\xi}$. Hence $e(\pi) \notin \operatorname{Sym}_\kappa(\Omega)$. $\square$

Claim 7.2.10. *f is a group homomorphism.*

Proof of Claim 7.2.10. Let π_1, $\pi_2 \in W$. Let $p = (\delta_p, H_p, E_p) \in F$ be a condition such that π_1, $\pi_2 \in H_p$. Let ξ be any ordinal such that $\delta_p \leq \xi < \kappa$ and let $q \in F$ be a condition such that $q \leq p$ and $\xi < \delta_q$. Then

$$e^q_{\pi_1,\xi} \circ e^q_{\pi_2,\xi} = e^q_{\pi_1\circ\pi_2,\xi}$$

and it follows that

$$e(\pi_1)\operatorname{Sym}_\kappa(\Omega) \circ e(\pi_2)\operatorname{Sym}_\kappa(\Omega) = e(\pi_1 \circ \pi_2)\operatorname{Sym}_\kappa(\Omega).$$

$\square$

Finally it is easily checked that $|\mathbb{Q}| = \max\{\kappa, \theta^{<\kappa}\}$. This completes the proof of Lemma 7.2.4. $\square$

Summing up our work in this section, we can now easily obtain the following result, which was promised in Section 7.1. (The final observation in the statement of Theorem 7.2.11 will be required in Section 7.4.)

Theorem 7.2.11. *Suppose that κ, $\lambda \in V$ are cardinals such that $\omega < \kappa^{<\kappa} = \kappa < \lambda = \lambda^{\kappa}$. Let α be any ordinal such that $\alpha < \lambda^+$. Then there exists a notion of forcing $\mathbb{Q}$ such that*

(1) *$\mathbb{Q}$ is κ-closed;*
(2) *$\mathbb{Q}$ has the κ^+-c.c.;*

and such that the following statements are true in the generic extension $V^{\mathbb{Q}}$.

(a) *$2^{\kappa} = \lambda$.*
(b) *There exist groups $H \leqslant G < \mathrm{Sym}(\kappa)$ such that $|H| = \kappa$ and the normaliser tower of H in G terminates after exactly α steps.*

Furthermore, if $V \vDash GCH$ and $\lambda = \kappa^+$, then $V^{\mathbb{Q}} \vDash GCH$.

Proof. Let γ be an ordinal such that $\max\{\alpha, \lambda\} \leq \gamma < \lambda^+$ and let $\mathbb{Q}$ be the notion of forcing obtained by applying Lemma 7.2.4 to $W = W_\gamma$. Then $|W_\gamma| = \lambda$ and so $|\mathbb{Q}| = \lambda^{<\kappa} = \lambda$. Since W_γ embeds in $\mathrm{Sym}(\kappa)/\mathrm{Sym}_\kappa(\kappa)$ in $V^{\mathbb{Q}}$, it follows that $V^{\mathbb{Q}} \vDash 2^{\kappa} \geq \lambda$. On the other hand, since $|\mathbb{Q}| = \lambda$ and $\mathbb{Q}$ has the κ^+-*c.c.*, it follows that there are at most $\lambda^{\kappa} = \lambda$ nice names for subsets of κ. Hence Lemma 6.4.2(b) implies that $V^{\mathbb{Q}} \vDash 2^{\kappa} \leq \lambda$. (If $V \vDash GCH$ and $\lambda = \kappa^+$, then a similar argument shows that $V^{\mathbb{Q}} \vDash 2^{\mu} = \mu^+$ for all cardinals $\mu \geq \kappa$. Since $\mathbb{Q}$ is κ-closed, if $\mu < \kappa$, then $({}^{\mu}2)^{V^{\mathbb{Q}}} = ({}^{\mu}2)^{V}$ and hence $V^{\mathbb{Q}} \vDash 2^{\mu} = \mu^+$.) Finally Lemma 7.2.2 implies that there exist groups $H \leqslant G < \mathrm{Sym}(\kappa)$ in $V^{\mathbb{Q}}$ such that $|H| = \kappa$ and the normaliser tower of H in G terminates after exactly α steps. $\square$

7.3. Closed groups of uncountable degree

In this section, we shall prove Theorem 7.1.6. Let κ be a regular uncountable cardinal such that $\kappa^{<\kappa} = \kappa$ and let G be a group of cardinality θ which satisfies the κ^+-compatibility condition. Let L be a first-order language consisting of κ binary relation symbols. The following notion of forcing $\mathbb{P}$ is designed to adjoin a structure $\mathcal{M}$ of cardinality κ for the language L such that $G \simeq \mathrm{Aut}(\mathcal{M})$.

Definition 7.3.1. Suppose that $L_0 \subseteq L$ and that $\mathcal{N}$ is a structure for the language L_0. Then a *restricted atomic type* in the free variable v for the language L_0 using parameters from $\mathcal{N}$ is a set t of formulas of the form $R(v, a)$, where $R \in L_0$ and $a \in \mathcal{N}$. An element $c \in \mathcal{N}$ is said to *realise* t if $\mathcal{N} \vDash \varphi[c]$ for every formula

$\varphi(v) \in t$. If no element of $\mathcal{N}$ realises t, then t is said to be *omitted* in $\mathcal{N}$. Notice that every element of $\mathcal{N}$ realises the *trivial restricted atomic type* $\emptyset$. Hence if t is omitted in $\mathcal{N}$, then $t \neq \emptyset$.

DEFINITION 7.3.2. Let $\mathbb{P}$ be the notion of forcing consisting of the conditions

$$p = (H, \pi, \mathcal{N}, T)$$

such that the following hold.

(a) H is a subgroup of G such that $|H| < \kappa$.
(b) There exists an ordinal $0 < \delta < \kappa$ and a subset $L(\mathcal{N}) \in [L]^{<\kappa}$ such that $\mathcal{N}$ is a structure with universe δ for the language $L(\mathcal{N})$.
(c) $\pi : H \to \mathrm{Aut}(\mathcal{N})$ is a group homomorphism.
(d) T is a set of restricted atomic types in the free variable v for the language $L(\mathcal{N})$ using parameters from $\mathcal{N}$. Furthermore, $|T| < \kappa$ and each $t \in T$ is omitted in $\mathcal{N}$.

We set $(H_2, \pi_2, \mathcal{N}_2, T_2) \leq (H_1, \pi_1, \mathcal{N}_1, T_1)$ if and only if

(1) $H_1 \leqslant H_2$.
(2) $\mathcal{N}_1$ is a substructure of $\mathcal{N}_2$.
(3) For all $h \in H_1$ and $\alpha \in \mathcal{N}_1$, $\pi_2(h)(\alpha) = \pi_1(h)(\alpha)$.
(4) $T_1 \subseteq T_2$.

It should be clear that the components $(H, \pi, \mathcal{N})$ in each condition $p \in \mathbb{P}$ are designed to generically adjoin a structure $\mathcal{M}$ of cardinality κ for the language L, together with an embedding π^* of G into $\mathrm{Aut}(\mathcal{M})$. The set T of restricted atomic types is needed to kill off potential extra automorphisms $g \in \mathrm{Aut}(\mathcal{M}) \smallsetminus \pi^*[G]$ and thus ensure that π^* is surjective.

LEMMA 7.3.3. *For each $p = (H, \pi, \mathcal{N}, T) \in \mathbb{P}$, there exists a condition*

$$p^+ = \left(H, \pi^+, \mathcal{N}^+, T\right) \leq p$$

such that $\pi^+ : H \to \mathrm{Aut}(\mathcal{N}^+)$ is an embedding.

PROOF. Let $\mathcal{N}^+$ be the structure for the language $L(\mathcal{N})$ such that

(a) the universe of $\mathcal{N}^+$ is the disjoint union $\mathcal{N} \sqcup H$;

(b) for each relation $R \in L(\mathcal{N})$, $R^{\mathcal{N}^+} = R^{\mathcal{N}}$.

Clearly none of the restricted atomic types in T is realised in $\mathcal{N}^+$. Let $\pi^+ : H \to \mathrm{Aut}(\mathcal{N}^+)$ be the embedding such that for each $h \in H$,

(i) $\pi^+(h)(x) = \pi(h)(x)$ for all $x \in \mathcal{N}$; and

(ii) $\pi^+(h)(x) = hx$ for all $x \in H$.

Then $p^+ = (H, \pi^+, \mathcal{N}^+, T) \leq p$. □

There is a slight inaccuracy in the proof of Lemma 7.3.3, as the universe of $\mathcal{N}^+$ should really be an ordinal $\delta < \kappa$. However, the proof can easily be repaired: simply replace $\mathcal{N}^+$ by a suitable isomorphic structure. Similar remarks apply to the proofs of Lemmas 7.3.6 and 7.3.8.

LEMMA 7.3.4. *$\mathbb{P}$ is κ-closed.*

PROOF. Suppose that $\delta < \kappa$ and that

$$p_0 \geq p_1 \geq \cdots \geq p_\xi \geq \cdots$$

is a descending δ-sequence of elements of $\mathbb{P}$. Then $\{p_\xi \mid \xi < \delta\}$ has a greatest lower bound in $\mathbb{P}$. To see this, let $p_\xi = (H_\xi, \pi_\xi, \mathcal{N}_\xi, T_\xi)$ for each $\xi < \delta$ and define $H = \bigcup_{\xi<\delta} H_\xi$, $\pi = \bigcup_{\xi<\delta} \pi_\xi$, $\mathcal{N} = \bigcup_{\xi<\delta} \mathcal{N}_\xi$, and $T = \bigcup_{\xi<\delta} T_\xi$. Since each $t \in T$ is a restricted atomic type, it follows that each $t \in T$ is also omitted in $\mathcal{N}$. Thus $p = (H, \pi, \mathcal{N}, T) \in \mathbb{P}$ and clearly p is the greatest lower bound of $\{p_\xi \mid \xi < \delta\}$ in $\mathbb{P}$. □

LEMMA 7.3.5. *$\mathbb{P}$ has the κ^+-c.c.*

PROOF. Suppose that $p_i = (H_i, \pi_i, \mathcal{N}_i, T_i) \in \mathbb{P}$ for $i < \kappa^+$. By Lemma 7.3.3, we can assume that $\pi_i : H_i \to \mathrm{Aut}(\mathcal{N}_i)$ is an embedding for each $i < \kappa$. Clearly there is a cardinal $\lambda < \kappa$ and a subset $I \in [\kappa^+]^{\kappa^+}$ such that $|H_i| = \lambda$ for all $i \in I$. Since there are only 2^λ groups of cardinality λ up to isomorphism and $2^\lambda \leq \kappa$, there exists a subset $J \in [I]^{\kappa^+}$ such that $H_i \simeq H_j$ for all $i, j \in J$. Thus, after passing to a suitable subsequence of $\langle p_i \mid i < \kappa^+ \rangle$, we can suppose that the following condition holds.

(1) There is a fixed group H such that for each $i < \kappa^+$, there exists an isomorphism $f_i : H \simeq H_i$.

Similarly, we can suppose that the following conditions also hold.

(2) There exists a fixed structure $\mathcal{N}$ such that $\mathcal{N}_i = \mathcal{N}$ for all $i < \kappa^+$.

(3) There exists a fixed set of restricted atomic types T such that $T_i = T$ for all $i < \kappa^+$.

(4) For each $i < \kappa^+$, let $\psi_i : H \to \mathrm{Aut}(\mathcal{N})$ be the embedding defined by $\psi_i = \pi_i \circ f_i$. Then $\psi_i = \psi_j$ for all $i < j < \kappa^+$.

Since G satisfies the κ^+-compatibility condition, there exist ordinals $i < j < \kappa^+$ and a surjective homomorphism $\varphi : \langle H_i, H_j \rangle \to H_i$ such that

(a) $\varphi \circ f_j = f_i$; and

(b) $\varphi \restriction H_i = id_{H_i}$.

Let $\langle H_i, H_j \rangle \to \mathrm{Aut}(\mathcal{N})$ be the homomorphism defined by $\pi = \pi_i \circ \varphi$. Clearly $\pi_i \subseteq \pi$. Note that if $x \in H_j$, then

$$\begin{aligned}\pi_i \circ \varphi(x) &= \pi_i \circ (\varphi \circ f_j) \circ f_j^{-1}(x) \\ &= \pi_i \circ f_i \circ f_j^{-1}(x) \\ &= \pi_j \circ f_j \circ f_j^{-1}(x) \\ &= \pi_j(x).\end{aligned}$$

Thus we also have that $\pi_j \subseteq \pi$. Consequently, we can define a condition $p \leq p_i$, p_j by

$$p = (\langle H_i, H_j \rangle, \pi, \mathcal{N}, T)\,.$$

□

Lemma 7.3.6. *For each $a \in G$,*

$$D_a = \{(H, \pi, \mathcal{N}, T) \mid a \in H\}$$

is a dense subset of $\mathbb{P}$.

Proof. Let $a \in G$ and $p = (H, \pi, \mathcal{N}, T) \in \mathbb{P}$. We can suppose that $a \notin H$. Let $H^+ = \langle H, a \rangle$. Let $C = \{g_i \mid i \in I\}$ be a set of left coset representatives for H in H^+, chosen so that $1 \in C$. Let $\mathcal{N}^+$ be the structure for the language $L(\mathcal{N})$ such that

(a) the universe of $\mathcal{N}^+$ is the cartesian product $C \times \mathcal{N}$; and

(b) for each relation $R \in L(\mathcal{N})$,

$$((g_i, x), (g_j, y)) \in R^{\mathcal{N}^+} \text{ iff } i = j \text{ and } (x, y) \in R^{\mathcal{N}}.$$

By identifying each $x \in \mathcal{N}$ with the element $(1, x) \in \mathcal{N}^+$, we can regard $\mathcal{N}$ as a substructure of $\mathcal{N}^+$. Again it is clear that none of the restricted atomic types in T is realised in $\mathcal{N}^+$.

Define an action of H^+ on $\mathcal{N}^+$ as follows. If $g \in H^+$ and $(g_i, x) \in \mathcal{N}^+$, then

$$g(g_i, x) = (g_j, \pi(h)(x)),$$

where $j \in I$ and $h \in H$ are such that $gg_i = g_j h$. It is easily checked that this action yields a homomorphism $\pi^+ : H^+ \to \mathrm{Aut}(\mathcal{N}^+)$ and that $(H^+, \pi^+, \mathcal{N}^+, T) \leq p$. □

LEMMA 7.3.7. *For each* $\alpha < \kappa$,

$$E_\alpha = \{(H, \pi, \mathcal{N}, T) \mid \alpha \in \mathcal{N}\}$$

is a dense subset of $\mathbb{P}$.

PROOF. Left to the reader. □

Let F be a $\mathbb{P}$-generic filter over V and let $V^{\mathbb{P}} = V[F]$ be the corresponding generic extension. Working within $V^{\mathbb{P}}$, define

$$\mathcal{M} = \bigcup \{\mathcal{N} \mid \text{ There exists } p = (H, \pi, \mathcal{N}, T) \in F\}$$

and

$$\pi^* = \bigcup \{\pi \mid \text{ There exists } p = (H, \pi, \mathcal{N}, T) \in F\}.$$

Then the above lemmas imply that $\mathcal{M}$ is a structure for L of cardinality κ and that π^* is an embedding of G into $\mathrm{Aut}(\mathcal{M})$. So the following lemma completes the proof of Theorem 7.1.6.

LEMMA 7.3.8. $\pi^* : G \to \mathrm{Aut}(\mathcal{M})$ *is a surjective homomorphism.*

PROOF. Suppose that $g \in \mathrm{Aut}(\mathcal{M}) \setminus \pi^*[G]$. Let $\widetilde{\mathcal{M}}$, $\widetilde{\pi}$ be the canonical $\mathbb{P}$-names for $\mathcal{M}$, π^* respectively and let $\widetilde{g}$ be a $\mathbb{P}$-name for g. Then there exists a condition $p \in F$ such that

$$p \Vdash \widetilde{g} \in \mathrm{Aut}(\widetilde{\mathcal{M}}) \text{ and } \widetilde{g} \neq \widetilde{\pi}(h) \text{ for all } h \in G.$$

Let $p' \leq p$. We shall inductively construct a descending sequence of conditions $p_m = (H_m, \pi_m, \mathcal{N}_m, T_m)$ for $m \in \omega$ such that the following hold.

(a) $p_0 = p'$.

(b) For all $x \in \mathcal{N}_m$, there exists $y \in \mathcal{N}_{m+1}$ such that $p_{m+1} \Vdash \widetilde{g}(x) = y$.

(c) For all $h \in H_m$, there exists $z \in \mathcal{N}_{m+1}$ such that

$$p_{m+1} \Vdash \widetilde{g}(z) \neq \pi_{m+1}(h)(z).$$

Suppose that $m \geq 0$ and that $p_m = (H_m, \pi_m, \mathcal{N}_m, T_m)$ has been constructed. Let $\mathcal{N}_m = \{x_\xi \mid \xi < \lambda\}$. Using the fact that $\mathbb{P}$ is κ-closed, we can inductively construct an auxillary descending sequence of conditions $r_\xi = \left(H'_\xi, \pi'_\xi, \mathcal{N}'_\xi, T'_\xi\right)$ for $\xi \leq \lambda$ such that the following hold.

(i) $r_0 = p_m$.

(ii) There exists $y_\xi \in \mathcal{N}'_{\xi+1}$ such that $r_{\xi+1} \Vdash \widetilde{g}(x_\xi) = y_\xi$.

(iii) If δ is a limit ordinal, then r_δ is the greatest lower bound of $\{r_\xi \mid \xi < \delta\}$.

Note that for each $x \in \mathcal{N}_m$, there exists $y \in \mathcal{N}'_\lambda$ such that $r_\lambda \Vdash \widetilde{g}(x) = y$. Using a similar argument, we can construct a condition

$$p_{m+1} = (H_{m+1}, \pi_{m+1}, \mathcal{N}_{m+1}, T_{m+1}) \leq r_\lambda$$

such that for all $h \in H_m$, there exists an element $z \in \mathcal{N}_{m+1}$ such that

$$p_{m+1} \Vdash \widetilde{g}(z) \neq \pi_{m+1}(h)(z).$$

Clearly p_{m+1} satisfies our requirements.

Now let $q = (H, \pi, \mathcal{N}, T)$ be the greatest lower bound of $\{p_m \mid m \in \omega\}$ in $\mathbb{P}$. Then $q \leq p'$ and there exists $g^* \in \operatorname{Aut}(\mathcal{N}) \setminus \pi[H]$ such that $q \Vdash \widetilde{g} \upharpoonright \mathcal{N} = g^*$. (In the remainder of this book, we shall refer to the above argument as the *bootstrap argument*.) Let $\mathcal{N}^+$ be the structure defined as follows.

(1) The universe of $\mathcal{N}^+$ is the disjoint union $\mathcal{N} \sqcup H$.

(2) For each relation $R \in L(\mathcal{N})$, $R^{\mathcal{N}^+} = R^{\mathcal{N}}$.

(3) For each $x \in \mathcal{N}$, let $R_x \in L \setminus L(\mathcal{N})$ be a new binary relation symbol. Then we set $(h, y) \in R_x^{\mathcal{N}^+}$ iff $h \in H$, $y \in \mathcal{N}$ and $\pi(h)(x) = y$.

Once again, it is clear that none of the restricted atomic types in T is realised in $\mathcal{N}^+$. Let $\pi^+ : H \to \operatorname{Sym} \mathcal{N}^+$ be the embedding such that

(i) $\pi^+(h)(x) = \pi(h)(x)$ for all $x \in \mathcal{N}$; and

(ii) $\pi^+(h)(x) = hx$ for all $x \in H$.

Then it is easily checked that $\pi^+[H] \leqslant \mathrm{Aut}(\mathcal{N}^+)$. Finally let t be the partial type defined by

$$t = \{R_x(v, g^*(x)) \mid x \in \mathcal{N}\}$$

and let $T^+ = T \cup \{t\}$.

CLAIM 7.3.9. *t is omitted in $\mathcal{N}^+$.*

PROOF OF CLAIM 7.3.9. Suppose that $h \in \mathcal{N}^+$ realises t. Then we must have that $h \in H$. Since $\mathcal{N}^+ \vDash R_x(h, g^*(x))$ for all $x \in \mathcal{N}$, it follows that $\pi(h)(x) = g^*(x)$ for all $x \in \mathcal{N}$. But this contradicts the fact that $g^* \in \mathrm{Aut}(\mathcal{N}) \smallsetminus \pi[H]$. $\square$

Thus $q^+ = (H, \pi^+, \mathcal{N}^+, T^+) \in \mathbb{P}$ and $q^+ \leq p'$. We have just shown that for each condition $p' \leq p$, there exists a corresponding strengthening $q^+ \leq p'$. In other words, the set of such conditions is dense below p. By Lemma 6.3.1, we can suppose that $q^+ \in F$ and hence that that $g^* \subseteq g$. Clearly $\mathcal{M} \vDash R_x(1, x)$ for each $x \in \mathcal{N}$. So applying the automorphism $g \in \mathrm{Aut}(\mathcal{M})$, we obtain that $\mathcal{M} \vDash R_x(g(1), g^*(x))$ for each $x \in \mathcal{N}$. But this means that $g(1) \in \mathcal{M}$ realises t, which is the final contradiction. $\square$

7.4. τ_κ can be strictly increasing

In this section, we shall prove that it is consistent that τ_κ is a strictly increasing function of κ. Beginning with a ground model V which satisfies GCH, we shall use a suitable reverse Easton forcing to construct a generic extension M in which the following statements are true.

(a) GCH holds.

(b) For each regular uncountable cardinal κ, there exists a centreless group T of cardinality κ such that $\tau(T) = \kappa^+$.

As we explained at the beginning of Section 6.8, it follows easily that τ_κ is strictly increasing in M. During our iteration, for each regular uncountable cardinal κ, we shall use the following result to adjoin a centreless group T of cardinality κ such that $\tau(T) = \kappa^+$.

LEMMA 7.4.1. *Assume GCH. If κ is a regular uncountable cardinal, then there exists a notion of forcing $\mathbb{R}_\kappa$ of cardinality κ^+ such that*

(a) $\mathbb{R}_\kappa$ *is* κ*-closed and has the* κ^+*-c.c.; and*

(b) $\mathbb{R}_\kappa$ *adjoins a centreless group* T *of cardinality* κ *such that* $\tau(T) = \kappa^+$*.*

Proof. Throughout this proof, the ground model will be denoted by N. By Theorem 7.2.11, there exists a notion of forcing $\mathbb{Q}$ of cardinality κ^+ such that

(1) $\mathbb{Q}$ is κ-closed and has the κ^+-*c.c.*; and

(2) in the generic extension $N^{\mathbb{Q}}$, there exist groups $H \leqslant G < \mathrm{Sym}(\kappa)$ such that $|H| = \kappa$ and the normaliser tower of H in G terminates after exactly κ^+ steps.

It is easily checked that $N^{\mathbb{Q}} \vDash GCH$. By Proposition 7.1.5, $G \in N^{\mathbb{Q}}$ satisfies the κ^+-compatibility condition.

Let L be a first order language consisting of κ binary relation symbols. By Theorem 7.1.6, there exists a notion of forcing $\mathbb{P} \in N^{\mathbb{Q}}$ of cardinality κ^+ such that

(i) $\mathbb{P}$ is κ-closed and has the κ^+-*c.c.*; and

(ii) $\mathbb{P}$ adjoins an L-structure $\mathcal{M}$ of cardinality κ such that $G \simeq \mathrm{Aut}(\mathcal{M})$.

By Theorem 4.1.9, there exists a centreless group $T \in (N^{\mathbb{Q}})^{\mathbb{P}}$ of cardinality κ such that $\tau(T) = \kappa^+$. Applying Theorem 6.8.4, it follows easily that if $\widetilde{\mathbb{P}}$ is a $\mathbb{Q}$-name of the notion of forcing $\mathbb{P} \in N^{\mathbb{Q}}$, then $\mathbb{R}_\kappa = \mathbb{Q} * \widetilde{\mathbb{P}}$ satisfies our requirements. □

Now let $V \vDash GCH$. We shall inductively construct a sequence $\langle \mathbb{Q}_\beta \mid \beta \in On \rangle$ of forcing notions satisfying Hypothesis 6.8.9. By the remark following Theorem 6.8.10, at successor stages β of the construction, we can assume inductively that $V^{\mathbb{Q}_{\beta-1}} \vDash GCH$.

Case 1. If $\beta = 0$, then $\mathbb{Q}_0 = \{\emptyset\}$ is the trivial notion of forcing.

Case 2. If β is a limit ordinal which is not inaccessible, then $\mathbb{Q}_\beta$ is the inverse limit of $\langle \mathbb{Q}_\gamma \mid \gamma < \beta \rangle$.

Case 3. If β is inaccessible, then $\mathbb{Q}_\beta$ is the direct limit of $\langle \mathbb{Q}_\gamma \mid \gamma < \beta \rangle$.

Case 4. Finally suppose that $\beta = \gamma + 1$ is a successor ordinal. First suppose that $\gamma = \kappa$ is a regular uncountable cardinal. Then we can assume inductively that $V^{\mathbb{Q}_\kappa} \vDash GCH$. Let $\mathbb{R}_\kappa \in V^{\mathbb{Q}_\kappa}$ be the notion of forcing, given by Lemma 7.4.1, which adjoins a centreless group T of cardinality κ such that $\tau(T) = \kappa^+$. Then we set $\mathbb{Q}_{\kappa+1} = \mathbb{Q}_\kappa * \widetilde{\mathbb{R}}_\kappa$. (As usual, $\widetilde{\mathbb{R}}_\kappa$ denotes a $\mathbb{Q}_\kappa$-name of the notion of forcing $\mathbb{R}_\kappa \in V^{\mathbb{Q}_\kappa}$.) Finally if γ is not a regular uncountable cardinal, then we set $\mathbb{Q}_{\gamma+1} = \mathbb{Q}_\gamma * \check{\mathbb{Q}}_0$.

Let $\mathbb{Q}_\infty$ be the direct limit of $\langle \mathbb{Q}_\beta \mid \beta \in On \rangle$; and for each $\beta \in On$, let $\tilde{\mathbb{Q}}_{\beta\infty}$ be the canonically chosen $\mathbb{Q}_\beta$-name for a proper class notion of forcing such that $\mathbb{Q}_\infty$ is isomorphic to a dense sub-order of $\mathbb{Q}_\beta * \tilde{\mathbb{Q}}_{\beta\infty}$. Let H be a $\mathbb{Q}_\infty$-generic filter over V and let $M = V[H]$ be the corresponding generic extension. For each $\beta \in On$, let $H_\beta = H \cap \mathbb{Q}_\beta$ and let $\mathbb{Q}_{\beta\infty} = (\tilde{\mathbb{Q}}_{\beta\infty})_{H_\beta}$. Then the following result is an immediate consequence of Theorem 6.8.10.

Lemma 7.4.2. (a) *$\mathbb{Q}_\infty$ preserves cofinalities and cardinals.*

(b) *If κ is a regular uncountable cardinal, then $V[H_{\kappa+1}] \vDash \mathbb{Q}_{\kappa+1\infty}$ is κ^+-closed.*

(c) *M is a model of $ZFC + GCH$.*

□

Now let $\kappa \in M$ be any regular uncountable cardinal. By construction, there exists a centreless group $T \in V^{\mathbb{Q}_{\kappa+1}}$ of cardinality κ such that $V^{\mathbb{Q}_{\kappa+1}} \vDash \tau(T) = \kappa^+$. Since $V[H_{\kappa+1}] \vDash \mathbb{Q}_{\kappa+1\infty}$ is κ^+-closed, Theorem 6.7.1 implies that $M \vDash \tau(T) = \kappa^+$. This completes the proof of the following result.

Theorem 7.4.3. *It is consistent with ZFC that τ_κ is a strictly increasing function of κ.*

□

It seems almost certain that it is also consistent that τ_κ is not a strictly increasing function of κ. However, as we shall explain in Section 7.6, this seems to be a much more difficult problem.

Conjecture 7.4.4. It is consistent with ZFC that $\tau_{\omega_1} = \tau_{\omega_2}$.

7.5. Two more applications

In this section, we shall present two more applications of Theorem 7.1.6. In both applications, we shall be working with a group $G \in V$ which satisfies the κ^+-compatibility condition for some uncountable regular cardinal such that $\kappa^{<\kappa} = \kappa$ and we shall use Theorem 7.1.6 to generically adjoin a graph $\Gamma \in V^{\mathbb{P}}$ of cardinality κ such that $G \simeq \mathrm{Aut}(\Gamma)$. Of course, the actual statement of Theorem 7.1.6 only gives a structure $\mathcal{M} \in V^{\mathbb{P}}$ of cardinality κ for a binary relational language L of cardinality κ such that $G \simeq \mathrm{Aut}(\mathcal{M})$. But this is no problem. For then we can use Theorem 4.1.8 to find a graph $\Gamma \in V^{\mathbb{P}}$ of cardinality κ such that $\mathrm{Aut}(\mathcal{M}) \simeq \mathrm{Aut}(\Gamma)$.

APPLICATION 7.5.1. In 1961, Dudley [**5**] proved that there does not exist an uncountable free Polish group. Consequently, since the automorphism group of a countable structure is a Polish group, it follows that there does not exist a countable structure $\mathcal{M}$ such that $\mathrm{Aut}(\mathcal{M})$ is the free group on 2^{ω} generators. (This result was independently rediscovered by Shelah [**45**].) In contrast, using Theorem 7.1.6, it is easy to establish the consistency of the existence of a structure $\mathcal{N}$ of cardinality ω_1 such that $\mathrm{Aut}(\mathcal{N})$ is the free group on 2^{ω_1} generators. It is not known whether the existence of such a structure can be proved in ZFC.

THEOREM 7.5.2. *Suppose that κ, λ, $\theta \in V$ are uncountable cardinals such that $\kappa^{<\kappa} = \kappa < \lambda \leq \theta = \theta^{\kappa}$. Then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that the following statements are true in $V^{\mathbb{P}}$.*

(a) *$2^{\kappa} = \theta$; and*

(b) *there exists a graph Γ of cardinality κ such that $\mathrm{Aut}(\Gamma)$ is the free group on λ generators.*

We shall make use of the following result, which is due to de Bruijn [**3**].

LEMMA 7.5.3. *Let κ be an infinite cardinal and let $F_{2^{\kappa}}$ be the free group on 2^{κ} generators. Then there exists an embedding of $F_{2^{\kappa}}$ into $\mathrm{Sym}(\kappa)$.*

PROOF. First let $\prod_{\alpha<\kappa} \mathrm{Sym}(\kappa)$ be the full direct product of κ copies of $\mathrm{Sym}(\kappa)$. Then $\prod_{\alpha<\kappa} \mathrm{Sym}(\kappa)$ embeds into $\mathrm{Sym}(\kappa)$. To see this, express $\kappa = \bigsqcup_{\alpha<\kappa} \Delta_{\alpha}$ as a disjoint union such that $|\Delta_{\alpha}| = \kappa$ for all $\alpha < \kappa$ and note that the full direct product $\prod_{\alpha<\kappa} \mathrm{Sym}(\Delta_{\alpha})$ can be regarded as a subgroup of $\mathrm{Sym}(\kappa)$.

Clearly the free group F_{κ} on κ generators embeds into $\mathrm{Sym}(F_{\kappa}) \simeq \mathrm{Sym}(\kappa)$. So $\prod_{\alpha<\kappa} F_{\kappa}$ embeds into $\prod_{\alpha<\kappa} \mathrm{Sym}(\kappa)$ and hence into $\mathrm{Sym}(\kappa)$. Thus it is enough to show that $F_{2^{\kappa}}$ embeds into $\prod_{\alpha<\kappa} F_{\kappa}$. To accomplish this, we shall show that there exists a set $\{\varphi_{\alpha} \mid \alpha < \kappa\}$ of homomorphisms $\varphi_{\alpha} : F_{2^{\kappa}} \to F_{\kappa}$ such that $\bigcap\{\ker \varphi_{\alpha} \mid \alpha < \kappa\} = 1$. Then we can define an embedding $\varphi : F_{2^{\kappa}} \to \prod_{\alpha<\kappa} F_{\kappa}$ by $\varphi(w) = \langle \varphi_{\alpha}(w) \mid \alpha < \kappa \rangle$.

Let F be the group which is freely generated by the set $\{x_A \mid A \subseteq \kappa\}$ and let G be the group which is freely generated by the set $\{y_B \mid B \in [\kappa]^{<\omega}\}$. Then $F \simeq F_{2^\kappa}$ and $G \simeq F_\kappa$. For each $C \in [\kappa]^{<\omega}$, let $\pi_C : F \to G$ be the homomorphism such that $\pi_C(x_A) = y_{A \cap C}$ for all $A \subseteq \kappa$. By the previous paragraph , it is enough to show that $\bigcap\{\ker \pi_C \mid C \in [\kappa]^{<\omega}\} = 1$. To see this, suppose that $1 \neq w \in F$; say, $w \in \langle x_{A_1}, \dots, x_{A_n}\rangle$. Then there exists $C \in [\kappa]^{<\omega}$ such that $A_i \cap C \neq A_j \cap C$ for all $1 \leq i < j \leq n$. Clearly $\pi_C \restriction \langle x_{A_1}, \dots, x_{A_n}\rangle$ is an embedding and so $\pi_C(w) \neq 1$. $\square$

Proof of Theorem 7.5.2. First we shall perform a preliminary forcing to obtain a c.t.m. M in which $\kappa^{<\kappa} = \kappa$ and $2^\kappa = \theta$. Let $\mathbb{R} = \mathrm{Fn}(\theta \times \kappa, 2, \kappa)$. By Theorem 6.4.4, $\mathbb{R}$ is κ-closed and has the κ^+-*c.c.* Thus $\mathbb{R}$ preserves cofinalities and cardinals. Let H be an $\mathbb{R}$-generic filter over $V$and let $M = V[H]$ be the corresponding generic extension. By Theorem 6.4.4, $M \models 2^\kappa = \theta$. Since $\mathbb{R}$ is κ-closed, it follows that $\mathbb{R}$ does not adjoin any new subsets S of κ with $|S| < \kappa$ and so $M \models \kappa^{<\kappa} = \kappa$.

From now on, we shall work within M. Let G be the free group on λ generators. By Lemma 7.5.3, there exists an embedding of G into $\mathrm{Sym}(\kappa)$ and so F satisfies the κ^+-compatibility condition. Let $\mathbb{P}$ be the notion of forcing, given by Theorem 7.1.6, which adjoins a graph Γ of cardinality κ such that $G \simeq \mathrm{Aut}(\Gamma)$. Then the notion of forcing $\mathbb{R} * \tilde{\mathbb{P}} \in V$ satisfies our requirements. (Here $\tilde{\mathbb{P}}$ denotes an $\mathbb{R}$-name of the notion of forcing $\mathbb{P} \in M = V^{\mathbb{R}}$.) $\square$

Application 7.5.4. Theorem 3.4.1 says that if G is a finitely generated centreless group, then the automorphism tower of G terminates after countably many steps. It is conceivable that a much more general result holds; namely, that the automorphism tower of G terminates after countably many steps, whenever G is a countable centreless group such that $\mathrm{Aut}(G)$ is also countable. To see why this might be true, let G be such a group. Then, by Kueker [**24**], there exists a finite subset $F \subseteq G$ such that each automorphism $\pi \in \mathrm{Aut}(G)$ is uniquely determined by its restriction $\pi \restriction F$. In terms of the automorphism tower of G, this says that there is a finite subset $F \subseteq G$ such that $C_{G_1}(F) = 1$. Suppose that the “rigidity” of F within $G = G_0$ is propagated along the automorphism tower of G; i.e. that $C_{G_\alpha}(F) = 1$ for all ordinals α. Then a routine modification of the proof of Theorem 3.4.1 shows that the automorphism tower of G terminates in countably many steps.

QUESTION 7.5.5. Let G be a centreless group such that $|\mathrm{Aut}(G)| = \omega$. Does there exist a finite subset $F \subseteq G$ such that $C_{G_\alpha}(F) = 1$ for all ordinals α?

The following weak form of Question 7.5.5 is also open.

QUESTION 7.5.6. Does there exist a centreless group G such that $|\mathrm{Aut}(G)| = \omega$ and $|\mathrm{Aut}(\mathrm{Aut}(G))| = 2^\omega$?

Of course, a positive answer to Question 7.5.5 implies a negative answer to Question 7.5.6. Using Theorem 7.1.6, it is easy to establish the consistency of the existence of a centreless group G of uncountable cardinality κ such that $|\mathrm{Aut}(G)| = \kappa$ and $|\mathrm{Aut}(\mathrm{Aut}(G))| = 2^\kappa$. Once again, it is not known whether the existence of such a group can be proved in ZFC.

THEOREM 7.5.7. *Suppose that $\kappa \in V$ is a regular uncountable cardinal such that $\kappa^{<\kappa} = \kappa$. Then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, which adjoins a centreless group G of cardinality κ such that*

(a) $|\mathrm{Aut}(G)| = \kappa$*; and*
(b) $|\mathrm{Aut}(\mathrm{Aut}(G))| = 2^\kappa$.

PROOF. For each α, $\xi < \kappa$, let $Z_\xi^\alpha = \langle z_\xi^\alpha \rangle$ be an infinite cyclic group. For each $\xi < \kappa$, let $A_\xi = \bigoplus_{\alpha<\kappa} Z_\xi^\alpha$ and let $B = \bigoplus_{\xi<\kappa} A_\xi$. Define an action of $\mathrm{Sym}(\kappa)$ on B by $\pi z_\xi^\alpha \pi^{-1} = z_{\pi(\xi)}^\alpha$ for all α, $\xi < \kappa$ and let $W = B \rtimes \mathrm{Sym}(\kappa)$ be the corresponding semidirect product. Let $H = \bigoplus_{\xi<\kappa} Z_\xi^\xi$. Then the members of the normaliser tower of H in W are

(a) $N_0(H) = H$;
(b) $N_1(H) = B$;
(c) $N_2(H) = W$.

Clearly W is embeddable in $\mathrm{Sym}(\kappa)$ and so W satisfies the κ^+-compatibility condition. Let $\mathbb{P}$ be the notion of forcing, given by Theorem 7.1.6, which adjoins a graph Γ of cardinality κ such that $W \simeq \mathrm{Aut}(\Gamma)$. Let $K_\Gamma \in V^{\mathbb{P}}$ be the corresponding field, which is given by Theorem 4.1.7. Then $G = PGL(2, K_\Gamma) \rtimes H \in V^{\mathbb{P}}$ is a group such that

$$|\mathrm{Aut}(G)| = |PGL(2, K_\Gamma) \rtimes B| = \kappa$$

and

$$|\mathrm{Aut}(\mathrm{Aut}(G))| = |PGL(2, K_\Gamma) \rtimes W| = 2^\kappa.$$

□

7.6. The main gap

A common feature of the consistency results in this chapter was that, in each case, it was only necessary to generically adjoin a single centreless group. For example, to prove the consistency of $\tau_{\omega_1} < \tau_{\omega_2}$, it was enough to adjoin a centreless group G of cardinality ω_2 such that $\tau(G) \geq (2^{\omega_1})^+$. On the other hand, if we wish to prove the consistency of $\tau_{\omega_1} = \tau_{\omega_2}$, then we must construct a generic extension in which we have some understanding of every centreless group G such that $\omega_1 \leq |G| \leq \omega_2$. (Of course, a similar remark applies to the problem of proving the consistency of a statement such as $\tau_{\omega_1} = \omega_3$, etc.)

In order to understand the difficulties involved, consider the following plausible approach to proving the consistency of $\tau_{\omega_1} = \tau_{\omega_2}$. Let $V \vDash GCH$ and let $(\mathbb{P}_\beta \mid \beta \leq \omega_4)$ be the countable support iteration such that

- if $\beta = \alpha + 1$ is a successor ordinal, then $\mathbb{P}_\beta = \mathbb{P}_\alpha * \widetilde{\mathbb{Q}}_\alpha$, where $\mathbb{Q}_\alpha \in V^{\mathbb{P}_\alpha}$ is the obvious notion of forcing which adjoins a centreless group $G^{(\alpha)}$ of cardinality ω_1 such that $\tau(G^{(\alpha)}) = \alpha$.

(I.e., $\mathbb{Q}_\alpha$ is the notion of forcing obtained by iterating those introduced in Sections 7.2 and 7.3.) Then it can be shown that the following statements hold in the resulting generic extension $M = V^{\mathbb{P}_{\omega_4}}$.

(a) $2^{\omega_1} = 2^{\omega_2} = \omega_4$.
(b) For each $\alpha < \omega_4$, $\tau(G^{(\alpha)}) = \alpha$.

It is natural to conjecture that $M \vDash \tau_{\omega_1} = \tau_{\omega_2} = \omega_4$. Consider an arbitrary centreless group $G \in M$ such that $\omega_1 \leq |G| \leq \omega_2$. Then there exists an ordinal $\beta < \omega_4$ such that $G \in V^{\mathbb{P}_\beta}$; and since $V^{\mathbb{P}_\beta} \vDash 2^{\omega_2} = \omega_3$, we have that

$$V^{\mathbb{P}_\beta} \vDash \tau(G) < (2^{\omega_2})^+ = \omega_4.$$

Unfortunately, in Section 6.7, we saw that the height $\tau(G)$ is not an upwards absolute concept and so it remains conceivable that $M \vDash \tau(G) \geq \omega_4$.

In Chapter 8, we shall prove that if G is an infinite centreless group, then it is impossible to prove any nontrivial results concerning the relationship between the values of $\tau(G)$ in the ground model and an arbitrary generic extension. More precisely, we shall prove that it is consistent that for every infinite cardinal λ and every ordinal $\alpha < \lambda$, there exists a centreless group G with the following properties.

(a) $\tau(G) = \alpha$.

(b) If β is any ordinal such that $1 \leq \beta < \lambda$, then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that $\tau^{V^{\mathbb{P}}}(G) = \beta$.

However, as the reader might expect, both the groups G and the notions of forcing $\mathbb{P}$ are specifically designed to witness the extreme nonabsoluteness of the height function. It should also be stressed that the main theorem of Chapter 8 is a consistency result and it is open whether such groups and notions of forcing can be proved to exist in ZFC. In particular, it remains unclear whether "naturally occurring" notions of forcing can significantly change the heights of pre-existing centreless groups. In Chapter 9, we shall study this question for the very simplest notions of forcing. More specifically, suppose that $V \vDash GCH$ and that $\kappa \leq \lambda < \theta$ are regular cardinals. Then we shall show that if $\mathbb{P} = \mathrm{Fn}(\theta, 2, \kappa)$ and $\mathbb{Q} = \mathrm{Fn}(\lambda^+, 2, \kappa)$, then for every centreless group $G \in V$ of cardinality λ,

$$\tau^{V^{\mathbb{P}}}(G) = \tau^{V^{\mathbb{Q}}}(G).$$

Since $V^{\mathbb{Q}} \vDash 2^\lambda = \lambda^+$, this implies that

$$V^{\mathbb{P}} \vDash \tau(G) < \lambda^{++}$$

for every centreless group $G \in V$ of cardinality λ. Hence if we choose $\theta \geq \lambda^{++}$, then we obtain that

$$V^{\mathbb{P}} \vDash \tau(G) < \lambda^{++} \leq \theta = 2^\lambda$$

for every centreless group $G \in V$ of cardinality λ. With a little more effort, we shall then prove that it is consistent that $\tau_\lambda < 2^\lambda$ for all regular cardinals λ.

7.7. Notes

The material in this chapter first appeared in Just-Shelah-Thomas [**22**]. The problem of determining which groups embed into $\mathrm{Sym}(\kappa)$ and its homomorphic images has been studied by many authors, including de Bruijn [**3**], McKenzie [**30**], Shelah [**43**], Clare [**4**], Rabinovič [**37**] and Felgner-Haug [**8**]. In particular, Felgner-Haug [**8**] proved in ZFC that if κ is a regular cardinal, then $\mathrm{Sym}_{\kappa^+}(\kappa^+)$ can be embedded in $\mathrm{Sym}(\kappa)/\mathrm{Sym}_\kappa(\kappa)$. Of course, this implies that both $\mathrm{Alt}(\kappa^+)$ and W_{κ^+} can be embedded in $\mathrm{Sym}(\kappa)/\mathrm{Sym}_\kappa(\kappa)$.

CHAPTER 8

Changing The Heights Of Automorphism Towers

In Section 6.7, we saw that the height $\tau(G)$ of the automorphism tower of an infinite centreless group G is not an absolute concept. For example, if $\mathcal{U}$ is a nonprincipal ultrafilter on ω, then the setwise stabiliser $S_{\mathcal{U}}$ is a complete group, but there exists a *c.c.c.* notion of forcing $\mathbb{P}$ which adjoins an outer automorphism to $S_{\mathcal{U}}$. Thus $0 = \tau(S_{\mathcal{U}}) < \tau^{V^{\mathbb{P}}}(S_{\mathcal{U}})$. Perhaps more surprisingly, there also exists an infinite centreless group G such that $\tau(G) = 2$ and such that if $\mathbb{Q}$ is any notion of forcing which adjoins a new real, then $\tau^{V^{\mathbb{Q}}}(G) = 1$. Thus the height $\tau(G)$ of the automorphism tower of an infinite centreless group G may either increase or decrease in a generic extension. In fact, if G is an infinite centreless group and $\mathbb{P}$ is a notion of forcing, then it is difficult to think of any nontrivial statement concerning the relationship between $\tau(G)$ and $\tau^{V^{\mathbb{P}}}(G)$. Of course, there is a trivial observation which can be made; namely, that since an outer automorphism of G remains an outer automorphism in $V^{\mathbb{P}}$, $\tau(G) \geq 1$ implies that $\tau^{V^{\mathbb{P}}}(G) \geq 1$. In this chapter, we shall prove that it is consistent that for every infinite cardinal λ and every ordinal $\alpha < \lambda$, there exists a centreless group G with the following properties.

(a) $\tau(G) = \alpha$.
(b) If β is any ordinal such that $1 \leq \beta < \lambda$, then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that $\tau^{V^{\mathbb{P}}}(G) = \beta$.

This result suggests that no nontrivial results concerning the relationship between $\tau(G)$ and $\tau^{V^{\mathbb{P}}}(G)$ are provable in ZFC. However, it should be pointed out that the cardinality of G depends on λ; and it remains an open question whether it is consistent that there exists a fixed centreless group G such that for unboundedly many ordinals α, there exists a cardinal-preserving notion of forcing $\mathbb{P}$ such that $\tau^{V^{\mathbb{P}}}(G) = \alpha$.

8.1. Changing the heights of automorphism towers

In Sections 8.2 and 8.3, we shall present a proof of the following theorem.

THEOREM 8.1.1. *It is consistent that for every infinite cardinal λ and every ordinal $\alpha < \lambda$, there exists a centreless group G with the following properties.*

(a) $\tau(G) = \alpha$.

(b) *If β is any ordinal such that $1 \le \beta < \lambda$, then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that $\tau^{V^{\mathbb{P}}}(G) = \beta$.*

In this section, we shall give a heuristic account of the main ideas of the proof of Theorem 8.1.1. For the sake of concreteness, suppose that we wish to construct a centreless group G with the following properties.

(a) $\tau(G) = 3$.

(b) There exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that $\tau^{V^{\mathbb{P}}}(G) = 4$.

(c) There exists a notion of forcing $\mathbb{Q}$, which preserves cofinalities and cardinals, such that $\tau^{V^{\mathbb{Q}}}(G) = 2$.

Applying Theorems 4.1.6 and 6.6.18, it is enough to find a graph Γ, together with a subgroup $H \leqslant \operatorname{Aut}(\Gamma)$, such that the following conditions hold. (We shall spell out the details of this reduction in the proof of Corollary 8.2.3.)

$(a)'$ The normaliser tower of H in $\operatorname{Aut}(\Gamma)$ terminates after exactly 3 steps.

$(b)'$ There exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that the normaliser tower of H in $\operatorname{Aut}^{V^{\mathbb{P}}}(\Gamma)$ terminates after exactly 4 steps.

$(c)'$ There exists a notion of forcing $\mathbb{Q}$, which preserves cofinalities and cardinals, such that the normaliser tower of H in $\operatorname{Aut}^{V^{\mathbb{Q}}}(\Gamma)$ terminates after exactly 2 steps.

Let P, Q and R be unary relation symbols and let L be the first-order language $\{P, Q, R\}$. As a first approximation to the graph Γ, we shall take the structure

$$\mathcal{N} = \langle N; \mathcal{N}^P, \mathcal{N}^Q, \mathcal{N}^R \rangle$$

for the language L such that $|\mathcal{N}^P| = |\mathcal{N}^Q| = |\mathcal{N}^R| = 8$. Thus

$$\mathrm{Aut}(\mathcal{N}) = \mathrm{Sym}(\mathcal{N}^P) \times \mathrm{Sym}(\mathcal{N}^Q) \times \mathrm{Sym}(\mathcal{N}^R).$$

(At first glance, it might seem strange to approximate the desired graph Γ by the finite structure $\mathcal{N}$, since the automorphism group of $\mathcal{N}$ will obviously remain unchanged in every generic extension of the universe. However, this objection only applies if we intend to code $\mathcal{N}$ within a finite graph. Instead we shall take Γ to be a suitable infinite graph which encodes $\mathcal{N}$.) We shall define the subgroup $H \leqslant \mathrm{Aut}(\mathcal{N})$ to be a suitable product of iterated wreath products. But first we need to define the general notion of the direct product of a collection of permutation groups.

Definition 8.1.2. Suppose that $G_i \leqslant \mathrm{Sym}(\Omega_i)$ for each $i \in I$. Then the *direct product* of the permutation groups $\{(G_i, \Omega_i) \mid i \in I\}$ is defined to be the permutation group

$$\prod_{i \in I} (G_i, \Omega_i) = \left(\prod_{i \in I} G_i, \bigsqcup_{i \in I} \Omega_i \right).$$

Here $\bigsqcup_{i \in I} \Omega_i$ denotes the *disjoint union* of the sets $\{\Omega_i \mid i \in I\}$ and $\prod_{i \in I} G_i$ acts on $\bigsqcup_{i \in I} \Omega_i$ in the natural manner; i.e. if $\pi = (g_i) \in \prod_{i \in I} G_i$, then $\pi \restriction \Omega_i = g_i$ for each $i \in I$. If $I = \{1, 2\}$, then we write

$$\prod_{i \in I} (G_i, \Omega_i) = (G_1, \Omega_1) \times (G_2, \Omega_2) = (G_1 \times G_2, \Omega_1 \sqcup \Omega_2).$$

Next we define the natural action of the nth iterated wreath product

$$P_n = \underbrace{(\dots (C_2 \mathrm{\ Wr\ } C_2) \mathrm{\ Wr\ } C_2) \mathrm{\ Wr\ } \cdots \mathrm{\ Wr\ } C_2) \dots)}_{n \text{ times}}$$

on the set $\Sigma_n = \{\ell \mid 0 \leq \ell < 2^n\}$ inductively as follows.

- $P_0 = 1$ acts trivially on $\Sigma_0 = \{0\}$.
- Suppose that the permutation group (P_n, Σ_n) has been defined. Then $P_{n+1} = [P_n \times P_n] \rtimes \langle \sigma_n \rangle$ is the wreath product P_n Wr C_2 acting naturally on $\Sigma_{n+1} = \Sigma_n \sqcup \Sigma'_n$, where $\Sigma'_n = \{\ell \mid 2^n \leq \ell < 2^{n+1}\}$. More explicitly,

  $$(P_n \times P_n, \Sigma_n \sqcup \Sigma'_n) \simeq (P_n, \Sigma_n) \times (P_n, \Sigma_n)$$

  and σ_n is an involution which interchanges the sets Σ_n and Σ'_n.

An easy induction shows that if $n \geq 1$, then $|P_n| = 2^{2^n-1}$ and that 2^{2^n-1} is the highest power of 2 which divides $(2^n)! = |\operatorname{Sym}(\Sigma_n)|$. Thus P_n is a Sylow 2-subgroup of $\operatorname{Sym}(\Sigma_n)$. It is also reasonably straightforward to show that P_n is self-normalising in $\operatorname{Sym}(\Sigma_n)$. (A generalisation of this result is included within the proof of Lemma 8.2.6.)

We are now ready to begin the definition of the subgroup

$$H \leqslant \operatorname{Aut}(\mathcal{N}) = \operatorname{Sym}(\mathcal{N}^P) \times \operatorname{Sym}(\mathcal{N}^Q) \times \operatorname{Sym}(\mathcal{N}^R).$$

First let

$$\begin{aligned}
\left(H^P, \mathcal{N}^P\right) &\simeq (P_0, \Sigma_0) \times (P_0, \Sigma_0) \times (P_1, \Sigma_1) \times (P_2, \Sigma_2) \\
\left(H^Q, \mathcal{N}^Q\right) &\simeq (P_3, \Sigma_3) \\
\left(H^R, \mathcal{N}^R\right) &\simeq (P_2, \Sigma_2) \times (P_2, \Sigma_2).
\end{aligned}$$

Then we define

$$\left(H, \mathcal{N}^P \sqcup \mathcal{N}^Q \sqcup \mathcal{N}^R\right) \simeq \left(H^P, \mathcal{N}^P\right) \times \left(H^Q, \mathcal{N}^Q\right) \times \left(H^R, \mathcal{N}^R\right).$$

At this point, it is helpful to introduce a pictorial representation of the above permutation groups and the groups which occur in their normaliser towers. It will be much clearer if we illustrate the representation with a couple of examples, instead of attempting to give a formal definition. As a first example, the permutation group $\left(H^P, \mathcal{N}^P\right)$ will be represented by the diagram:

$$\boxed{\blacktriangle}\ \boxed{\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}$$

Here the triangles represent the elements of $\mathcal{N}^P$; and the 4 boxes represent the 4 orbits of H^P on $\mathcal{N}^P$, on each of which H^P acts as a suitable iterated wreath product. By Lemma 1.3.15, if (G, Ω) is a permutation group and $\pi \in \operatorname{Sym}(\Omega)$ normalises G, then π permutes the orbits of G. Using this observation, together with the fact that P_n is self-normalising in $\operatorname{Sym}(\Sigma_n)$ for each $0 \leq n \leq 3$, we see that the normaliser tower of H^P of $\operatorname{Sym}(\mathcal{N}^P)$ is represented by the following sequence of diagrams:

$$\boxed{\blacktriangle}\ \boxed{\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}$$

$$\boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}$$

$$\boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}$$

$$\boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}$$

In order to indicate that each of the sets $\mathcal{N}^P$, $\mathcal{N}^Q$ and $\mathcal{N}^R$ must be preserved setwise by the elements of $\mathrm{Aut}(\mathcal{N})$, we shall represent the elements of $\mathcal{N}^Q$ by stars and the elements of $\mathcal{N}^R$ by diamonds. Thus the following diagram represents the permutation group $\big(H, \mathcal{N}^P \sqcup \mathcal{N}^Q \sqcup \mathcal{N}^R\big)$:

$$\boxed{\blacktriangle}\ \boxed{\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

and the normaliser tower of H in $\mathrm{Aut}(\mathcal{N})$ is represented by the following sequence of diagrams:

$$\boxed{\blacktriangle}\ \boxed{\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

$$\boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

$$\boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

$$\boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

In particular, the normaliser tower of H in $\mathrm{Aut}(\mathcal{N})$ terminates after exactly 3 steps.

Now imagine that $\mathbb{P}$ is a miraculous notion of forcing which manages to convert each star into a triangle. Then H is represented by the following diagram in $V^{\mathbb{P}}$:

$$\boxed{\blacktriangle}\ \boxed{\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

and the normaliser tower of H in $\mathrm{Aut}^{V^{\mathbb{P}}}(\mathcal{N}) = \mathrm{Sym}(\mathcal{N}^P \sqcup \mathcal{N}^Q) \times \mathrm{Sym}(\mathcal{N}^R)$ is represented by the following sequence of diagrams:

$$\boxed{\blacktriangle}\ \boxed{\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

$$\boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

$$\boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

$$\boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

$$\boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge\blacklozenge}$$

Hence the normaliser tower of H in $\mathrm{Aut}^{V^{\mathbb{P}}}(\mathcal{N})$ terminates after exactly 4 steps.

Finally imagine that $\mathbb{Q}$ is an equally miraculous notion of forcing which converts each diamond into a triangle. Then H is represented by the following diagram in $V^{\mathbb{Q}}$:

$$\boxed{\blacktriangle}\ \boxed{\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar}$$

and the normaliser tower of H in $\mathrm{Aut}^{V^{\mathbb{Q}}}(\mathcal{N}) = \mathrm{Sym}(\mathcal{N}^P \sqcup \mathcal{N}^R) \times \mathrm{Sym}(\mathcal{N}^Q)$ is represented by the following sequence of diagrams:

$$\boxed{\blacktriangle}\ \boxed{\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar}$$

$$\boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar}$$

$$\boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle\blacktriangle}\ \boxed{\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar\bigstar}$$

Hence the normaliser tower of H in $\mathrm{Aut}^{V^{\mathbb{Q}}}(\mathcal{N})$ terminates after exactly 2 steps. Of course, the induced action on the orbit of size 12 is not isomorphic to one of the form (P_n, Σ_n) for any $n \geq 0$. Instead the induced action is isomorphic to the natural action of $P_2 \operatorname{Wr} \mathrm{Sym}(3)$ on $\Sigma_2 \sqcup \Sigma_2 \sqcup \Sigma_2$.

While the existence of the notions of forcing $\mathbb{P}$ and $\mathbb{Q}$ is a pure fantasy, it turns out to be possible to convert the above argument into a valid proof. We simply replace each point of $\mathcal{N}$ by a suitably chosen connected rigid graph. (Recall that a structure $\mathcal{M}$ is said to be rigid iff the identity map $id_{\mathcal{M}}$ is the only automorphism of $\mathcal{M}$.) In more detail, let Γ_1, Γ_2 and Γ_3 be pairwise nonisomorphic connected rigid graphs and let Γ be the graph consisting of 8 copies of Γ_i for each $1 \leq i \leq 3$. Let

$$\Phi = \Phi_1 \sqcup \Phi_2 \sqcup \Phi_3$$

be the set of connected components of Γ, where each Φ_i consists of the 8 copies of Γ_i. Then $\mathrm{Aut}(\Gamma)$ can be naturally identified with the group

$$\mathrm{Sym}(\Phi_1) \times \mathrm{Sym}(\Phi_2) \times \mathrm{Sym}(\Phi_3) \simeq \mathrm{Aut}(\mathcal{N}).$$

Using this identification, we can regard H as a subgroup of $\mathrm{Aut}(\Gamma)$; and we have already seen that:

- $(a)''$ The normaliser tower of H in $\mathrm{Sym}(\Phi_1) \times \mathrm{Sym}(\Phi_2) \times \mathrm{Sym}(\Phi_3)$ terminates after exactly 3 steps;
- $(b)''$ The normaliser tower of H in $\mathrm{Sym}(\Phi_1 \cup \Phi_2) \times \mathrm{Sym}(\Phi_3)$ terminates after exactly 4 steps; and
- $(c)''$ The normaliser tower of H in $\mathrm{Sym}(\Phi_1 \cup \Phi_3) \times \mathrm{Sym}(\Phi_2)$ terminates after exactly 2 steps.

So the pair $H \leqslant \mathrm{Aut}(\Gamma)$ will satisfy conditions $(a)'$, $(b)'$ and $(c)'$ provided that the rigid connected graphs Γ_1, Γ_2 and Γ_3 can also be chosen to satisfy the following properties.

$(b)'''$ There exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that the following statements are true in $V^{\mathbb{P}}$:

(i) $\Gamma_1 \simeq \Gamma_2$.

(ii) $\Gamma_1 \not\simeq \Gamma_3$.

(iii) Γ_1, Γ_2 and Γ_3 remain rigid.

$(c)'''$ There exists a notion of forcing $\mathbb{Q}$, which preserves cofinalities and cardinals, such that the following statements are true in $V^{\mathbb{Q}}$:

(i) $\Gamma_1 \simeq \Gamma_3$.

(ii) $\Gamma_1 \not\simeq \Gamma_2$.

(iii) Γ_1, Γ_2 and Γ_3 remain rigid.

Of course, it is far from clear that such graphs exist. For example, since Γ_1 and Γ_2 remain rigid, there must exist a unique isomorphism $\pi : \Gamma_1 \to \Gamma_2$ in $V^{\mathbb{P}}$; and it is natural to suspect that if a unique isomorphism π exists, then π should be "canonical" and hence should already exist in V. Fortunately the following result shows that it is at least consistent that such graphs exist.

Theorem 8.1.3. *It is consistent that for every regular cardinal $\kappa \geq \omega$, there exists a set $\{\Gamma_\alpha \mid \alpha < \kappa^+\}$ of pairwise nonisomorphic connected rigid graphs with the following property. If E is any equivalence relation on κ^+, then there exists a notion of forcing $\mathbb{P}$ such that*

(a) *$\mathbb{P}$ preserves cofinalities and cardinals;*

(b) *$\mathbb{P}$ does not adjoin any new κ-sequences of ordinals;*

(c) *each graph Γ_α remains rigid in $V^{\mathbb{P}}$;*

(d) *$\Gamma_\alpha \simeq \Gamma_\beta$ in $V^{\mathbb{P}}$ iff $\alpha \mathrel{E} \beta$.*

In Section 8.2, we shall show how to derive Theorem 8.1.1 from Theorem 8.1.3. The argument, which is essentially algebraic in nature, follows the method sketched in this section; but unfortunately, because of the extra difficulties which arise at limit stages, the analysis of the relevant normaliser towers is substantially more complicated. Finally in Section 8.3, we shall prove Theorem 8.1.3.

8.2. More normaliser towers

In this section, we shall derive Theorem 8.1.1 from Theorem 8.1.3. So throughout this section, we will assume that the following hypothesis holds in the ground model V.

HYPOTHESIS 8.2.1. For every regular cardinal $\kappa \geq \omega$, there exists a set of pairwise nonisomorphic connected rigid graphs $\{\Gamma_\alpha \mid \alpha < \kappa^+\}$ with the following property. If E is any equivalence relation on κ^+, then there exists a notion of forcing $\mathbb{P}$ such that

(a) $\mathbb{P}$ preserves cofinalities and cardinals;
(b) $\mathbb{P}$ does not adjoin any new κ-sequences of ordinals;
(c) each graph Γ_α remains rigid in $V^{\mathbb{P}}$;
(d) $\Gamma_\alpha \simeq \Gamma_\beta$ in $V^{\mathbb{P}}$ iff $\alpha \mathrel{E} \beta$.

Most of our effort will go into proving the following result.

THEOREM 8.2.2. *Assume Hypothesis 8.2.1. Then for every infinite cardinal λ and every ordinal $\alpha < \lambda$, there exist a graph Γ and a subgroup $H \leqslant \mathrm{Aut}(\Gamma)$ with the following properties.*

(a) *The normaliser tower of H in $\mathrm{Aut}(\Gamma)$ terminates after exactly α steps.*
(b) *If β is any ordinal such that $1 \leq \beta < \lambda$, then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that the normaliser tower of H in $\mathrm{Aut}^{V^{\mathbb{P}}}(\Gamma)$ terminates after exactly β steps.*

COROLLARY 8.2.3. *Assume Hypothesis 8.2.1. Then for every infinite cardinal λ and every ordinal $\alpha < \lambda$, there exists a centreless group G with the following properties.*

(a) $\tau(G) = \alpha$.
(b) *If β is any ordinal such that $1 \leq \beta < \lambda$, then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that $\tau^{V^{\mathbb{P}}}(G) = \beta$.*

PROOF OF COROLLARY 8.2.3. Let $\Gamma = \langle X, E \rangle$ be the graph and $H \leqslant \mathrm{Aut}(\Gamma)$ be the subgroup given by Theorem 8.2.2. By Theorem 6.6.18, there exists a field

K_Γ of cardinality $\max\{|X|, \omega\}$ such that whenever $\mathbb{P}$ is a (possibly trivial) notion of forcing and $M = V^{\mathbb{P}}$, then the following conditions are satisfied.

(a) X is an $\mathrm{Aut}^M(K_\Gamma)$-invariant subset of K_Γ.

(b) The restriction mapping, $\pi \mapsto \pi \restriction X$, is an isomorphism from $\mathrm{Aut}^M(K_\Gamma)$ onto $\mathrm{Aut}^M(\Gamma)$.

Let $G = PGL(2, K_\Gamma) \rtimes H$. By Theorem 4.1.6, $\tau^M(G)$ is the height of the normaliser tower of H in $\mathrm{Aut}^M(\Gamma)$. Thus G satisfies our requirements. $\square$

In the last section, we saw that the normaliser tower of the permutation group

$$(H_3, \Delta_3) \simeq (P_0, \Sigma_0) \times (P_0, \Sigma_0) \times (P_1, \Sigma_1) \times (P_2, \Sigma_2)$$

in $(\mathrm{Sym}(\Delta_3), \Delta_3)$ terminates after exactly 3 steps. Now we shall define an analogous permutation group $(H_\alpha, \Delta_\alpha)$ for each ordinal α.

DEFINITION 8.2.4. For each ordinal α, let $(H_\alpha, \Delta_\alpha)$ and $(F_\alpha, \Delta_\alpha)$ be the permutation groups defined inductively as follows.

(a) $|\Delta_0| = 1$ and $H_0 = F_0 = \{id_{\Delta_0}\}$.

(b) If $\alpha > 0$, then we define

$$(H_\alpha, \Delta_\alpha) = (H_0, \Delta_0) \times \prod_{\beta < \alpha} (F_\beta, \Delta_\beta)$$

and we define F_α to be the terminal group of the normaliser tower of H_α in $\mathrm{Sym}(\Delta_\alpha)$.

In particular, it follows that for each successor ordinal $\alpha = \beta + 1$, we have that

$$(H_{\beta+1}, \Delta_{\beta+1}) = (H_\beta, \Delta_\beta) \times (F_\beta, \Delta_\beta)\,.$$

During the proof of Lemma 8.2.6, we shall need a more explicit notation for the set Δ_α. So for each successor ordinal $\alpha = \beta + 1$, let Δ^1_β be the set such that $\Delta_{\beta+1} = \Delta_\beta \sqcup \Delta^1_\beta$; i.e.

$$(H_{\beta+1}, \Delta_{\beta+1}) = \left(H_\beta \times F_\beta, \Delta_\beta \sqcup \Delta^1_\beta\right).$$

Note that if δ is a limit ordinal, then

$$\Delta_\delta = \bigcup_{\xi < \delta} \Delta_\xi.$$

Also if α is an arbitrary ordinal and $\gamma < \alpha$, then

$$\Delta_\alpha = \Delta_\gamma \sqcup \bigsqcup_{\gamma \leq \xi < \alpha} \Delta^1_\xi.$$

NOTATION 8.2.5. Throughout this section, $N_\beta(H_\alpha)$ always denotes the βth group in the normaliser tower of H_α in $\mathrm{Sym}(\Delta_\alpha)$.

Lemma 8.2.6 says that the normaliser tower of H_α in $\mathrm{Sym}(\Delta_\alpha)$ terminates after exactly α steps. As the proof of Lemma 8.2.6 is rather involved, we will first explain the basic idea for the case when α is a successor ordinal. Let $\alpha = \beta + 1$, so that

$$(H_{\beta+1}, \Delta_{\beta+1}) = (H_\beta, \Delta_\beta) \times \left(F_\beta, \Delta^1_\beta\right).$$

We shall show that for each $\gamma \leq \beta$,

$$(N_\gamma(H_{\beta+1}), \Delta_{\beta+1}) = (N_\gamma(H_\beta), \Delta_\beta) \times \left(F_\beta, \Delta^1_\beta\right).$$

In particular,

$$(N_\beta(H_{\beta+1}), \Delta_{\beta+1}) = (F_\beta, \Delta_\beta) \times \left(F_\beta, \Delta^1_\beta\right).$$

It is now easily seen

$$(N_{\beta+1}(H_{\beta+1}), \Delta_{\beta+1}) = \left([F_\beta \times F_\beta] \rtimes \langle \sigma \rangle, \Delta_\beta \sqcup \Delta^1_\beta\right),$$

where σ is an element of order 2 which interchanges the sets Δ_β and Δ^1_β; and it turns out that this group is self-normalising in $\mathrm{Sym}(\Delta_{\beta+1})$.

In the proof of Lemma 8.2.6, we shall need to study the blocks of imprimitivity of F_α in its action on Δ_α. Recall that if (G, Ω) is a transitive permutation group, then the nonempty subset Z of Ω is a *block of imprimitivity* if for each $g \in G$, either $g[Z] = Z$ or $g[Z] \cap Z = \emptyset$. In this case, the distinct elements of the set $\{g[Z] \mid g \in G\}$ form a partition $\bigsqcup_{j \in J} Z_j$ of Ω; and we obtain a corresponding G-invariant equivalence relation E on Ω by defining $x \mathrel{E} y$ iff there exists $j \in J$ such that x, $y \in Z_j$. Conversely, if E is a G-invariant equivalence relation on Ω, then each E-equivalence class is a block of imprimitivity.

LEMMA 8.2.6. *For each ordinal α, the normaliser tower of H_α in* $\mathrm{Sym}(\Delta_\alpha)$ *terminates after exactly α steps.*

Proof. We shall prove the following statements by a simultaneous induction on $\alpha \geq 0$.

(1_α) F_α acts transitively on Δ_α.

(2_α) Let $\Delta_0 = \{v_0\}$. Then $\{\Delta_\beta \mid \beta \leq \alpha\}$ is the set of blocks Z of imprimitivity of $(F_\alpha, \Delta_\alpha)$ such that $v_0 \in Z$.

(3_α) For each $\beta \leq \alpha$, let E^α_β be the F_α-invariant equivalence relation on Δ_α corresponding to the block of imprimitivity Δ_β. Then for each $\beta \leq \gamma < \alpha$, the set Δ^1_γ is a union of E^α_β-equivalence classes.

(4_α) If $\beta < \alpha$, then

$$(N_\beta(H_\alpha), \Delta_\alpha) = (F_\beta, \Delta_\beta) \times \prod_{\beta \leq \gamma < \alpha} \left(F_\gamma, \Delta^1_\gamma\right).$$

(5_α) Furthermore, if $\beta < \alpha$, then $N_\beta(H_\alpha)$ is the subgroup of F_α consisting of those elements which fix setwise each of the sets in the partition

$$\{\Delta_\beta\} \cup \{\Delta^1_\gamma \mid \beta \leq \gamma < \alpha\}$$

of Δ_α.

(6_α) $N_\alpha(H_\alpha)$ is self-normalising in $\mathrm{Sym}(\Delta_\alpha)$ and so $F_\alpha = N_\alpha(H_\alpha)$.

It is easily checked that the result holds for $\alpha = 0$. Next we shall deal with the successor step of the induction. So suppose that $\alpha \geq 0$ and that the result holds for all $\beta \leq \alpha$. The following claim will be used in the proof that condition ($4_{\alpha+1}$) holds.

Claim 8.2.7. *If $\beta < \gamma \leq \alpha$, then (F_β, Δ_β) and $(F_\gamma, \Delta_\gamma)$ are nonisomorphic permutation groups.*

Proof of Claim 8.2.7. Let $\beta \leq \alpha$ and let v be any point of Δ_β. Let $\mathcal{B}_\beta(v)$ be the set of blocks Z of imprimitivity of (F_β, Δ_β) such that $v \in Z$. Then conditions (1_β) and (2_β) imply that $(\mathcal{B}_\beta(v), \subset)$ is a well-ordering of order-type $\beta + 1$. Now suppose that $\beta < \gamma \leq \alpha$ and that (f, φ) is a permutation group isomorphism from (F_β, Δ_β) onto $(F_\gamma, \Delta_\gamma)$. Thus

(i) $f : F_\beta \to F_\gamma$ is a group isomorphism;

(ii) $\varphi : \Delta_\beta \to \Delta_\gamma$ is a bijection; and

(iii) for all $g \in F_\beta$ and $v \in \Delta_\beta$, $f(g)(\varphi(v)) = \varphi(g(v))$.

It is easily checked that if $v \in \Delta_\beta$ and $X \subseteq \Delta_\beta$, then

$$X \in \mathcal{B}_\beta(v) \text{ iff } \varphi[X] \in \mathcal{B}_\gamma(\varphi(v)).$$

But this means that φ induces an order-isomorphism between the well-orderings $(\mathcal{B}_\beta(v), \subset)$ and $(\mathcal{B}_\gamma(\varphi(v)), \subset)$, which is a contradiction. □

By Lemma 1.3.15, if (G, Ω) is a permutation group and $\pi \in \mathrm{Sym}(\Omega)$ normalises G, then π permutes the orbits of G. Furthermore, if X and Y are G-orbits and $\pi[X] = Y$, then G must induce isomorphic permutation groups via its actions on X and Y. Using these observations, it is now easy to see that condition $(4_{\alpha+1})$ holds. First note that

$$(H_{\alpha+1}, \Delta_{\alpha+1}) = (H_\alpha, \Delta_\alpha) \times \left(F_\alpha, \Delta^1_\alpha\right).$$

Then, since the permutation group $(F_\alpha, \Delta_\alpha)$ is not isomorphic to $(F_\gamma, \Delta_\gamma)$ for any $\gamma < \alpha$, we see inductively that

$$\begin{aligned}(N_\beta(H_{\alpha+1}), \Delta_{\alpha+1}) &= (N_\beta(H_\alpha), \Delta_\alpha) \times \left(F_\alpha, \Delta^1_\alpha\right) \\ &= (F_\beta, \Delta_\beta) \times \prod_{\beta \leq \gamma < \alpha+1} \left(F_\gamma, \Delta^1_\gamma\right)\end{aligned}$$

for all $\beta \leq \alpha$. Next we shall check that condition $(1_{\alpha+1})$ holds. To see this, first note that

$$(N_\alpha(H_{\alpha+1}), \Delta_{\alpha+1}) = (F_\alpha, \Delta_\alpha) \times \left(F_\alpha, \Delta^1_\alpha\right).$$

Clearly $N_{\alpha+1}(H_{\alpha+1})$ contains an involution σ which interchanges Δ_α and Δ^1_α. Thus

$$[F_\alpha \times F_\alpha] \rtimes \langle \sigma \rangle = F_\alpha \text{ Wr } \mathrm{Sym}(2) \leqslant N_{\alpha+1}(H_{\alpha+1}) \leqslant F_{\alpha+1}$$

and so $F_{\alpha+1}$ acts transitively on $\Delta_{\alpha+1}$. For later use, we shall next check that

$$(N_{\alpha+1}(H_{\alpha+1}), \Delta_{\alpha+1}) = \left(F_\alpha \text{ Wr } \mathrm{Sym}(2), \Delta_\alpha \sqcup \Delta^1_\alpha\right).$$

To see this, suppose that $g \in N_{\alpha+1}(H_{\alpha+1})$ is arbitrary. Then g must permute the two orbits Δ_α and Δ^1_α of $N_\alpha(H_{\alpha+1})$. After replacing g by σg if necessary, we can suppose that $g[\Delta_\alpha] = \Delta_\alpha$ and $g[\Delta^1_\alpha] = \Delta^1_\alpha$; i.e. $g \in \mathrm{Sym}(\Delta_\alpha) \times \mathrm{Sym}(\Delta^1_\alpha)$. By condition (6_α), F_α is self-normalising in $\mathrm{Sym}(\Delta_\alpha)$. It follows that $g \in F_\alpha \times F_\alpha$.

Next we shall check that $\{\Delta_\beta \mid \beta \leq \alpha+1\}$ is the set of blocks Z of imprimitivity of

$$(N_{\alpha+1}(H_{\alpha+1}), \Delta_{\alpha+1}) = \left(F_\alpha \text{ Wr } \mathrm{Sym}(2), \Delta_\alpha \sqcup \Delta^1_\alpha\right)$$

such that $v_0 \in Z$. It is easily checked that Δ_β is a block of imprimitivity of $N_{\alpha+1}(H_{\alpha+1})$ for each $\beta \leq \alpha+1$. Now let $Z \subseteq \Delta_{\alpha+1}$ be any block of imprimitivity such that $v_0 \in Z$. First consider the case when $Z \cap \Delta^1_\alpha \neq \emptyset$. Since the subgroup $1 \times F_\alpha$ fixes v_0 and acts transitively on Δ^1_α, it follows that $\Delta^1_\alpha \subseteq Z$. Similarly $\Delta_\alpha \subseteq Z$ and so $Z = \Delta_{\alpha+1}$. So we can suppose that $Z \subseteq \Delta_\alpha$. In this case, Z must also be a block of imprimitivity of $(F_\alpha, \Delta_\alpha)$ and so $Z = \Delta_\beta$ for some $\beta \leq \alpha$.

Now suppose that $g \in \mathrm{Sym}(\Delta_{\alpha+1})$ normalises $N_{\alpha+1}(H_{\alpha+1}) = F_\alpha \text{ Wr } \mathrm{Sym}(2)$. Then g must permute the set $\{E^{\alpha+1}_\beta \mid \beta \leq \alpha+1\}$ of $N_{\alpha+1}(H_{\alpha+1})$-invariant equivalence relations on $\Delta_{\alpha+1}$. Since the set $\{E^{\alpha+1}_\beta \mid \beta \leq \alpha+1\}$ is well-ordered under inclusion, it follows that $E^{\alpha+1}_\beta$ is also g-invariant for each $\beta \leq \alpha+1$. In particular, $E^{\alpha+1}_\alpha$ is g-invariant and so g must induce a permutation of the set $\{\Delta_\alpha, \Delta^1_\alpha\}$ of $E^{\alpha+1}_\alpha$-equivalence classes. This implies that g must normalise the setwise stabiliser of Δ_α in $N_{\alpha+1}(H_{\alpha+1})$. In other words, g normalises $F_\alpha \times F_\alpha = N_\alpha(H_{\alpha+1})$ and so $g \in N_{\alpha+1}(H_{\alpha+1})$. So we have now established that conditions $(6_{\alpha+1})$ and $(2_{\alpha+1})$ also hold.

Next we shall check that condition $(5_{\alpha+1})$ holds. Let $\beta \leq \alpha$. Since

$$(N_\beta(H_{\alpha+1}), \Delta_{\alpha+1}) = (F_\beta, \Delta_\beta) \times \prod_{\beta \leq \gamma < \alpha+1} \left(F_\gamma, \Delta^1_\gamma\right),$$

it follows that $N_\beta(H_{\alpha+1})$ fixes setwise each of the sets in the partition

$$\{\Delta_\beta\} \cup \{\Delta^1_\gamma \mid \beta \leq \gamma < \alpha+1\}$$

of $\Delta_{\alpha+1}$. Conversely suppose $g \in F_{\alpha+1} = F_\alpha \text{ Wr } \mathrm{Sym}(2)$ fixes setwise each of the sets in the partition $\{\Delta_\beta\} \cup \{\Delta^1_\gamma \mid \beta \leq \gamma < \alpha+1\}$ of $\Delta_{\alpha+1}$. Then $g[\Delta_\alpha] \cap \Delta_\alpha \neq \emptyset$ and so $g \in F_\alpha \times F_\alpha$. Furthermore, $g \restriction \Delta_\alpha$ fixes setwise each of the sets in the partition $\{\Delta_\beta\} \cup \{\Delta^1_\gamma \mid \beta \leq \gamma < \alpha\}$ of Δ_α. So condition (5_α) implies that $g \restriction \Delta_\alpha \in N_\beta(H_\alpha)$. It follows that $g \in N_\beta(H_{\alpha+1})$.

Finally we shall check that condition $(3_{\alpha+1})$ holds and thus complete the proof that the inductive hypotheses hold for $\alpha+1$. Let $\beta \leq \gamma < \alpha+1$. If $\gamma < \alpha$, then condition (3_α) says that Δ^1_γ is a union of sets of the form $g[\Delta_\beta]$ for various $g \in F_\alpha$; and this implies that Δ^1_γ is a union of $E^{\alpha+1}_\beta$-equivalence classes. So we can suppose that $\gamma = \alpha$. Since Δ_α is a union of $E^{\alpha+1}_\beta$-equivalence classes and there exists an element $g \in F_{\alpha+1}$ such that $g[\Delta_\alpha] = \Delta^1_\alpha$, it follows that Δ^1_α is also a union of $E^{\alpha+1}_\beta$-equivalence classes.

Now suppose that λ is a limit ordinal and that the result holds for all $\alpha < \lambda$. Arguing as above, it is easy to see that conditions (4_λ) and (1_λ) hold; and it is also easy to check that the following statements hold.

$(2_\lambda)'$ For each $\beta \leq \lambda$, the set Δ_β is a block of imprimitivity of $(N_\lambda(H_\lambda), \Delta_\lambda)$.

$(3_\lambda)'$ For each $\beta \leq \lambda$, let E^λ_β be the $N_\lambda(H_\lambda)$-invariant equivalence relation on Δ_λ corresponding to the block of imprimitivity Δ_β. Then for each $\beta \leq \gamma < \lambda$, the set Δ^1_γ is a union of E^λ_β-equivalence classes.

$(5_\lambda)'$ If $\beta < \lambda$, then $N_\beta(H_\lambda)$ is the subgroup of $N_\lambda(H_\lambda)$ consisting of those elements which fix setwise each of the sets in the partition

$$\{\Delta_\beta\} \cup \{\Delta^1_\gamma \mid \beta \leq \gamma < \lambda\}$$

of Δ_λ.

Thus it is enough to prove the following two claims.

CLAIM 8.2.8. *If Z is a block of imprimitivity of $(N_\lambda(H_\lambda), \Delta_\lambda)$ with $v_0 \in Z$, then $Z = \Delta_\beta$ for some $\beta \leq \lambda$.*

CLAIM 8.2.9. *$N_\lambda(H_\lambda)$ is self-normalising in* $\mathrm{Sym}(\Delta_\lambda)$ *and so $F_\lambda = N_\lambda(H_\lambda)$.*

PROOF OF CLAIM 8.2.8. If there exists $\gamma < \lambda$ such that $Z \subseteq \Delta_\gamma$, then Z must also be a block of imprimitivity of $(F_\gamma, \Delta_\gamma)$ and so $Z = \Delta_\beta$ for some $\beta \leq \gamma$. Hence we can suppose that the set $I = \{\gamma < \lambda \mid Z \cap \Delta^1_\gamma \neq \emptyset\}$ is cofinal in λ. Fix some $\gamma \in I$. By condition (4_λ),

$$(N_\gamma(H_\lambda), \Delta_\lambda) = (F_\gamma, \Delta_\gamma) \times \prod_{\gamma \leq \xi < \lambda} \left(F_\xi, \Delta^1_\xi\right).$$

Hence for each $x \in \Delta_\gamma$, there exists an element $g \in N_\gamma(H_\lambda) \leqslant N_\lambda(H_\lambda)$ such that

(i) $g(v_0) = x$; and

(ii) $g(y) = y$ for all $y \in \Delta^1_\gamma$.

By condition (ii), $g[Z] \cap Z \neq \emptyset$ and hence $x \in g[Z] = Z$. Thus $\Delta_\gamma \subseteq Z$ for each $\gamma \in I$ and so $Z = \Delta_\lambda$. $\square$

PROOF OF CLAIM 8.2.9. Note that for each $\pi \in N_\lambda(H_\lambda)$, there exists $\beta < \lambda$ such that $\pi \in N_\beta(H_\lambda)$ and so $\pi\left[\Delta^1_\gamma\right] = \Delta^1_\gamma$ for all $\beta \leq \gamma < \lambda$. Suppose that $g \in \mathrm{Sym}(\Delta_\lambda)$ normalises $N_\lambda(H_\lambda)$. Then arguing as in the successor stage of the argument, it follows that E^λ_β is also g-invariant for each $\beta \leq \lambda$.

First we shall show that there exists an ordinal $\beta < \lambda$ such that $g\left[\Delta^1_\gamma\right] = \Delta^1_\gamma$ for all $\beta \leq \gamma < \lambda$. To see this, for each $\gamma < \lambda$, let $C^1_\gamma = g\left[\Delta^1_\gamma\right]$ and suppose that there exists a cofinal subset $I \subseteq \lambda$ such that $C^1_\gamma \neq \Delta^1_\gamma$ for all $\gamma \in I$. Fix some $\gamma \in I$. Since Δ^1_γ is an E^λ_γ-equivalence class, it follows that $C^1_\gamma = g\left[\Delta^1_\gamma\right]$ is also an E^λ_γ-equivalence class. Using condition $(3_\lambda)'$, it follows that either

(i) $C^1_\gamma = \Delta_\gamma$, or

(ii) there exists $f(\gamma) > \gamma$ such that $C^1_\gamma \subsetneq \Delta^1_{f(\gamma)}$.

If $\gamma \neq \gamma' \in I$, then $\Delta^1_\gamma \cap \Delta^1_{\gamma'} = \emptyset$ and so $C^1_\gamma \cap C^1_{\gamma'} = \emptyset$. Hence condition (i) can hold for at most one element $\gamma \in I$. Without loss of generality, we can assume that condition (ii) holds for all $\gamma \in I$. Furthermore, by passing to a suitable subset of I if necessary, we can assume that the resulting function $f : I \to \lambda$ is injective. Remember that

$$(H_\lambda, \Delta_\lambda) = (F_0, \Delta_0) \times \prod_{0 \leq \xi < \lambda} \left(F_\xi, \Delta^1_\xi\right).$$

Since each F_ξ acts transitively on Δ^1_ξ, there exists an element $\psi \in H_\lambda \leqslant N_\lambda(H_\lambda)$ such that $\psi\left[C^1_\gamma\right] \neq C^1_\gamma$ for all $\gamma \in I$. Let $\pi = g^{-1}\psi g \in N_\lambda(H_\lambda)$. Then $\pi\left[\Delta^1_\gamma\right] \neq \Delta^1_\gamma$ for all $\gamma \in I$, which is a contradiction.

Thus there exists $\beta < \lambda$ such that $g \in \mathrm{Sym}(\Delta_\beta) \times \prod_{\beta \leq \gamma < \lambda} \mathrm{Sym}(\Delta^1_\gamma)$. Since $N_\beta(H_\lambda)$ is the subgroup of $N_\lambda(H_\lambda)$ consisting of those elements which fix setwise each of the sets in the partition $\{\Delta_\beta\} \cup \{\Delta^1_\gamma \mid \beta \leq \gamma < \lambda\}$ of Δ_λ, it follows that g normalises $N_\beta(H_\lambda)$. Thus $g \in N_{\beta+1}(H_\lambda) \leqslant N_\lambda(H_\lambda)$. □

This completes the proof of Lemma 8.2.6. □

Definition 8.2.10. If Γ is a connected rigid graph and α is an ordinal, then we define $\mathcal{G}_\alpha(\Gamma)$ to be the graph obtained by replacing each element of Δ_α by a copy of Γ.

Thus Δ_α is essentially the set of connected components of $\mathcal{G}_\alpha(\Gamma)$ and $\mathrm{Aut}(\mathcal{G}_\alpha(\Gamma))$ can be naturally identified with $\mathrm{Sym}(\Delta_\alpha)$. Let $H_\alpha(\Gamma)$, $F_\alpha(\Gamma)$ be the subgroups of $\mathrm{Aut}(\mathcal{G}_\alpha(\Gamma))$ which correspond to the subgroups H_α, F_α of $\mathrm{Sym}(\Delta_\alpha)$. Then the following result is an immediate consequence of Lemma 8.2.6.

Lemma 8.2.11. *If Γ is a rigid connected graph, then normaliser tower of $H_\alpha(\Gamma)$ in $\mathrm{Aut}(\mathcal{G}_\alpha(\Gamma))$ terminates after exactly α steps.*

□

Now let β be any ordinal such that $1 \leq \beta < \alpha$. We shall next introduce a "device" which essentially has the effect of halting the normaliser tower of $H_\alpha(\Gamma)$ in $\operatorname{Aut}(\mathcal{G}_\alpha(\Gamma))$ after only β steps.

Definition 8.2.12. If Γ is a connected rigid graph and $1 \leq \beta < \alpha$, then we define

$$\left(D^\alpha_\beta(\Gamma), \mathcal{G}^\alpha_\beta(\Gamma)\right) = (H_\alpha(\Gamma), \mathcal{G}_\alpha(\Gamma)) \times (F_\beta(\Gamma), \mathcal{G}_\beta(\Gamma)) \times (F_\beta(\Gamma), \mathcal{G}_\beta(\Gamma)).$$

Lemma 8.2.13. *If Γ is a connected rigid graph, then the normaliser tower of $D^\alpha_\beta(\Gamma)$ in $\operatorname{Aut}(\mathcal{G}^\alpha_\beta(\Gamma))$ terminates after exactly β steps.*

Proof. Let $(B, \Gamma') = (F_\beta(\Gamma), \mathcal{G}_\beta(\Gamma)) \times (F_\beta(\Gamma), \mathcal{G}_\beta(\Gamma)) \times (F_\beta(\Gamma), \mathcal{G}_\beta(\Gamma))$ and let $F_\beta(\Gamma) \operatorname{Wr} \operatorname{Sym}(3) = B \rtimes \operatorname{Sym}(3)$ be the associated wreath product. By rearranging the order of its factors, we can identify $\left(D^\alpha_\beta(\Gamma), \mathcal{G}^\alpha_\beta(\Gamma)\right)$ with

$$(H_\beta(\Gamma), \mathcal{G}_\beta(\Gamma)) \times (B, \Gamma') \times \prod_{\beta<\gamma<\alpha} (F_\gamma(\Gamma), \mathcal{G}_\gamma(\Gamma)).$$

Arguing as in the proof of Lemma 8.2.6, we find that the βth element of the normaliser tower of $D^\alpha_\beta(\Gamma)$ in $\operatorname{Aut}(\mathcal{G}^\alpha_\beta(\Gamma))$ is

$$(F_\beta(\Gamma), \mathcal{G}_\beta(\Gamma)) \times (F_\beta(\Gamma) \operatorname{Wr} \operatorname{Sym}(3), \Gamma') \times \prod_{\beta<\gamma<\alpha} (F_\gamma(\Gamma), \mathcal{G}_\gamma(\Gamma))$$

and that this group is self-normalising in $\operatorname{Aut}(\mathcal{G}^\alpha_\beta(\Gamma))$. □

We are now ready to complete the proof of Theorem 8.2.2. So let λ be any infinite cardinal and let α be any ordinal such that $\alpha < \lambda$. Choose a regular cardinal κ such that $\kappa \geq \lambda$. Let $\{\Gamma_\gamma \mid \gamma < \kappa^+\}$ be the set of pairwise nonisomorphic connected rigid graphs given by Hypothesis 8.2.1. If $\alpha \geq 1$, then we define

$$(B_\alpha, \Gamma^\alpha) = \prod_{1\leq\beta<\alpha} \Big((F_\beta(\Gamma_\beta), \mathcal{G}_\beta(\Gamma_\beta)) \times (F_\beta(\Gamma_\beta), \mathcal{G}_\beta(\Gamma_\beta)) \Big)$$

and

$$(H, \Gamma) = (B_\alpha, \Gamma^\alpha) \times (H_\alpha(\Gamma_\alpha), \mathcal{G}_\alpha(\Gamma_\alpha)) \times \prod_{\alpha\leq\gamma<\lambda} (F_\gamma(\Gamma_{\gamma+1}), \mathcal{G}_\gamma(\Gamma_{\gamma+1})).$$

If $\alpha = 0$, then we define

$$(H, \Gamma) = (F_0(\Gamma_0), \mathcal{G}_0(\Gamma_0)) \times \prod_{0\leq\gamma<\lambda} (F_\gamma(\Gamma_{\gamma+1}), \mathcal{G}_\gamma(\Gamma_{\gamma+1})).$$

We shall show that the graph Γ and the subgroup $H \leqslant \operatorname{Aut}(\Gamma)$ satisfy the requirements of Theorem 8.2.2.

LEMMA 8.2.14. *The normaliser tower of H in $\operatorname{Aut}(\Gamma)$ terminates after exactly α steps.*

PROOF. For example, suppose that $\alpha \geq 1$. Then

$$\operatorname{Aut}(\Gamma) = \operatorname{Aut}(\Gamma^{\alpha}) \times \operatorname{Aut}(\mathcal{G}_{\alpha}(\Gamma_{\alpha})) \times \prod_{\alpha \leq \gamma < \lambda} \operatorname{Aut}(\mathcal{G}_{\gamma}(\Gamma_{\gamma+1})).$$

First by Lemma 8.2.11, the normaliser tower of $H_{\alpha}(\Gamma_{\alpha})$ in $\operatorname{Aut}(\mathcal{G}_{\alpha}(\Gamma_{\alpha}))$ terminates after exactly α steps. Next since $F_{\gamma}(\Gamma_{\gamma+1})$ is self-normalising in $\operatorname{Aut}(\mathcal{G}_{\gamma}(\Gamma_{\gamma+1}))$, it follows that $\prod_{\alpha \leq \gamma < \lambda} F_{\gamma}(\Gamma_{\gamma+1})$ is self-normalising in $\prod_{\alpha \leq \gamma < \lambda} \operatorname{Aut}(\mathcal{G}_{\gamma}(\Gamma_{\gamma+1}))$. Finally note that

$$(F_{\beta}(\Gamma_{\beta}), \mathcal{G}_{\beta}(\Gamma_{\beta})) \times (F_{\beta}(\Gamma_{\beta}), \mathcal{G}_{\beta}(\Gamma_{\beta})) \simeq (N_{\beta}(H_{\beta+1}(\Gamma_{\beta})), \mathcal{G}_{\beta+1}(\Gamma_{\beta}))$$

and so the normaliser tower of B_{α} in $\operatorname{Aut}(\Gamma^{\alpha})$ terminates after exactly 1 step. It follows that the normaliser tower of H in $\operatorname{Aut}(\Gamma)$ terminates after exactly α steps. $\square$

The following lemma completes the proof of Theorem 8.2.2.

LEMMA 8.2.15. *If β is any ordinal such that $1 \leq \beta < \lambda$, then there exists a notion of forcing $\mathbb{P}$, which preserves cofinalities and cardinals, such that the normaliser tower of H in $\operatorname{Aut}^{V^{\mathbb{P}}}(\Gamma)$ terminates after exactly β steps.*

PROOF. Clearly we can suppose that $\beta \neq \alpha$. There are two cases to consider. First suppose that $1 \leq \beta < \alpha$. Let E be the equivalence relation on κ^{+} such that

$$\gamma \; E \; \delta \quad \text{iff} \quad \{\gamma, \delta\} = \{\alpha, \beta\} \text{ or } \gamma = \delta$$

and let $\mathbb{P}$ be the corresponding notion of forcing, given by Hypothesis 8.2.1. Using the facts that

(a) each graph Γ_{δ} remains rigid in $V^{\mathbb{P}}$, and
(b) $\mathbb{P}$ does not adjoin any new κ-sequences of ordinals,

we see that

$$(H_{\gamma}(\Gamma_{\delta}), \mathcal{G}_{\gamma}(\Gamma_{\delta}))^{V^{\mathbb{P}}} = (H_{\gamma}(\Gamma_{\delta}), \mathcal{G}_{\gamma}(\Gamma_{\delta}))$$

and

$$(F_\gamma(\Gamma_\delta), \mathcal{G}_\gamma(\Gamma_\delta))^{V^{\mathbb{P}}} = (F_\gamma(\Gamma_\delta), \mathcal{G}_\gamma(\Gamma_\delta))$$

for all γ, $\delta < \kappa^+$. Let

$$(B', \Gamma') = \prod_{\substack{1 \leq \gamma < \alpha \\ \gamma \neq \beta}} ((F_\gamma(\Gamma_\gamma), \mathcal{G}_\gamma(\Gamma_\gamma)) \times (F_\gamma(\Gamma_\gamma), \mathcal{G}_\gamma(\Gamma_\gamma)))\,.$$

Then in $V^{\mathbb{P}}$, the permutation group (H, Γ) is isomorphic to

$$(B', \Gamma') \times \left(D^\alpha_\beta(\Gamma_\alpha), \mathcal{G}^\alpha_\beta(\Gamma_\alpha)\right) \times \prod_{\alpha \leq \gamma < \lambda} (F_\gamma(\Gamma_{\gamma+1}), \mathcal{G}_\gamma(\Gamma_{\gamma+1}))\,.$$

Hence the normaliser tower of H in $\mathrm{Aut}^{V^{\mathbb{P}}}(\Gamma)$ terminates after exactly β steps.

Now suppose that $\alpha < \beta < \lambda$. We will only deal with the case when $\alpha \geq 1$. (The case when $\alpha = 0$ is almost identical.) Let E be the equivalence relation on κ^+ such that $\gamma \; E \; \delta$ iff either

(i) $\alpha \leq \gamma, \delta < \beta + 1$; or
(ii) $\gamma = \delta$;

and let $\mathbb{P}$ be the corresponding notion of forcing, given by Hypothesis 8.2.1. Then in $V^{\mathbb{P}}$, the permutation group (H, Γ) is isomorphic to

$$(B_\alpha, \Gamma^\alpha) \times (H_\beta(\Gamma_\alpha), \mathcal{G}_\beta(\Gamma_\alpha)) \times \prod_{\beta \leq \gamma < \lambda} (F_\gamma(\Gamma_{\gamma+1}), \mathcal{G}_\gamma(\Gamma_{\gamma+1}))\,.$$

Hence the normaliser tower of H in $\mathrm{Aut}^{V^{\mathbb{P}}}(\Gamma)$ terminates after exactly β steps. $\square$

8.3. Rigid trees

In this section, we shall prove Theorem 8.1.3. Instead of working directly with graphs, we will find it more convenient to prove the following analogous theorem for trees.

THEOREM 8.3.1. *It is consistent that for every regular cardinal $\kappa \geq \omega$, there exists a set $\{T_\alpha \mid \alpha < \kappa^+\}$ of pairwise nonisomorphic rigid trees of height κ^+ with the following property. If E is any equivalence relation on κ^+, then there exists a notion of forcing $\mathbb{P}$ such that*

(a) *$\mathbb{P}$ preserves cofinalities and cardinals;*
(b) *$\mathbb{P}$ does not adjoin any new κ-sequences of ordinals;*
(c) *each tree T_α remains rigid in $V^{\mathbb{P}}$;*

(d) $T_\alpha \simeq T_\beta$ *in* $V^{\mathbb{P}}$ *iff* $\alpha \mathrel{E} \beta$.

First we shall show how to derive Theorem 8.1.3 from Theorem 8.3.1.

Proof of Theorem 8.1.3. Let $\mathcal{T}$ be the category of trees and let $\mathcal{G}$ be the category of graphs. For each tree $T \in \mathcal{T}$, let Γ_T be the corresponding connected graph which encodes T, as given in the proof of Lemma 4.2.2. Thus if T, $T' \in \mathcal{T}$, then

(a) $T \simeq T'$ iff $\Gamma_T \simeq \Gamma_{T'}$; and
(b) $\mathrm{Aut}(T) \simeq \mathrm{Aut}(\Gamma_T)$.

Furthermore, an examination of the proof of Lemma 4.2.2 shows that the construction of Γ_T from T is upwards absolute. Hence properties (a) and (b) will continue to hold in any generic extension of the ground model.

Now suppose that for every regular cardinal $\kappa \geq \omega$, the ground model contains a set $\{T_\alpha \mid \alpha < \kappa^+\}$ of trees satisfying the conditions of Theorem 8.3.1. By the above remarks, it follows that the corresponding set $\{\Gamma_{T_\alpha} \mid \alpha < \kappa^+\}$ of graphs satisfies the conditions of Theorem 8.1.3. $\square$

Before we begin the proof of Theorem 8.3.1, we need to introduce some notions from the theory of trees.

Definition 8.3.2. (a) A *tree* is a partially ordered set $\langle T, < \rangle$ such that for every $x \in T$, the set $\mathrm{pred}_T(x) = \{y \in T \mid y < x\}$ is well-ordered by $<$.
(b) If $x \in T$, then the *height* of x in T, denoted $\mathrm{ht}_T(x)$, is the order-type of $\mathrm{pred}_T(x)$ under $<$.
(c) If α is an ordinal, then the αth *level* of T is the set

$$\mathrm{Lev}_\alpha(T) = \{x \in T \mid \mathrm{ht}_T(x) = \alpha\}$$

and $T \restriction \alpha = \bigcup_{\beta < \alpha} \mathrm{Lev}_\beta(T)$. The *height* of the tree T is the least ordinal α such that $\mathrm{Lev}_\alpha(T) = \emptyset$.
(d) A *branch* of T is a maximal linearly ordered subset of T. The *length* of a branch B is the order-type of B. An α-*branch* is a branch of length α.

Definition 8.3.3. Let δ be an ordinal and let λ be a cardinal. A tree T is said to be a (δ, λ)-*tree* iff

(i) for all $\alpha < \delta$, $0 < |\mathrm{Lev}_\alpha(T)| < \lambda$; and

(ii) $\mathrm{Lev}_\delta(T) = \emptyset$.

A (δ, λ)-tree T is *normal* if each of the following conditions is satisfied.

(a) If $\delta > 0$, then $|\mathrm{Lev}_0(T)| = 1$.

(b) If $\alpha + 1 < \delta$ and $x \in \mathrm{Lev}_\alpha(T)$, then there exist exactly two elements y_1, $y_2 \in \mathrm{Lev}_{\alpha+1}(T)$ such that $x < y_1$ and $x < y_2$.

(c) If $\alpha < \beta < \delta$ and $x \in \mathrm{Lev}_\alpha(T)$, then there exists $y \in \mathrm{Lev}_\beta(T)$ such that $x < y$.

(d) If α is a limit ordinal and x, $y \in \mathrm{Lev}_\alpha(T)$ are distinct elements, then $\mathrm{pred}_T(x) \neq \mathrm{pred}_T(y)$.

Let T be a (δ, λ)-tree. Then the tree T^+ is an *end-extension* of T, written $T \lessdot T^+$, if $T^+ \upharpoonright \delta = T$. The tree T^+ is a *proper* end-extension if $T \lessdot T^+$ and $T \neq T^+$.

Definition 8.3.4. Let $\kappa \geq \omega$ be a regular cardinal such that $\kappa^{<\kappa} = \kappa$ and let $\alpha < \kappa^+$. A normal (α, κ^+)-tree T is *κ-closed* if for each $\beta < \alpha$ such that $\mathrm{cf}(\beta) < \kappa$ and each increasing sequence of elements of T

$$x_0 < x_1 < \cdots < x_\xi < \cdots, \qquad \xi < \beta,$$

such that $x_\xi \in \mathrm{Lev}_\xi(T)$ for each $\xi < \beta$, there exists an element $y \in \mathrm{Lev}_\beta(T)$ such that $\mathrm{pred}_T(y) = \{x_\xi \mid \xi < \beta\}$.

In the proof of Theorem 8.3.1, we shall make repeated use of the following simple observation.

Lemma 8.3.5. *Let $\kappa \geq \omega$ be a regular cardinal such that $\kappa^{<\kappa} = \kappa$ and let $\alpha < \kappa^+$. If T is a κ-closed normal (α, κ^+)-tree and $x \in T$, then there exists an α-branch B of T such that $x \in B$.*

□

Now we are ready to discuss the notion of forcing $\mathbb{Q}_\kappa$ which will adjoin a set $\{T_\alpha \mid \alpha < \kappa^+\}$ of pairwise nonisomorphic rigid trees of height κ^+ satisfying the conditions of Theorem 8.3.1. Until further notice, M will denote the ground model and $\kappa \geq \omega$ will be a regular cardinal such that $\kappa^{<\kappa} = \kappa$ and $2^\kappa = \kappa^+$.

Definition 8.3.6. Let $\kappa \geq \omega$ be a regular cardinal such that $\kappa^{<\kappa} = \kappa$. Then $\mathbb{Q}_\kappa$ is the notion of forcing consisting of all conditions $p = \langle t^p_\alpha \mid \alpha < \kappa^+ \rangle$, where

(a) each t^p_α is a κ-closed normal (β_α, κ^+)-tree for some $\beta_\alpha < \kappa^+$;

(b) there exists an ordinal $\gamma < \kappa^+$ such that $t^p_\alpha = \emptyset$ for all $\gamma \le \alpha < \kappa^+$; and

(c) if $t^p_\alpha \neq \emptyset$, then the universe of t^p_α is an ordinal $\eta_\alpha < \kappa^+$.

We define $q \le p$ iff $t^p_\alpha \lessdot t^q_\alpha$ for all $\alpha < \kappa^+$.

It is easily checked that $\mathbb{Q}_\kappa$ is κ^+-closed and that $|\mathbb{Q}_\kappa| = \kappa^+$. (Clause (c) of Definition 8.3.6 was only included to ensure that $|\mathbb{Q}_\kappa| = \kappa^+$. The reader can safely ignore clause (c) from this point onwards.) By Theorem 6.3.21, $\mathbb{Q}_\kappa$ preserves cofinalities and cardinals. The meaning of a condition $p = \langle t^p_\alpha \mid \alpha < \kappa^+ \rangle \in \mathbb{Q}_\kappa$ should be clear: for each $\alpha < \kappa^+$, the tree t^p_α is intended to be an initial approximation to the tree $T_\alpha \in M^{\mathbb{Q}_\kappa}$. The following result implies that whenever $\beta_\alpha < \gamma < \kappa^+$, there exists a κ-closed normal (γ, κ^+)-tree t^+ such that $t^p_\alpha \lessdot t^+$. Consequently, the tree $T_\alpha \in M^{\mathbb{Q}_\kappa}$ will have height κ^+.

Lemma 8.3.7. *Let $\kappa \ge \omega$ be a regular cardinal such that $\kappa^{<\kappa} = \kappa$.*

(a) *For each $\beta < \kappa^+$, there exists a κ-closed normal (β, κ^+)-tree.*

(b) *If $\beta < \gamma < \kappa^+$ and S is a κ-closed normal (β, κ^+)-tree, then there exists a κ-closed normal (γ, κ^+)-tree S' such that $S \lessdot S'$.*

Proof. It is enough to prove (b). So suppose that S is a κ-closed normal (β, κ^+)-tree. We shall prove inductively that there exists a sequence $\langle S_\gamma \mid \beta \le \gamma < \kappa^+ \rangle$ satisfying:

(i) $S = S_\beta$;

(ii) if $\beta \le \xi \le \gamma$, then $S_\xi \lessdot S_\gamma$; and

(iii) S_γ is a κ-closed normal (γ, κ^+)-tree.

Clearly there is no difficulty when γ is a limit ordinal or when $\gamma = \xi + 1$ for some successor ordinal ξ. Hence we can assume that $\gamma = \delta + 1$, where δ is a limit ordinal.

First suppose that $\mathrm{cf}(\delta) < \kappa$ and let $\langle \xi_i \mid i < \mathrm{cf}(\delta) \rangle$ be a strictly increasing sequence of ordinals such that $\sup_{i<\mathrm{cf}(\delta)} \xi_i = \delta$. Let $\mathcal{B}$ be the set of δ-branches of S_δ. Notice that each branch $B \in \mathcal{B}$ is uniquely determined by its $\mathrm{cf}(\delta)$-subset

$$B \cap \bigcup_{i<\mathrm{cf}(\delta)} \mathrm{Lev}_{\xi_i}(S_\delta).$$

Since $\kappa^{<\kappa} = \kappa$, it follows that $|\mathcal{B}| \le \kappa$. Hence we can take

$$S_{\delta+1} = S_\delta \cup \{b_B \mid B \in \mathcal{B}\},$$

where $y < b_B$ iff $y \in B$.

Now suppose that $\mathrm{cf}(\delta) = \kappa$. For each $x \in S_\delta$, fix some δ-branch $B(x)$ of S_δ such that $x \in B(x)$. Then we can take

$$S_{\delta+1} = S_\delta \cup \{b_{B(x)} \mid x \in S_\delta\},$$

where $y < b_{B(x)}$ iff $y \in B(x)$. □

Notice that there was no real choice in the construction of S_γ, except in the case where $\gamma = \delta + 1$ for a limit ordinal δ of cofinality κ, when we get to select which branches of S_δ are the predecessors of elements of $S_{\delta+1}$. The following easy observation will play a crucial role in the proof that the trees $\{T_\alpha \mid \alpha < \kappa^+\}$ are rigid and pairwise nonisomorphic, since it will enable us to "kill off" potential isomorphisms.

LEMMA 8.3.8. *Suppose that $\delta < \kappa^+$ is a limit ordinal of cofinality κ and that S is a κ-closed normal (δ, κ^+)-tree. If $B \neq C$ are δ-branches of S, then there exists a κ-closed normal $(\delta + 1, \kappa^+)$-tree S' such that:*

(1) *$S \lessdot S'$;*
(2) *there exists $b \in S'$ such that* $\mathrm{pred}_{S'}(b) = B$*; and*
(3) *there does* not *exist $c \in S'$ such that* $\mathrm{pred}_{S'}(c) = C$.

PROOF. Clearly we can choose a collection $\mathcal{B} = \{B(x) \mid x \in S\}$ of δ-branches of S such that:

(i) $x \in B(x)$ for each $x \in S$;
(ii) $B \in \mathcal{B}$; and
(iii) $C \notin \mathcal{B}$.

Then we can take

$$S' = S \cup \{b_{B(x)} \mid x \in S\},$$

where $y < b_{B(x)}$ iff $y \in B(x)$. □

While the trees $\{T_\alpha \mid \alpha < \kappa^+\}$ adjoined by $\mathbb{Q}_\kappa$ are supposed to be pairwise nonisomorphic, they are also supposed to be "potentially isomorphic" by means of a further notion of forcing. In order for this to be possible, it is necessary that $T_\alpha \restriction \gamma \simeq T_\beta \restriction \gamma$ for all $\alpha < \beta < \kappa^+$ and $\gamma < \kappa^+$. This point is dealt with by the following lemma.

Lemma 8.3.9. *Let $\kappa \geq \omega$ be a regular cardinal such that $\kappa^{<\kappa} = \kappa$ and let $\alpha < \kappa^+$. If S and T are κ-closed normal (α, κ^+)-trees, then $S \simeq T$. Furthermore, if $\delta + 1 \leq \alpha$, then for each isomorphism $\varphi : S \restriction (\delta + 1) \to T \restriction (\delta + 1)$, there exists an isomorphism $\pi : S \to T$ such that $\varphi \subseteq \pi$.*

Proof. If $\alpha \leq \kappa$, then S and T are both complete binary trees of height α and so $S \simeq T$. Hence we can suppose that $\kappa < \alpha < \kappa^+$. Thus $|S| = |T| = \kappa$. We will define an isomorphism $\pi = \bigcup_{\xi<\kappa} \pi_\xi : S \to T$ via a back-and-forth argument.

Suppose that we have defined π_ξ for some $\xi < \kappa$. Assume inductively that there exists a set $\{B_i \mid i \in I\}$ of α-branches of S such that

(i) $|I| < \kappa$; and
(ii) $\operatorname{dom} \pi_\xi = \bigcup_{i \in I} B_i$.

Let s be any element of $S \smallsetminus \operatorname{dom} \pi_\xi$. Choose an α-branch B of S such that $s \in B$. Let $B = \{b_\tau \mid \tau < \alpha\}$, where $b_\tau \in B \cap \operatorname{Lev}_\tau(S)$. Let β be the least ordinal such that $b_\beta \notin \operatorname{dom} \pi_\xi$. First suppose that $\beta = \gamma + 1$ is a successor ordinal. Then there exists an $i \in I$ such that $b_\gamma \in B_i$. Let $C_i = \pi_\xi [B_i]$. Then there exists a unique element $c \in \operatorname{Lev}_\beta(T) \smallsetminus C_i$ such that $\pi_\xi(b_\gamma) < c$. Let C be an α-branch of T such that $c \in C$ and let $\psi : B \to C$ be the unique order-preserving bijection. Then $\pi_{\xi+1} = \pi_\xi \cup \psi$ is a partial isomorphism such that $s \in \operatorname{dom} \pi_{\xi+1}$. Now suppose that β is a limit ordinal. Since $\{b_\tau \mid \tau < \beta\}$ is covered by the set of branches $\{B_i \mid i \in I\}$, it follows that $\operatorname{cf}(\beta) < \kappa$. Hence there exists an element $c \in \operatorname{Lev}_\beta(T)$ such that $\operatorname{pred}_T(c) = \{\pi_\xi(b_\tau) \mid \tau < \beta\}$. Let C be an α-branch of T such that $c \in C$ and let $\psi : B \to C$ be the unique order-preserving bijection. Once again, $\pi_{\xi+1} = \pi_\xi \cup \psi$ is a partial isomorphism such that $s \in \operatorname{dom} \pi_{\xi+1}$. By a similar argument, if t is any element of $T \smallsetminus \operatorname{ran}(\pi_{\xi+1})$, then we can find a partial isomorphism $\pi_{\xi+2} \supset \pi_{\xi+1}$ such that $t \in \operatorname{ran} \pi_{\xi+2}$. Hence we can ensure that $\pi = \bigcup_{\xi<\kappa} \pi_\xi$ is an isomorphism from S onto T.

Finally suppose that $\delta + 1 \leq \alpha$ and that $\varphi : S \restriction (\delta + 1) \to T \restriction (\delta + 1)$ is an isomorphism. For each $s \in S$ and $t \in T$, define $S[s] = \{x \in S \mid s \leq x\}$ and $T[t] = \{y \in T \mid t \leq y\}$. Let γ be the ordinal such that $\alpha = \delta + \gamma$. Then for each $s \in \operatorname{Lev}_\delta(S)$, both $S[s]$ and $T[\varphi(s)]$ are κ-closed normal (γ, κ^+)-trees, and so $S[s] \simeq T[\varphi(s)]$. Hence φ can be extended to an isomorphism $\pi : S \to T$. $\square$

Now let G be a $\mathbb{Q}_\kappa$-generic filter over M and let $M[G]$ be the corresponding generic extension. Since $\mathbb{Q}_\kappa$ is κ^+-closed, we also have that $\kappa^{<\kappa} = \kappa$ and $2^\kappa = \kappa^+$ within $M[G]$. For each $\alpha < \kappa^+$, let $T_\alpha = \bigcup\{t^p_\alpha \mid p \in G\} \in M[G]$.

LEMMA 8.3.10. *In $M[G]$, $\{T_\alpha \mid \alpha < \kappa^+\}$ is a set of distinct pairwise nonisomorphic rigid trees of height κ^+.*

PROOF. For each $\alpha < \kappa^+$, let $\widetilde{T}_\alpha$ be the canonical $\mathbb{Q}_\kappa$-name for T_α. First suppose that for some $\alpha < \kappa^+$, there exists a nonidentity automorphism f of T_α in $M[G]$. Let $\widetilde{f}$ be a $\mathbb{Q}_\kappa$-name for f. Then there exists a condition $p \in \mathbb{Q}_\kappa$ and an element $a \in t^p_\alpha$ such that

$$p \Vdash \widetilde{f} : \widetilde{T}_\alpha \to \widetilde{T}_\alpha \text{ is an isomorphism such that } \widetilde{f}(a) \neq a.$$

Since $\mathbb{Q}_\kappa$ is κ^+-closed, we can inductively define a descending sequence of conditions $\langle p_\xi \mid \xi < \kappa\rangle$ such that

(1) $p_0 = p$;
(2) $t^{p_{\xi+1}}_\alpha$ is a proper end-extension of $t^{p_\xi}_\alpha$; and
(3) $p_{\xi+1}$ decides $\widetilde{f} \restriction t^{p_\xi}_\alpha$.

Let $q \in \mathbb{Q}_\kappa$ be the greatest lower bound of the sequence $\langle p_\xi \mid \xi < \kappa\rangle$. Then $t^q_\alpha = \bigcup_{\xi<\kappa} t^{p_\xi}_\alpha$ and so q decides $\widetilde{f} \restriction t^q_\alpha$. (This is another instance of the *bootstrap argument*.) Note that t^q_α is a κ-closed normal (γ, κ^+)-tree for some ordinal γ such that $\mathrm{cf}(\gamma) = \kappa$. Let B be a γ-branch of t^q_α such that $a \in B$ and let C be the γ-branch of t^q_α such that $q \Vdash \widetilde{f}[B] = C$. Then $B \neq C$. Since $\mathrm{cf}(\gamma) = \kappa$, there exists a κ-closed normal $(\gamma + 1, \kappa^+)$-tree t^+_α such that

(i) t^+_α is a proper end-extension of t^q_α;
(ii) there exists $x \in t^+_\alpha$ such that $\mathrm{pred}_{t^+_\alpha}(x) = B$; and
(iii) there does *not* exist $y \in t^+_\alpha$ such that $\mathrm{pred}_{t^+_\alpha}(y) = C$.

Let $r \le q$ be a condition such that $t^+_\alpha \lessdot t^r_\alpha$. Then

$$r \Vdash \widetilde{f} \restriction t^q_\alpha \text{ cannot be extended to an automorphism of } t^+_\alpha,$$

which is a contradiction.

Now suppose that for some ordinals $\alpha < \beta < \kappa^+$, there exists an isomorphism $g : T_\alpha \to T_\beta$ in $M[G]$. Let $\widetilde{g}$ be a $\mathbb{Q}_\kappa$-name for g. Then there exists a condition

$p \in \mathbb{Q}_\kappa$ such that

$$p \Vdash \widetilde{g} : \widetilde{T}_\alpha \to \widetilde{T}_\beta \text{ is an isomorphism.}$$

By the bootstrap argument, there exists a condition $q \leq p$ such that

(a) t^q_α and t^q_β are κ-closed normal (γ, κ^+)-trees for some γ such that $\mathrm{cf}(\gamma) = \kappa$; and

(b) q decides $\widetilde{g} \restriction t^q_\alpha$.

Let B be a γ-branch of t^q_α and let C be the γ-branch of t^q_β such that $q \Vdash \widetilde{g}[B] = C$. Since $\mathrm{cf}(\gamma) = \kappa$, there exist κ-closed normal $(\gamma + 1, \kappa^+)$-trees t^+_α and t^+_β such that

(1) t^+_α and t^+_β are proper end-extensions of t^q_α, t^q_β respectively;

(2) there exists $x \in t^+_\alpha$ such that $\mathrm{pred}_{t^+_\alpha}(x) = B$; and

(3) there does *not* exist $y \in t^+_\beta$ such that $\mathrm{pred}_{t^+_\beta}(y) = C$.

But this means that $\widetilde{g} \restriction t^q_\alpha$ cannot be extended to an isomorphism from t^+_α onto t^+_β. Once again, this yields a contradiction. $\square$

Next suppose that $E \in M[G]$ is any equivalence relation on κ^+. Let $A \subseteq \kappa^+$ be the set of E-equivalence class representatives obtained by selecting the least element of each class.

Definition 8.3.11. $\mathbb{P}_E$ is the notion of forcing in $M[G]$ consisting of all conditions $p = \langle f_{\alpha\beta} \mid \alpha < \beta < \kappa^+ \rangle$ such that for some $\gamma < \kappa^+$,

(a) if $\alpha \in A$, $\beta < \gamma$ and $\alpha \; E \; \beta$, then there exists $\delta < \kappa^+$ such that $f_{\alpha\beta}$ is an isomorphism from $T_\alpha \restriction (\delta + 1)$ onto $T_\beta \restriction (\delta + 1)$;

(b) otherwise, $f_{\alpha\beta} = \emptyset$.

If $p = \langle f_{\alpha\beta} \mid \alpha < \beta < \kappa^+ \rangle$, $q = \langle g_{\alpha\beta} \mid \alpha < \beta < \kappa^+ \rangle \in \mathbb{P}_E$, then we define $q \leq p$ iff $f_{\alpha\beta} \subseteq g_{\alpha\beta}$ for all $\alpha < \beta < \kappa^+$.

Remark 8.3.12. Some readers may be wondering why we have introduced the set A of E-equivalence class representatives. Consider the slightly simpler notion of forcing $\mathbb{P}'_E$ consisting of all conditions $p = \langle f_{\alpha\beta} \mid \alpha < \beta < \kappa^+ \rangle$ such that for some $\gamma < \kappa^+$,

(a) if $\alpha < \beta < \gamma$ and $\alpha \; E \; \beta$, then there exists $\delta < \kappa^+$ such that $f_{\alpha\beta}$ is an isomorphism from $T_\alpha \restriction (\delta + 1)$ onto $T_\beta \restriction (\delta + 1)$;

(b) otherwise, $f_{\alpha\beta} = \emptyset$.

Using Lemma 8.3.9, it is easily seen that the notion of forcing $\mathbb{P}'_E$ adjoins a generic isomorphism $g_{\alpha\beta} : T_\alpha \to T_\beta$ for each $\alpha < \beta < \kappa^+$ such that $\alpha \; E \; \beta$. Fix such a pair $\alpha < \beta$ and suppose that there exists an ordinal γ such that $\beta < \gamma < \kappa^+$ and $\beta \; E \; \gamma$. Then it is easily checked that $g_{\alpha\gamma}$ and $g_{\beta\gamma} \circ g_{\alpha\beta}$ will be *distinct* isomorphisms from T_α onto T_γ and so T_α will no longer be rigid. The set A was introduced to deal with precisely this problem. For example, suppose that $\alpha \in A$. Then the notion of forcing $\mathbb{P}_E$ will only directly adjoin isomorphisms $g_{\alpha\beta} : T_\alpha \to T_\beta$ and $g_{\alpha\gamma} : T_\alpha \to T_\gamma$. Of course, we can then obtain an isomorphism from T_β onto T_γ by forming the composition $g_{\alpha\gamma} \circ g_{\alpha\beta}^{-1}$.

Let H be a $\mathbb{P}_E$-generic filter over $M[G]$ and let $M[G][H]$ be the corresponding generic extension. The following result is an immediate consequence of the discussion in Remark 8.3.12.

LEMMA 8.3.13. *If $\alpha < \beta < \kappa^+$ and $\alpha \; E \; \beta$, then there exists an isomorphism $g_{\alpha\beta} : T_\alpha \to T_\beta$ in $M[G][H]$.*

□

LEMMA 8.3.14. *$\mathbb{P}_E$ preserves cofinalities and cardinals, and does not adjoin any new κ-sequences of ordinals. Furthermore, the following statements hold in $M[G][H]$.*

(a) *T_α is rigid for each $\alpha < \kappa^+$.*
(b) *If $\alpha < \beta < \kappa^+$, then $T_\alpha \simeq T_\beta$ iff $\alpha \; E \; \beta$.*

PROOF. Let $\widetilde{E}$, $\widetilde{A}$ and $\widetilde{\mathbb{P}}_E$ be $\mathbb{Q}_\kappa$-names for E, A and $\mathbb{P}_E$ respectively. Let $\mathbb{R}$ be the subset of $\mathbb{Q}_\kappa * \widetilde{\mathbb{P}}_E$ consisting of those conditions

$$\langle p, \widetilde{q} \rangle = \langle \langle t^p_\alpha \mid \alpha < \kappa^+ \rangle, \langle f_{\alpha\beta} \mid \alpha < \beta < \kappa^+ \rangle \rangle$$

such that for some γ, $\delta < \kappa^+$,

(1) p decides $\widetilde{E} \restriction (\gamma \times \gamma)$ and hence p also decides $\widetilde{A} \cap \gamma$;
(2) if $\alpha < \gamma$, then t^p_α is a κ-closed normal $(\delta + 1, \kappa^+)$-tree; and
(3) (i) if $\alpha < \beta < \gamma$ and $p \Vdash \alpha \in \widetilde{A}$ and $\alpha \; \widetilde{E} \; \beta$, then $f_{\alpha\beta}$ is an isomorphism from t^p_α onto t^p_β;
  (ii) otherwise, $f_{\alpha\beta} = \emptyset$.

(Here we are identifying each isomorphism $f_{\alpha\beta}$ with its canonical $\mathbb{Q}_\kappa$-name $\check{f}_{\alpha\beta}$.)

CLAIM 8.3.15. *$\mathbb{R}$ is a dense subset of $\mathbb{Q}_\kappa * \widetilde{\mathbb{P}}_E$.*

PROOF OF CLAIM 8.3.15. Let $\langle p, \widetilde{q}\rangle = \langle p, \langle \widetilde{f}_{\alpha\beta} \mid \alpha < \beta < \kappa^+\rangle\rangle$ be any element of $\mathbb{Q}_\kappa * \widetilde{\mathbb{P}}_E$. Then there exists $p' \leq p$ and $\gamma < \kappa^+$ such that p' forces

(a) $\widetilde{f}_{\alpha\beta} = \emptyset$ for all $\beta \geq \gamma$; and

(b) if $\alpha \in \widetilde{A}$, $\beta < \gamma$ and $\alpha \mathrel{\widetilde{E}} \beta$, then there exists $\tau < \gamma$ such that $\operatorname{dom} \widetilde{f}_{\alpha\beta}$ is a κ-closed normal $(\tau+1, \kappa^+)$-tree.

Since $\mathbb{Q}_\kappa$ is κ^+-closed, there exists $r \leq p'$ such that

(c) r decides $\widetilde{E} \restriction (\gamma \times \gamma)$ and hence r also decides $\widetilde{A} \cap \gamma$;

(d) there exists $\delta \geq \gamma$ such that t^r_α is a κ-closed normal $(\delta+1, \kappa^+)$-tree for each $\alpha < \gamma$; and

(e) if $\alpha < \beta < \gamma$ and $r \Vdash \alpha \in \widetilde{A}$ and $\alpha \mathrel{\widetilde{E}} \beta$, then there exists $\tau < \gamma$ and an isomorphism $f_{\alpha\beta} : t^r_\alpha \restriction (\tau+1) \to t^r_\beta \restriction (\tau+1)$ such that $r \Vdash \widetilde{f}_{\alpha\beta} = f_{\alpha\beta}$.

By Lemma 8.3.9, for each pair $\alpha < \beta < \gamma$ such that $r \Vdash \alpha \in \widetilde{A}$ and $\alpha \mathrel{\widetilde{E}} \beta$, there exists an isomorphism $g_{\alpha\beta} : t^r_\alpha \to t^r_\beta$ such that $f_{\alpha\beta} \subset g_{\alpha\beta}$. Let $g_{\alpha\beta} = \emptyset$ for all other pairs $\alpha < \beta < \kappa^+$. Then $\langle r, \langle g_{\alpha\beta} \mid \alpha < \beta < \kappa^+\rangle\rangle \in \mathbb{R}$ is a strengthening of $\langle p, \widetilde{q}\rangle$. $\square$

Thus the forcing notions $\mathbb{Q}_\kappa * \widetilde{\mathbb{P}}_E$ and $\mathbb{R}$ are equivalent. It is easily checked that $|\mathbb{R}| = \kappa^+$.

CLAIM 8.3.16. *$\mathbb{R}$ is κ^+-closed.*

PROOF OF CLAIM 8.3.16. Suppose that $\lambda < \kappa^+$ and that $\langle \langle p_\xi, \widetilde{q}_\xi\rangle \mid \xi < \lambda\rangle$ is a descending sequence of elements of $\mathbb{R}$. For each $\xi < \lambda$, let $p_\xi = \langle t^{p_\xi}_\alpha \mid \alpha < \kappa^+\rangle$ and $\widetilde{q}_\xi = \langle f^\xi_{\alpha\beta} \mid \alpha < \beta < \kappa^+\rangle$. For each $\alpha < \beta < \kappa^+$, let $t_\alpha = \bigcup_{\xi<\lambda} t^{p_\xi}_\alpha$ and $f_{\alpha\beta} = \bigcup_{\xi<\lambda} f^\xi_{\alpha\beta}$. Then $p = \langle t_\alpha \mid \alpha < \kappa^+\rangle \in \mathbb{Q}_\kappa$ and there exist ordinals γ, $\delta < \kappa^+$ such that

(1) p decides $\widetilde{E} \restriction (\gamma \times \gamma)$ and hence p also decides $\widetilde{A} \cap \gamma$;

(2) if $\alpha < \gamma$, then t_α is a κ-closed normal (δ, κ^+)-tree; and

(3) (i) if $\alpha < \beta < \gamma$ and $p \Vdash \alpha \in \widetilde{A}$ and $\alpha \mathrel{\widetilde{E}} \beta$, then $f_{\alpha\beta}$ is an isomorphism from t_α onto t_β;

(ii) otherwise, $f_{\alpha\beta} = \emptyset$.

If δ is a successor ordinal, then

$$\langle p, \widetilde{q} \rangle = \langle p, \langle f_{\alpha\beta} \mid \alpha < \beta < \kappa^+ \rangle \rangle \in \mathbb{R}$$

and $\langle p, \widetilde{q} \rangle \leq \langle p_\xi, \widetilde{q}_\xi \rangle$ for all $\xi < \lambda$. So suppose that δ is a limit ordinal. In order to construct a condition $\langle p^+, \widetilde{q}^+ \rangle \in \mathbb{R}$ such that $\langle p^+, \widetilde{q}^+ \rangle \leq \langle p_\xi, \widetilde{q}_\xi \rangle$ for all $\xi < \lambda$, it is enough if we can simultaneously solve the following extension problems.

(8.3.16 Suppose that $\alpha < \beta < \gamma$ and that $p \Vdash \alpha \in \widetilde{A}$ and $\alpha\ \widetilde{E}\ \beta$. Then we must extend t_α, t_β to κ-closed normal $(\delta+1, \kappa^+)$-trees t_α^+, t_β^+ such that $f_{\alpha\beta}$ can be extended to an isomorphism $g_{\alpha\beta} : t_\alpha^+ \to t_\beta^+$.

If $\mathrm{cf}(\delta) < \kappa$, then there is no difficulty, since then both t_α^+ and t_β^+ will be obtained by adjoining elements above *every* δ-branch of t_α, t_β respectively. So suppose that $\mathrm{cf}(\delta) = \kappa$. First we extend each of the trees t_α such that $p \Vdash \alpha \in \widetilde{A}$ to a κ-closed normal $(\delta + 1, \kappa^+)$-trees t_α^+. If $\beta < \gamma$ and $p \nVdash \beta \in \widetilde{A}$, then there exists a unique $\alpha < \beta$ such that $p \Vdash \alpha\ \widetilde{E}\ \beta$. We now choose t_β^+ so that $f_{\alpha\beta}$ can be extended to an isomorphism $g_{\alpha\beta} : t_\alpha^+ \to t_\beta^+$. □

Hence $\mathbb{R}$ preserves cofinalities and cardinals, and does not adjoin any new κ-sequences of ordinals. It follows that no new κ-sequences of ordinals are adjoined if we force with $\mathbb{P}_E$ over $M[G]$.

Now suppose that for some $\mu < \kappa^+$, there exists a nonidentity automorphism φ of T_μ in $M[G][H]$. By Lemma 8.3.13, we can assume that $\mu \in A$. Let $\widetilde{\varphi}$ be an $\mathbb{R}$-name for φ. Then there exists a condition $\langle p, \widetilde{q} \rangle \in \mathbb{R}$ and an element $a \in t_\mu^p$ such that

$$\langle p, \widetilde{q} \rangle \Vdash \widetilde{\varphi} : \widetilde{T}_\mu \to \widetilde{T}_\mu \text{ is an automorphism such that } \widetilde{\varphi}(a) \neq a.$$

Since $\mathbb{R}$ is κ^+-closed, we can inductively define a descending sequence of conditions $\langle \langle p_\xi, \widetilde{q}_\xi \rangle \mid \xi < \kappa \rangle$ such that

(i) $\langle p_0, \widetilde{q}_0 \rangle = \langle p, \widetilde{q} \rangle$;

(ii) $t_\mu^{p_{\xi+1}}$ is a proper end-extension of $t_\mu^{p_\xi}$; and

(iii) $\langle p_{\xi+1}, \widetilde{q}_{\xi+1} \rangle$ decides $\widetilde{\varphi} \restriction t_\mu^{p_\xi}$.

Let $t_\mu = \bigcup_{\xi<\kappa} t_\mu^{p_\xi}$ and let $\psi : t_\mu \to t_\mu$ be the nonidentity automorphism such that for each $\xi < \kappa$, $\langle p_{\xi+1}, \widetilde{q}_{\xi+1} \rangle \Vdash \widetilde{\varphi} \restriction t_\mu^{p_\xi} \subset \psi$. Note that t_μ is a κ-closed normal (η, κ^+)-tree for some η such that $\mathrm{cf}(\eta) = \kappa$. Arguing as in the proof of Lemma 8.3.10, we see that there exists a κ-closed normal $(\eta+1, \kappa^+)$-tree $t_\mu^+ \supset t_\mu$ such that

ψ cannot be extended to an automorphism of t^+_μ. And then arguing as in the proof of Claim 8.3.16, we see that there exists a condition $\langle p_\kappa, \widetilde{q}_\kappa\rangle \in \mathbb{R}$ such that

(1) $\langle p_\kappa, \widetilde{q}_\kappa\rangle \leq \langle p_\xi, \widetilde{q}_\xi\rangle$ for all $\xi < \kappa$; and

(2) $t^{p_\kappa}_\mu = t^+_\mu$.

But this is a contradiction.

A similar argument shows that if $\alpha < \beta < \kappa^+$ and α, β are not E-equivalent, then T_α and T_β remain nonisomorphic in $M[G][H]$. This completes the proof of Lemma 8.3.14. $\square$

Finally we will use a suitable reverse Easton iteration to complete the proof of Theorem 8.3.1. Let $V_0 \vDash GCH$. We shall inductively construct a sequence of forcing notions $\langle \mathbb{P}_\beta \mid \beta \in On\rangle$ satisfying Hypothesis 6.8.9. By the remark following Theorem 6.8.10, at successor stages β of the construction, we can assume inductively that $V^{\mathbb{P}_{\beta-1}} \vDash GCH$.

Case 1. If $\beta = 0$, then $\mathbb{P}_0 = \{\emptyset\}$ is the trivial notion of forcing.

Case 2. If β is a limit ordinal which is not inaccessible, then $\mathbb{P}_\beta$ is the inverse limit of $\langle \mathbb{P}_\gamma \mid \gamma < \beta\rangle$.

Case 3. If β is inaccessible, then $\mathbb{P}_\beta$ is the direct limit of $\langle \mathbb{P}_\gamma \mid \gamma < \beta\rangle$.

Case 4. Finally suppose that $\beta = \gamma + 1$ is a successor ordinal. First suppose that $\gamma = \kappa^+$ for some regular cardinal $\kappa \geq \omega$. Then we can assume inductively that $V^{\mathbb{P}_\gamma} \vDash GCH$. Let $\mathbb{Q}_\kappa \in V^{\mathbb{P}_\gamma}$ be the notion of forcing introduced in Definition 8.3.6. Then we set $\mathbb{P}_{\gamma+1} = \mathbb{P}_\gamma * \widetilde{\mathbb{Q}}_\kappa$. (As usual, $\widetilde{\mathbb{Q}}_\kappa$ denotes a $\mathbb{P}_\gamma$-name of the notion of forcing $\mathbb{Q}_\kappa \in V^{\mathbb{P}_\gamma}$.) Finally if γ does not have the form κ^+ for some regular cardinal $\kappa \geq \omega$, then we set $\mathbb{P}_{\gamma+1} = \mathbb{P}_\gamma * \check{\mathbb{P}}_0$.

Let $\mathbb{P}_\infty$ be the direct limit of $\langle \mathbb{P}_\beta \mid \beta \in On\rangle$; and for each $\beta \in On$, let $\tilde{\mathbb{P}}_{\beta\infty}$ be the canonically chosen $\mathbb{P}_\beta$-name for a proper class notion of forcing such that $\mathbb{P}_\infty$ is isomorphic to a dense sub-order of $\mathbb{P}_\beta * \tilde{\mathbb{P}}_{\beta\infty}$. Let G be a $\mathbb{P}_\infty$-generic filter over V_0 and let $V = V_0[G]$ be the corresponding generic extension. For each $\beta \in On$, let $G_\beta = G \cap \mathbb{P}_\beta$ and let $\mathbb{P}_{\beta\infty} = (\tilde{\mathbb{P}}_{\beta\infty})_{G_\beta}$. Then the following result is an immediate consequence of Theorem 6.8.10.

Lemma 8.3.17. (a) $\mathbb{P}_\infty$ *preserves cofinalities and cardinals.*

(b) *If* $\theta = \kappa^+$ *for some regular cardinal* $\kappa \geq \omega$, *then* $V[G_{\theta+1}] \vDash \mathbb{P}_{\theta+1\infty}$ *is* κ^{++}*-closed.*

(c) V *is a model of* $ZFC + GCH$.

□

Let κ be a regular cardinal and let $\theta = \kappa^+$. Let $\{T_\alpha \mid \alpha < \kappa^+\} \in V_0[G_{\theta+1}]$ be the set of trees which is adjoined by $\mathbb{Q}_\kappa$ at the θth stage of the iteration. Since $V[G_{\theta+1}] \vDash \mathbb{P}_{\theta+1\infty}$ is κ^{++}-closed, it follows that $\{T_\alpha \mid \alpha < \kappa^+\}$ remains a set of pairwise nonisomorphic rigid trees in V. Now let $E \in V$ be any equivalence relation on κ^+ and let $\mathbb{P}_E$ be the corresponding notion of forcing, which was introduced in Definition 8.3.11. Again using the fact that $V[G_{\theta+1}] \vDash \mathbb{P}_{\theta+1\infty}$ is κ^{++}-closed, we see that E, $\mathbb{P}_E \in V_0[G_{\theta+1}]$. We have already shown that $\mathbb{P}_E$ has the appropriate properties in $V_0[G_{\theta+1}]$. Thus it only remains to prove that these properties are preserved in V.

LEMMA 8.3.18. *In V, $\mathbb{P}_E$ preserves cofinalities and cardinals, and does not adjoin any new κ-sequences of ordinals. The following statements hold in $V^{\mathbb{P}_E}$.*

(a) *T_α is rigid for each $\alpha < \kappa^+$.*
(b) *If $\alpha < \beta < \kappa^+$, then $T_\alpha \simeq T_\beta$ iff $\alpha \mathrel{E} \beta$.*

PROOF. Since $|\mathbb{P}_E| = \kappa^+$, $\mathbb{P}_E$ preserves cofinalities and cardinals greater than κ^+. The remaining parts of the lemma correspond to combinatorial properties of $\mathbb{P}_E$ which are preserved under κ^{++}-closed forcing. For example, working inside V, suppose that $\alpha < \kappa^+$ and that $p \in \mathbb{P}_E$ satisfies

$$p \Vdash \widetilde{f} : T_\alpha \to T_\alpha \text{ is an automorphism.}$$

We can assume that $\widetilde{f}$ is a nice $\mathbb{P}_E$-name; i.e. that

$$\widetilde{f} = \bigcup \{\{\langle s, t\rangle\} \times A_{s,t} \mid s, t \in T_\alpha\},$$

where each $A_{s,t}$ is an antichain of $\mathbb{P}_E$. Since $V[G_{\theta+1}] \vDash \mathbb{P}_{\theta+1\infty}$ is κ^{++}-closed, it follows that $\widetilde{f} \in V_0[G_{\theta+1}]$. A moment's thought shows that, working within $V_0[G_{\theta+1}]$, we also have that

$$p \Vdash \widetilde{f} : T_\alpha \to T_\alpha \text{ is an automorphism;}$$

and so there exists $q \leq p$ such that within $V_0[G_{\theta+1}]$,

$$q \Vdash \widetilde{f}(t) = t \text{ for all } t \in T_\alpha.$$

This means that whenever $s \neq t$ are distinct elements of T_α, then q is incompatible with every element of $A_{s,t}$. Consequently, working within V, we also have that

$$q \Vdash \widetilde{f}(t) = t \text{ for all } t \in T_\alpha.$$

Hence T_α is rigid in $V^{\mathbb{P}_E}$. $\square$

This completes the proof of Theorem 8.3.1.

8.4. Notes

The material in this chapter first appeared in Hamkins-Thomas [**15**]. The proof of Theorem 8.3.1 relies heavily on the ideas of Jech [**18**], who used a similar bootstrap argument to prove the consistency of the existence of Suslin trees T such that $2^\omega < |\operatorname{Aut}(T)| < 2^{\omega_1}$.

CHAPTER 9

Bounding The Heights Of Automorphism Towers

A common feature of the consistency results in the earlier chapters of this book was that, in each case, it was essentially only necessary to generically adjoin a single centreless group. (For example, to prove the consistency of $\tau_{\omega_1} < \tau_{\omega_2}$, it was enough to adjoin a centreless group G of cardinality ω_2 such that $\tau(G) \geq (2^{\omega_1})^+$.) Thus most of our effort went into the construction of suitable notions of forcing for adjoining centreless groups with appropriate properties. On the other hand, in this chapter, we shall prove that it is consistent that $\tau_\lambda < 2^\lambda$ for every regular cardinal λ; and this requires finding a generic extension in which we have some understanding of every centreless group G of regular cardinality λ. Consequently, in this chapter, our notions of forcing will be extremely simple and all of our effort will go into analysing the automorphism tower of an arbitrary centreless group of regular cardinality λ in the corresponding generic extension. More specifically, fix some regular cardinal λ. Suppose that $V \vDash GCH$ and that θ is a regular cardinal such that $\theta \geq \lambda^{++}$. Let $\mathbb{P} = \mathrm{Fn}(\theta, 2, \lambda)$ and for each subset $X \subseteq \theta$, let $\mathbb{P} \restriction X = \mathrm{Fn}(X, 2, \lambda)$. Then we shall show that if $G \in V^{\mathbb{P}}$ is a centreless group of cardinality λ, then there exists a subset $X \in [\theta]^{\lambda^+}$ such that $G \in V^{\mathbb{P}\restriction X}$ and

$$\tau^{V^{\mathbb{P}}}(G) = \tau^{V^{\mathbb{P}\restriction X}}(G).$$

Since $V^{\mathbb{P}\restriction X} \vDash 2^\lambda = \lambda^+$, this implies that

$$V^{\mathbb{P}} \vDash \tau(G) < \lambda^{++} \leq \theta = 2^\lambda.$$

Then we shall use a suitable reverse Easton forcing to deal with all regular cardinals λ simultaneously.

9.1. The overspill argument

In Sections 9.2 and 9.3, we shall present a proof of the following result.

THEOREM 9.1.1. *It is consistent that $\tau_\lambda < 2^\lambda$ for all regular cardinals λ.*

In this section, we shall discuss the intuitive ideas which motivate the rather technical proof of Theorem 9.1.1. As usual, it is enough to find a notion of forcing which deals with a single regular cardinal λ, and then we can use a suitable reverse Easton forcing to complete the proof of Theorem 9.1.1. For the rest of this section, let M be a c.t.m. and let κ, λ, $\theta \in M$ be cardinals such that

(a) $\kappa^{<\kappa} = \kappa \leq \lambda$
(b) $2^\lambda = \lambda^+$
(c) $\lambda^{++} \leq \theta = \theta^\lambda$.

Recall that if if κ is an infinite cardinal and X is any set, then $\mathrm{Fn}(X, 2, \kappa)$ is the notion of forcing consisting of all functions p such that

(a) $\operatorname{dom} p \subseteq X$,
(b) $\operatorname{ran} p \subseteq 2$, and
(c) $|p| < \kappa$,

ordered by $q \leq p$ iff $q \supseteq p$. Let $\mathbb{P} = \mathrm{Fn}(\theta, 2, \kappa) \in M$. Clearly

$$\mathrm{Fn}(\theta, 2, \kappa) \simeq \mathrm{Fn}(\theta \times \lambda, 2, \kappa).$$

Hence, by Theorems 6.4.4 and 6.5.4, $M^{\mathbb{P}} \models 2^\lambda = \theta \geq \lambda^{++}$. In the rest of this section, we shall sketch a heuristic argument that if $G \in M^{\mathbb{P}}$ is a centreless group of cardinality λ, then $\tau^{M^{\mathbb{P}}}(G) < \lambda^{++}$. Our argument is based upon the fact that if X is any set such that $|X| \geq \kappa$, then $\mathbb{Q} = \mathrm{Fn}(X, 2, \kappa)$ is a weakly homogeneous notion of forcing. More specifically, we shall make use of Theorem 9.1.4. (The proof of Theorem 9.1.4 will be given at the end of this section.)

DEFINITION 9.1.2. A notion of forcing $\mathbb{Q}$ is *weakly homogeneous* if for each pair of conditions p, $q \in \mathbb{Q}$, there exists an automorphism $\psi \in \mathrm{Aut}(\mathbb{Q})$ such that $\psi(p)$ and q are compatible.

EXAMPLE 9.1.3. We shall show that if κ is an infinite cardinal and X is any set such that $|X| \geq \kappa$, then $\mathbb{Q} = \mathrm{Fn}(X, 2, \kappa)$ is a weakly homogeneous notion of forcing. To see this, note that there is a natural action of $\mathrm{Sym}(X)$ as a group of automorphisms of $\mathbb{Q}$, defined by

$$\psi(p) = p \circ \psi^{-1}$$

for each $\psi \in \mathrm{Sym}(X)$ and $p \in \mathbb{Q}$. Also note that if $\psi \in \mathrm{Sym}(X)$ and $p \in \mathbb{Q}$, then $\operatorname{dom} \psi(p) = \psi[\operatorname{dom} p]$. Now let p, $q \in \mathbb{Q}$ be an arbitrary pair of conditions. Then there exists a permutation $\psi \in \mathrm{Sym}(X)$ such that $\psi[\operatorname{dom} p] \cap \operatorname{dom} q = \emptyset$ and it follows that $\psi(p)$ and q are compatible.

Theorem 9.1.4. *Let M be a c.t.m. and let $\mathbb{Q} \in M$ be a weakly homogeneous notion of forcing. If $\varphi(v_1, \dots, v_n)$ is any formula and $a_1, \dots, a_n \in M$, then either*

$$1_{\mathbb{Q}} \Vdash \varphi(\check{a}_1, \dots, \check{a}_n) \quad \text{or} \quad 1_{\mathbb{Q}} \Vdash \neg\varphi(\check{a}_1, \dots, \check{a}_n).$$

Our argument will also make use of the following basic results on the product structure of $\mathbb{P} = \mathrm{Fn}(\theta, 2, \kappa)$.

Theorem 9.1.5. *With the above hypotheses, let $\mathbb{P} = \mathrm{Fn}(\theta, 2, \kappa) \in M$ and let H be a $\mathbb{P}$-generic filter over M. Suppose that $X \in M$ satisfies $X \subseteq \theta$ and let $Y = \theta \smallsetminus X$.*

(a) $\mathrm{Fn}(\theta, 2, \kappa) \simeq \mathrm{Fn}(X, 2, \kappa) \times \mathrm{Fn}(Y, 2, \kappa)$.
(b) *Let $H_X = H \cap \mathrm{Fn}(X, 2, \kappa)$ and $H_Y = H \cap \mathrm{Fn}(Y, 2, \kappa)$. Then H_X is a $\mathrm{Fn}(X, 2, \kappa)$-generic filter over M and H_Y is a $\mathrm{Fn}(Y, 2, \kappa)$-generic filter over $M[H_X]$. Furthermore, $M[H] = M[H_X][H_Y]$.*

Proof. (a) It is easily checked that the map

$$p \mapsto \langle p \restriction X, p \restriction Y \rangle$$

is a isomorphism from $\mathrm{Fn}(\theta, 2, \kappa)$ onto $\mathrm{Fn}(X, 2, \kappa) \times \mathrm{Fn}(Y, 2, \kappa)$.

(b) This is Theorem VIII.2.1 of Kunen [**26**]. □

Notice that, since $\mathrm{Fn}(X, 2, \kappa) \in M$ is κ-closed, it follows that

$$\mathrm{Fn}(Y, 2, \kappa)^{M[H_X]} = \mathrm{Fn}(Y, 2, \kappa)^M.$$

Lemma 9.1.6. *With the above hypotheses, let $\mathbb{P} = \mathrm{Fn}(\theta, 2, \kappa) \in M$ and let H be a $\mathbb{P}$-generic filter over M. Suppose that $\mu \in M$ is a cardinal such that $\mu \leq \theta$ and that $S \in M[H]$ satisfies $S \subseteq \mu$. Then there exists $X \in M$ of cardinality μ such that $X \subseteq \theta$ and $S \in M[H \cap \mathrm{Fn}(X, 2, \kappa)]$.*

PROOF. Since $\mathbb{P} = \mathrm{Fn}(\theta, 2, \kappa) \in M$ is κ-closed, it follows that if $\mu < \kappa$, then $S \in M$. Hence we can suppose that $\mu \geq \kappa$. Let

$$\tau = \bigcup \{\{\check{\alpha}\} \times A_\alpha \mid \alpha \in \mu\} \in M$$

be a nice $\mathbb{P}$-name for a subset of μ such that $\tau_H = S$. Since $\mathbb{P} \in M$ has the κ^+-*c.c*, it follows that each antichain A_α has cardinality at most κ. Hence there exists a subset $X \in \mathcal{P}^M(\theta)$ of cardinality μ such that $A_\alpha \subseteq \mathrm{Fn}(X, 2, \kappa)$ for each $\alpha \in \mu$. Let $H_X = H \cap \mathrm{Fn}(X, 2, \kappa)$. Then τ is a nice $\mathrm{Fn}(X, 2, \kappa)$-name for a subset of μ such that $\tau_{H_X} = S$. In particular, $S \in M[H_X] = M[H \cap \mathrm{Fn}(X, 2, \kappa)]$. □

We are now ready to complete our heuristic argument that if $G \in M^{\mathbb{P}}$ is a centreless group of cardinality λ, then $\tau^{M^{\mathbb{P}}}(G) < \lambda^{++}$. Let $M^{\mathbb{P}} = M[H]$, where $H \subseteq \mathbb{P}$ is a $\mathbb{P}$-generic filter over M, and let $G \in M^{\mathbb{P}}$ be a centreless group of cardinality λ. To simplify notation, we shall initially assume that $G \in M$. Now suppose that $M^{\mathbb{P}} \vDash \tau(G) \geq \lambda^{++}$ and let γ be any ordinal such that $\gamma < \lambda^{++}$. Then, in $M^{\mathbb{P}}$, we have a strictly increasing tower of groups

$$G = G_0^{M^{\mathbb{P}}} \lhd G_1^{M^{\mathbb{P}}} \lhd \cdots \lhd G_\alpha^{M^{\mathbb{P}}} \lhd \cdots \lhd G_\gamma^{M^{\mathbb{P}}} \lhd G_{\gamma+1}^{M^{\mathbb{P}}}$$

and it seems reasonable to suppose that this will be reflected down to some submodel of $M^{\mathbb{P}}$ of the form $N = M[H \cap \mathrm{Fn}(X_\gamma, 2, \kappa)]$, where $X_\gamma \in [\theta]^{\lambda^+} \cap M$. Roughly speaking, for each $\alpha \leq \gamma$, we should first choose an element $g_\alpha \in G_{\alpha+1}^{M^{\mathbb{P}}} \smallsetminus G_\alpha^{M^{\mathbb{P}}}$, and then we should find a subset $X_\gamma \in [\theta]^{\lambda^+} \cap M$ such that $g_\alpha \in N$ and G_α^N is a g_α-invariant subgroup of $G_\alpha^{\mathbb{P}}$ for each $\alpha \leq \gamma$. Then $g_\alpha \restriction G_\alpha^N \in G_{\alpha+1}^N \smallsetminus G_\alpha^N$ for all $\alpha \leq \gamma$ and so $N \vDash \tau(G) > \gamma$. (The reader is advised not to examine the previous two sentences too carefully. In particular, he should not be concerned if he notices that G_α^N is probably not even a subgroup of $G_\alpha^{M^{\mathbb{P}}}$ when $\alpha \geq 2$.) Thus for each ordinal $\gamma < \lambda^{++}$, there should exist a subset $X_\gamma \in [\theta]^{\lambda^+} \cap M$ such that

$$M[H \cap \mathrm{Fn}(X_\gamma, 2, \kappa)] \vDash \tau(G) > \gamma.$$

Now let $\mathbb{Q} = \mathrm{Fn}(\lambda^+, 2, \theta)$ and notice that $\mathrm{Fn}(X_\gamma, 2, \kappa) \simeq \mathbb{Q}$ for all $\gamma < \lambda^{++}$. Hence for each $\gamma < \lambda^{++}$, there exists a $\mathbb{Q}$-generic filter K_γ over M such that $M[H \cap \mathrm{Fn}(X_\gamma, 2, \kappa)] = M[K_\gamma]$. Hence by Theorem 9.1.4,

$$1_{\mathbb{Q}} \Vdash_{\mathbb{Q}} \tau(\check{G}) > \check{\gamma}$$

for all $\gamma < \lambda^{++}$. But this means that if K is any $\mathbb{Q}$-generic filter over M, then

$$M[K] \vDash \tau(G) \geq \lambda^{++} = \left(2^{\lambda}\right)^{+},$$

which contradicts Corollary 3.3.2. Thus $\tau^{M^{\mathbb{P}}}(G) < \lambda^{++}$ for every centreless group $G \in M$ of cardinality λ.

Now suppose that $G \in M^{\mathbb{P}}$ is any centreless group of cardinality λ. Then there exists a subset $Y \in [\theta]^{\lambda} \cap M$ such that $G \in N = M[H \cap \mathrm{Fn}(Y, 2, \kappa)]$ and κ, λ, θ continue to satisfy our hypotheses in N. Since $M^{\mathbb{P}} = N[H \cap \mathrm{Fn}(\theta \smallsetminus Y, 2, \kappa)]$ and $\mathrm{Fn}(\theta \smallsetminus Y, 2, \kappa) \simeq \mathbb{P}$, our previous argument shows that $\tau^{M^{\mathbb{P}}}(G) < \lambda^{++}$.

We shall present a rigorous proof of Theorem 9.1.1 in Sections 9.2 and 9.3. The main technical difficulty will concern the problem of comparing the values of $\tau(G)$ within various generic extensions $M^{\mathbb{P}}$ of the c.t.m. M. We have already dealt with an easy case of this problem in Section 6.7. There we showed that if $\mathbb{P}$ is $|G|^{+}$-closed, then $\tau^{M^{\mathbb{P}}}(G) = \tau^{M}(G)$. The proof was relatively easy because, with this hypothesis on $\mathbb{P}$, we could show inductively that $G_{\alpha}^{M^{\mathbb{P}}} = G_{\alpha}^{M}$ for all ordinals α. In Section 9.2, we shall develop a technique for comparing $\tau^{M^{\mathbb{P}}}(G)$ and $\tau^{M}(G)$ in the situation when $G_1^M = \mathrm{Aut}^M(G)$ is a proper subgroup of $G_1^{M^{\mathbb{P}}} = \mathrm{Aut}^{M^{\mathbb{P}}}(G)$. (In this case, there is already a difficulty in trying to relate $G_2^M = \mathrm{Aut}^M(G_1^M)$ and $G_2^{M^{\mathbb{P}}} = \mathrm{Aut}^{M^{\mathbb{P}}}(G_1^{M^{\mathbb{P}}})$.) Then we shall complete the proof of Theorem 9.1.1 in Section 9.3.

In the remainder of this section, we shall present a proof of Theorem 9.1.4. The proof is based on the observation that if $\mathbb{Q}$ is any notion of forcing, then $\mathrm{Aut}(\mathbb{Q})$ acts naturally on the entire $\mathbb{Q}$-forcing apparatus. For example, if $\psi \in \mathrm{Aut}(\mathbb{Q})$, then for each $\mathbb{Q}$-name τ, we can inductively define the associated $\mathbb{Q}$-name $\psi(\tau)$ by

$$\psi(\tau) = \{\langle \psi(\sigma), \psi(p) \rangle \mid \langle \sigma, p \rangle \in \tau\}.$$

LEMMA 9.1.7. *Let $\mathbb{Q}$ be a notion of forcing and let $\psi \in \mathrm{Aut}(\mathbb{Q})$. Suppose that $\varphi(v_1, \dots, v_n)$ is any formula and that $\tau_1, \dots, \tau_n$ are $\mathbb{Q}$-names. Then for all $p \in \mathbb{Q}$,*

$$p \Vdash \varphi(\tau_1, \dots, \tau_n) \quad \textit{iff} \quad \psi(p) \Vdash \varphi(\psi(\tau_1), \dots, \psi(\tau_n)).$$

The definition of the forcing relation $\Vdash$ in Chapter 6 involved the collection of all $\mathbb{Q}$-generic filters. Using this definition, Lemma 9.1.7 is not quite obvious. However, there is an alternative approach to the forcing relation for which the analogue of

Lemma 9.1.7 is completely obvious. (The relation $\Vdash^*$ is essentially just the result of a tedious inductive analysis of the forcing relation $\Vdash$, which can be carried out within the c.t.m. M; i.e. an analysis which avoids the use of any objects which lie outside M.)

DEFINITION 9.1.8. Let $\mathbb{Q}$ be a notion of forcing. If $p \in \mathbb{Q}$, $\varphi(v_1, \dots, v_n)$ is any formula and $\tau_1, \dots, \tau_n$ are $\mathbb{Q}$-names, then we define the relation

$$p \Vdash^* \varphi(\tau_1, \dots, \tau_n)$$

by induction on the complexity of $\varphi(v_1, \dots, v_n)$ as follows.

(a) $p \Vdash^* \tau_1 = \tau_2$ iff the following two clauses hold.
  (i) For all $\langle \pi_1, s_1 \rangle \in \tau_1$ and $q \leq p$, there exists $r \leq q$ such that either $r \perp s_1$ or there exists $\langle \pi_2, s_2 \rangle \in \tau_2$ with $r \leq s_2$ and $r \Vdash^* \pi_1 = \pi_2$.
  (ii) For all $\langle \pi_2, s_2 \rangle \in \tau_2$ and $q \leq p$, there exists $r \leq q$ such that either $r \perp s_2$ or there exists $\langle \pi_1, s_1 \rangle \in \tau_1$ with $r \leq s_1$ and $r \Vdash^* \pi_1 = \pi_2$.

(b) $p \Vdash^* \tau_1 \in \tau_2$ iff for all $q \leq p$, there exists $r \leq q$ and $\langle \pi, s \rangle \in \tau_2$ such that $r \leq s$ and $r \Vdash^* \pi = \tau_1$.

(c) $p \Vdash^* \varphi(\tau_1, \dots, \tau_n) \wedge \psi(\tau_1, \dots, \tau_n)$ iff $p \Vdash^* \varphi(\tau_1, \dots, \tau_n)$ and $p \Vdash^* \psi(\tau_1, \dots, \tau_n)$.

(d) $p \Vdash^* \neg\varphi(\tau_1, \dots, \tau_n)$ iff there does not exist a condition $q \leq p$ such that $q \Vdash^* \varphi(\tau_1, \dots, \tau_n)$.

(e) $p \Vdash^* \exists x \varphi(x, \tau_1, \dots, \tau_n)$ iff for all $q \leq p$, there exists $r \leq q$ and a $\mathbb{Q}$-name σ such that $r \Vdash^* \varphi(\sigma, \tau_1, \dots, \tau_n)$.

As we mentioned earlier, the following analogue of Lemma 9.1.7 is completely obvious.

LEMMA 9.1.9. *Let $\mathbb{Q}$ be a notion of forcing and let $\psi \in \mathrm{Aut}(\mathbb{Q})$. Suppose that $\varphi(v_1, \dots, v_n)$ is any formula and that $\tau_1, \dots, \tau_n$ are $\mathbb{Q}$-names. Then for all $p \in \mathbb{Q}$,*

$$p \Vdash^* \varphi(\tau_1, \dots, \tau_n) \quad \textit{iff} \quad \psi(p) \Vdash^* \varphi(\psi(\tau_1), \dots, \psi(\tau_n)).$$

$\square$

Clearly Lemma 9.1.7 is an immediate consequence of Lemma 9.1.9 and Theorem 9.1.10.

THEOREM 9.1.10. *Let $\mathbb{Q}$ be a notion of forcing. Let $\varphi(v_1, \dots, v_n)$ be any formula and let $\tau_1, \dots, \tau_n$ be $\mathbb{Q}$-names. Then for all $p \in \mathbb{Q}$,*

$$p \Vdash \varphi(\tau_1, \dots, \tau_n) \quad \textit{iff} \quad p \Vdash^* \varphi(\tau_1, \dots, \tau_n).$$

PROOF. For example, see Kunen [**26**, Section VIII.3]. $\square$

We are now ready to present the proof of Theorem 9.1.4. Let M be a c.t.m. and let $\mathbb{Q} \in M$ be a weakly homogeneous notion of forcing. Let $\varphi(v_1, \dots, v_n)$ be any formula and let $a_1, \dots, a_n \in M$. Suppose that

$$1_{\mathbb{Q}} \nVdash \varphi(\check{a}_1, \dots, \check{a}_n) \quad \text{and} \quad 1_{\mathbb{Q}} \nVdash \neg\varphi(\check{a}_1, \dots, \check{a}_n).$$

Then there exist conditions p, $q \in \mathbb{Q}$ such that

$$p \Vdash \varphi(\check{a}_1, \dots, \check{a}_n) \quad \text{and} \quad q \Vdash \neg\varphi(\check{a}_1, \dots, \check{a}_n).$$

Let $\psi \in \operatorname{Aut}^M(\mathbb{Q})$ be an automorphism such that $\psi(p)$ and q are compatible. An easy induction shows that $\psi(\check{b}) = \check{b}$ for all $b \in M$. Hence, by Lemma 9.1.7, $\psi(p) \Vdash \varphi(\check{a}_1, \dots, \check{a}_n)$. But this means that if $r \leq \psi(p), q$, then

$$r \Vdash \varphi(\check{a}_1, \dots, \check{a}_n) \wedge \neg\varphi(\check{a}_1, \dots, \check{a}_n),$$

which is a contradiction.

9.2. Conservative extensions

Suppose that M_1 is a c.t.m. and that $G \in M_1$ is a centreless group. Let $\mathbb{P} \in M_1$ be a notion of forcing and let $M_2 = M_1^{\mathbb{P}}$ be the corresponding generic extension. In this section, we shall develop a technique for comparing $\tau^{M_1}(G)$ and $\tau^{M_2}(G)$ in the case when $G_1^{M_1} = \operatorname{Aut}^{M_1}(G)$ is a proper subgroup of $G_1^{M_2} = \operatorname{Aut}^{M_2}(G)$. In this case, there is already a difficulty in trying to relate $G_2^{M_1} = \operatorname{Aut}^{M_1}(G_1^{M_1})$ and $G_2^{M_2} = \operatorname{Aut}^{M_2}(G_1^{M_2})$. However, suppose that every automorphism φ of $G_1^{M_1}$ extends to an automorphism φ^+ of $G_1^{M_2}$. Since $G \leqslant G_1^{M_1}$, Theorem 1.1.10 implies that there must be a unique such extension φ^+ for each $\varphi \in G_2^{M_1} = \operatorname{Aut}^{M_1}(G_1^{M_1})$; and so we can define a canonical embedding $\pi : G_2^{M_1} \to G_2^{M_2}$ by $\pi(\varphi) = \varphi^+$. If we now also suppose that every automorphism of $\pi[G_2^{M_1}]$ extends to an automorphism of $G_2^{M_2}$, then we next obtain a canonical embedding of $G_3^{M_1}$ into $G_3^{M_2}$. Continuing in this fashion, we arrive at the notion of a G-conservative extension.

DEFINITION 9.2.1. Let M_1 be a c.t.m. Let $\mathbb{P} \in M_1$ be a notion of forcing and let $M_2 = M_1^{\mathbb{P}}$ be the corresponding generic extension. If $G \in M_1$ is a centreless group, then M_2 is said to be a *G-conservative extension of* M_1 iff for each $\alpha \in On$, M_2 contains an embedding $\pi_\alpha : G_\alpha^{M_1} \to G_\alpha^{M_2}$ such that the following conditions are satisfied.

(a) If $\beta < \alpha$, then $\pi_\beta \subseteq \pi_\alpha$.
(b) $\pi_0 = id_G$.
(c) If α is a limit ordinal, then $\pi_\alpha = \bigcup_{\beta<\alpha} \pi_\beta$.
(d) Finally suppose that $\alpha = \beta + 1$ and that $g \in G_\alpha^{M_1} = \mathrm{Aut}^{M_1}(G_\beta^{M_1})$. Then $\pi_\alpha(g) \in G_\alpha^{M_2} = \mathrm{Aut}^{M_2}(G_\beta^{M_2})$ extends the automorphism $\pi_\beta g \pi_\beta^{-1}$ of the subgroup $\pi_\beta[G_\beta^{M_1}] \leqslant G_\beta^{M_2}$.

In this case, we shall say that π_α is the *canonical embedding* of $G_\alpha^{M_1}$ into $G_\alpha^{M_2}$.

REMARK 9.2.2. To get a better understanding of Definition 9.2.1 and to see that there is a unique canonical embedding $\pi_\alpha : G_\alpha^{M_1} \to G_\alpha^{M_2}$ for each $\alpha \in On$, we shall spell out exactly what needs to be checked when verifying that M_2 is a G-conservative extension of M_1. So suppose inductively that we have shown that there exists a canonical embedding $\pi_\beta : G_\beta^{M_1} \to G_\beta^{M_2}$ for each $\beta < \alpha$. Clearly no problems arise if α is a limit ordinal. So suppose that $\alpha = \beta + 1$. Then it is enough to check that for each $g \in G_\alpha = \mathrm{Aut}^{M_1}(G_\beta^{M_1})$, the automorphism $\pi_\beta g \pi_\beta^{-1}$ of $\pi_\beta[G_\beta^{M_1}]$ can be extended to an automorphism h of $G_\beta^{M_2}$. To see this, notice that

$$G_0 \leqslant \pi_\beta[G_\beta^{M_1}] \leqslant G_\beta^{M_2},$$

and so Theorem 1.1.10 implies that h is the unique automorphism of $G_\beta^{M_2}$ which extends $\pi_\beta g \pi_\beta^{-1}$. Hence we can define the embedding $\pi_\alpha : G_\alpha^{M_1} \to G_\alpha^{M_2}$ by

$$\pi_\alpha(g) = \text{ the unique } h \in G_\alpha^{M_2} \text{ such that } \pi_\beta g \pi_\beta^{-1} \subseteq h.$$

It only remains to verify that $\pi_\beta \subseteq \pi_\alpha$. To see this, let $a \in G_\beta^{M_1}$. Then a is identified with the corresponding inner automorphism $i_a \in G_\alpha^{M_1} = \mathrm{Aut}^{M_1}(G_\beta^{M_1})$, It is easily checked that $\pi_\beta i_a \pi_\beta^{-1} \subseteq i_{\pi_\beta(a)}$ and hence $\pi_\alpha(i_a) = i_{\pi_\beta(a)}$. Since $\pi_\beta(a)$ is identified with $i_{\pi_\beta(a)}$ within $G_\alpha^{M_2}$, it follows that $\pi_\beta \subseteq \pi_\alpha$.

LEMMA 9.2.3. *Suppose that M_2 is a G-conservative extension of M_1 and let $\pi_\alpha : G_\alpha^{M_1} \to G_\alpha^{M_2}$ be the canonical embedding for each $\alpha \in On$. If $\beta \leq \alpha$, then*

$$\pi_\alpha[G_\alpha^{M_1}] \cap G_\beta^{M_2} = \pi_\beta[G_\beta^{M_1}]$$

Hence $\tau^{M_1}(G) \leq \tau^{M_2}(G)$.

PROOF. Fix some ordinal α. We shall argue by induction on $\beta \leq \alpha$ that

$$\pi_\alpha[G_\alpha^{M_1}] \cap G_\beta^{M_2} = \pi_\beta[G_\beta^{M_1}].$$

Clearly the result holds for $\beta = 0$ and no problems arise when β is a limit ordinal. So suppose that $\beta = \gamma + 1$. First note that since $\pi_\beta \subseteq \pi_\alpha$, it follows that

$$\pi_\beta[G_\beta^{M_1}] \leqslant \pi_\alpha[G_\alpha^{M_1}] \cap G_\beta^{M_2}.$$

Now suppose that $h \in \pi_\alpha[G_\alpha^{M_1}] \cap G_\beta^{M_2}$. Then h normalises

$$\pi_\alpha[G_\alpha^{M_1}] \cap G_\gamma^{M_2} = \pi_\gamma[G_\gamma^{M_1}].$$

Let $g \in G_\alpha^{M_1}$ be the element such that $\pi_\alpha(g) = h$. Then g normalises $G_\gamma^{M_1}$. By Proposition 4.1.2, $G_{\gamma+1}^{M_1}$ is the normaliser of $G_\gamma^{M_1}$ in $G_\alpha^{M_1}$. Hence

$$h = \pi_\alpha(g) \in \pi_\alpha[G_\beta^{M_1}] = \pi_\beta[G_\beta^{M_1}].$$

Finally to see that $\tau^{M_1}(G) \leq \tau^{M_2}(G)$, let $\tau^{M_1}(G) = \tau$ and for each $\beta < \tau$, let $g_\beta \in G_{\beta+1}^{M_1} \smallsetminus G_\beta^{M_1}$. Then $\pi_{\beta+1}(g_\beta) \in G_{\beta+1}^{M_2} \smallsetminus G_\beta^{M_2}$ and hence $\tau^{M_2}(G) > \beta$. □

There exist examples of G-conservative extensions such that $\tau^{M_1}(G) < \tau^{M_2}(G)$. For example, suppose that $G \in M_1$ is a complete group. Then it is clear that if $\mathbb{P} \in M_1$ is any notion of forcing, then $M_2 = M_1^{\mathbb{P}}$ is a G-conservative extension of M_1. In particular, this is true if $\mathbb{P}$ is a notion of forcing which adjoins an outer automorphism of G; and in this case, we have that $\tau^{M_1}(G) < \tau^{M_2}(G)$.

Now suppose that M is a c.t.m. and that κ, λ, $\theta \in M$ are cardinals such that

(a) $\kappa^{<\kappa} = \kappa \leq \lambda$
(b) $2^\lambda = \lambda^+$
(c) $\lambda^{++} \leq \theta = \theta^\lambda$.

If $G \in M$ is a centreless group of cardinality λ and $\mathbb{P} = \mathrm{Fn}(\theta, 2, \kappa) \in M$, then there is no reason to suppose that $M^{\mathbb{P}}$ is a G-conservative extension of M. For example, if $\kappa = \omega$ and $\lambda = 2^\omega$, then Theorem 6.6.16 says that there exists a centreless group $H \in M$ of cardinality λ such that $\tau^M(H) > \tau^{M^{\mathbb{P}}}(H)$ and so $M^{\mathbb{P}}$ is not an H-conservative extension of M. However, in the next section, we shall prove that if $\mathbb{Q} = \mathrm{Fn}(\lambda^+, 2, \kappa)$ and $G \in M$ is any centreless group of cardinality λ, then $M^{\mathbb{P}}$ is a G-conservative extension of $M^{\mathbb{Q}}$. Furthermore, we shall also prove that $\tau^{M^{\mathbb{Q}}}(G) = \tau^{M^{\mathbb{P}}}(G)$. Since $M^{\mathbb{Q}} \vDash 2^\lambda = \lambda^+$, it follows that $\tau^{M^{\mathbb{Q}}}(G) < \lambda^{++}$ and hence that $M^{\mathbb{P}} \vDash \tau(G) < \lambda^{++} \leq \theta = 2^\lambda$.

We shall end this section with a slightly technical result, which will be used repeatedly in the proof of Theorem 9.3.1. We have already mentioned that if $\pi_\beta : G_\beta^{M_1} \to G_\beta^{M_2}$ is a canonical embedding, then there exists a canonical embedding $\pi_{\beta+1} : G_{\beta+1}^{M_1} \to G_{\beta+1}^{M_2}$ iff for each $g \in \mathrm{Aut}^{M_1}(G_\beta^{M_1})$, the automorphism $\pi_\beta g \pi_\beta^{-1}$ of $\pi_\beta[G_\beta^{M_1}]$ can be extended to an automorphism h of $G_\beta^{M_2}$. The next lemma says that it is enough to check that for each $g \in \mathrm{Aut}^{M_1}(G_\beta^{M_1})$, the embedding

$$\pi_\beta \circ (g \restriction G) \circ \pi_\beta^{-1} = \pi_\beta \circ (g \restriction G)$$

of $\pi_\beta[G] = G$ into $G_\beta^{M_2}$ can be extended to an automorphism h of $G_\beta^{M_2}$.

LEMMA 9.2.4. *Suppose that M_1 is a c.t.m. and that $G \in M_1$ is a centreless group. Let M_2 be a generic extension of M_1 and suppose that there exists a canonical embedding $\pi_\gamma : G_\gamma^{M_1} \to G_\gamma^{M_2}$ for each $\gamma \leq \beta$. If the elements $g \in \mathrm{Aut}^{M_1}(G_\beta^{M_1})$ and $h \in \mathrm{Aut}^{M_2}(G_\beta^{M_2})$ satisfy*

(a) *$h[G] \subseteq \pi_\beta[G_\beta^{M_1}]$ and*

(b) *$\pi_\beta^{-1} \circ (h \restriction G) \subseteq g$,*

then $\pi_\beta g \pi_\beta^{-1} \subseteq h$.

PROOF. This is equivalent to the statement that $g \subseteq \pi_\beta^{-1} h \pi_\beta$. We will prove this inclusion for $g \restriction G_\gamma^{M_1}$ by induction on $\gamma \leq \beta$. Since $\pi_\beta \restriction G$ is the identity map, clause (b) says that the result holds when $\gamma = 0$. Once again, no difficulties arise when γ is a limit ordinal. Suppose that $\gamma = \xi + 1$ and let $\varphi \in G_{\xi+1}^{M_1}$ be any element. For each $a \in G_\xi^{M_1}$, let $a_\varphi \in G_\xi^{M_1}$ be the element such that $\varphi a \varphi^{-1} = a_\varphi$; and for each $\psi \in G_{\xi+1}^{M_1}$, let $\overline{\psi} = \pi_\beta(\psi)$. Then for each $a \in G_\xi^{M_1}$, we have that

$$\overline{g(\varphi)} \cdot \overline{g(a)} \cdot \left(\overline{g(\varphi)}\right)^{-1} = \overline{g(a_\varphi)}$$

and that

$$h(\overline{\varphi}) \cdot h(\overline{a}) \cdot h(\overline{\varphi})^{-1} = h(\overline{a_\varphi}).$$

By the inductive hypothesis, $g \restriction G_\xi^{M_1} \subseteq \pi_\beta^{-1} h \pi_\beta$ and this implies that $h(\overline{a}) = \overline{g(a)}$ and that $h(\overline{a_\varphi}) = \overline{g(a_\varphi)}$. Thus for all elements $d \in (h \circ \pi_\beta)[G_\xi^{M_1}]$, we have that

$$h(\overline{\varphi})^{-1} \cdot \overline{g(\varphi)} \in C_{G_\beta^{M_2}}(d).$$

Since $G \leqslant \pi_\beta[G_\xi^{M_1}]$, this implies that

$$h(\overline{\varphi})^{-1} \cdot \overline{g(\varphi)} \in C_{G_\beta^{M_2}}(h[G]).$$

By Theorem 1.1.10, $C_{G_\beta^{M_2}}(G) = 1$. Applying the automorphism h of $G_\beta^{M_2}$, we obtain that $C_{G_\beta^{M_2}}(h[G]) = 1$ and hence $h(\overline{\varphi}) = \overline{g(\varphi)}$. Consequently the result also holds for $\xi + 1$. □

9.3. The reflection argument

In this section, we shall present the proof of Theorem 9.1.1. As we shall see, Theorem 9.1.1 is an easy consequence of the following result, together with a suitable reverse Easton forcing.

THEOREM 9.3.1. *Let M be a c.t.m. and let $\kappa \leq \lambda < \theta$ be regular cardinals which satisfy the following conditions.*

(a) $\kappa^{<\kappa} = \kappa$.

(b) $2^\lambda = \lambda^+$.

(c) $\theta \geq \lambda^{++}$.

Let $\mathbb{P} = \mathrm{Fn}(\theta, 2, \kappa)$ and let $\mathbb{Q} = \mathrm{Fn}(\lambda^+, 2, \kappa)$. Then for every centreless group $G \in M$ of cardinality λ:

(1) *$M^{\mathbb{P}}$ is a G-conservative extension of $M^{\mathbb{Q}}$.*

(2) $\tau^{M^{\mathbb{P}}}(G) = \tau^{M^{\mathbb{Q}}}(G) < \lambda^{++}$.

Throughout the proof of Theorem 9.3.1, the notation $\mathcal{P}(\theta)$ and $[\theta]^{\lambda^+}$ will always refer to $\mathcal{P}(\theta) \cap M$ and $[\theta]^{\lambda^+} \cap M$ respectively. Fix some $\mathbb{P}$-generic filter H over M. Then for each $X \in \mathcal{P}(\theta)$, we define $M(X) = M[H \cap \mathrm{Fn}(X, 2, \kappa)]$. We shall make repeated use of the following easy consequence of Theorem 9.1.4.

LEMMA 9.3.2. *Suppose that X, Y, $Z \in [\theta]^{\leq \lambda^+}$ are such that $X \subset Y$, Z and $|Y \smallsetminus X| = |Z \smallsetminus X| = \lambda^+$. If $\varphi(v_1, \dots, v_n)$ is any formula and $a_1, \dots, a_n \in M(X)$, then*

$$M(Y) \vDash \varphi(a_1, \dots, a_n) \quad \textit{iff} \quad M(Z) \vDash \varphi(a_1, \dots, a_n).$$

Proof. Note that $M(Y)$ and $M(Z)$ are generic extensions of $M(X)$ with respect to the notions of forcing $\mathrm{Fn}(Y \smallsetminus X, 2, \kappa)^M$, $\mathrm{Fn}(Z \smallsetminus X, 2, \kappa)^M$ respectively; and also that

$$\mathrm{Fn}(Y \smallsetminus X, 2, \kappa)^M = \mathrm{Fn}(Y \smallsetminus X, 2, \kappa)^{M(X)} \simeq \mathrm{Fn}(\lambda^+, 2, \kappa)^{M(X)}$$

and

$$\mathrm{Fn}(Z \smallsetminus X, 2, \kappa)^M = \mathrm{Fn}(Z \smallsetminus X, 2, \kappa)^{M(X)} \simeq \mathrm{Fn}(\lambda^+, 2, \kappa)^{M(X)}.$$

Working inside $M(X)$, we can apply Theorem 9.1.4 to the weakly homogeneous notion of forcing $\mathbb{Q} = \mathrm{Fn}(\lambda^+, 2, \kappa)^{M(X)}$ and hence obtain that

$$\begin{aligned} M(Y) \vDash \varphi(a_1, \ldots, a_n) \quad &\text{iff} \quad 1_{\mathbb{Q}} \Vdash_{\mathbb{Q}} \varphi(\check{a}_1, \ldots, \check{a}_n) \\ &\text{iff} \quad M(Z) \vDash \varphi(a_1, \ldots, a_n). \end{aligned}$$

□

In particular, by considering the special case when $X = \emptyset$ and $Z = \lambda^+$, we obtain the following result.

Lemma 9.3.3. *Suppose that $Y \in [\theta]^{\lambda^+}$. If $\varphi(v_1, \ldots, v_n)$ is any formula and $a_1, \ldots, a_n \in M$, then*

$$M(Y) \vDash \varphi(a_1, \ldots, a_n) \quad \textit{iff} \quad M(\lambda^+) \vDash \varphi(a_1, \ldots, a_n).$$

□

Theorem 9.3.1 is an easy consequence of the following result.

Lemma 9.3.4. *If $G \in M$ be a centreless group of cardinality λ, then for each $\alpha \in On$, the following statements are true.*

$(a)_\alpha$ *If $X \in [\theta]^{\lambda^+}$, there exists a canonical embedding $\pi_\alpha^{X,\theta} : G_\alpha^{M(X)} \to G_\alpha^{M(\theta)}$.*

$(b)_\alpha$ *Whenever X, $Y \in [\theta]^{\lambda^+}$ satisfy $X \subset Y$ and $|Y \smallsetminus X| = \lambda^+$, then there exists a canonical embedding $\pi_\alpha^{X,Y} : G_\alpha^{M(X)} \to G_\alpha^{M(Y)}$. Furthermore, the following diagram commutes.*

$$\begin{tikzcd} & G_\alpha^{M(\theta)} \\ G_\alpha^{M(X)} \arrow[ur, "\pi_\alpha^{X,\theta}"] \arrow[r, "\pi_\alpha^{X,Y}"] & G_\alpha^{M(Y)} \arrow[u, "\pi_\alpha^{Y,\theta}"'] \end{tikzcd}$$

$(c)_\alpha$ *For each* $g \in G_\alpha^{M(\theta)}$, *there exists* $X \in [\theta]^{\lambda^+}$ *and* $h \in G_\alpha^{M(X)}$ *such that* $\pi_\alpha^{X,\theta}(h) = g$.

$(d)_\alpha$ *Suppose that* $X = \bigcup_{i<\lambda^+} X_i$, *where* $\langle X_i \mid i < \lambda^+ \rangle \in M$ *is a smooth strictly increasing sequence of elements of* $[\theta]^{\lambda^+}$ *such that* $|X_{i+1} \smallsetminus X_i| = \lambda^+$ *for each* $i < \lambda^+$. *Then* $\pi_\alpha^{X,\theta}[G_\alpha^{M(X)}] = \bigcup_{i<\lambda^+} \pi_\alpha^{X_i,\theta}[G_\alpha^{M(X_i)}]$.

Before we prove Lemma 9.3.4, we shall show how to complete the proof of Theorem 9.3.1.

Proof of Theorem 9.3.1. Letting $X = \lambda^+$ in Lemma 9.3.4$(a)_\alpha$, we see that $M^{\mathbb{P}} = M(\theta)$ is a G-conservative extension of $M^{\mathbb{Q}} = M(\lambda^+)$. Hence by Lemma 9.2.3, we have that $\tau^{M(\lambda^+)}(G) \leq \tau^{M(\theta)}(G)$. To see that $\tau^{M(\theta)}(G) \leq \tau^{M(\lambda^+)}(G)$, suppose that $\beta < \tau^{M(\theta)}(G)$ and let $g \in G_{\beta+1}^{M(\theta)} \smallsetminus G_\beta^{M(\theta)}$. By Lemma 9.3.4$(c)_{\beta+1}$, there exists $Y \in [\theta]^{\lambda^+}$ and $h \in G_{\beta+1}^{M(Y)}$ such that $\pi_{\beta+1}^{Y,\theta}(h) = g$. Applying Lemma 9.2.3, we see that $h \in G_{\beta+1}^{M(Y)} \smallsetminus G_\beta^{M(Y)}$ and so $M(Y) \vDash \tau(G) > \beta$. By Lemma 9.3.3, $M(\lambda^+) \vDash \tau(G) > \beta$. Hence $\tau^{M(\theta)}(G) = \tau^{M(\lambda^+)}(G)$. Finally Theorem 6.4.4 implies that $M(\lambda^+) \vDash 2^\lambda = \lambda^+$ and so $\tau^{M(\lambda^+)}(G) < \lambda^{++}$. $\square$

Proof of Lemma 9.3.4. We shall prove that statements $(a)_\alpha$, $(b)_\alpha$, $(c)_\alpha$ and $(d)_\alpha$ hold by a simultaneous induction on α. The result is trivially true if $\alpha = 0$ and there are no difficulties at limit stages of the induction. So suppose that $\alpha = \beta + 1$ and that the result holds for β. For each $Z \in [\theta]^{\lambda^+}$, let $\Gamma^Z = \pi_\beta^{Z,\theta}[G_\beta^{M(Z)}]$. Notice that statement $(b)_\beta$ implies that if Y, $Z \in [\theta]^{\lambda^+}$ satisfy $Y \subset Z$ and $|Z \smallsetminus Y| = \lambda^+$, then $\Gamma^Y \leqslant \Gamma^Z$. It also follows easily from statement $(b)_\beta$ that if X, Y, $Z \in [\theta]^{\lambda^+}$ satisfy $X \subset Y \subset Z$ and $|Y \smallsetminus X| = |Z \smallsetminus Y| = \lambda^+$, then the following diagram commutes.

$$\begin{tikzcd}
 & G_\beta^{M(Z)} \\
G_\alpha^{M(X)} \arrow[ur, "\pi_\beta^{X,Z}"] \arrow[r, "\pi_\beta^{X,Y}"] & G_\alpha^{M(Y)} \arrow[u, "\pi_\beta^{Y,Z}"']
\end{tikzcd}$$

First we shall prove that statement $(a)_\alpha$ holds. Fix some $X \in [\theta]^{\lambda^+}$. As we noted in Remark 9.2.2, it is enough to show that whenever $g \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)})$, then there exists a corresponding automorphism $h \in \mathrm{Aut}^{M(\theta)}(G_\beta^{M(\theta)})$ such that $\pi_\beta^{X,\theta} \circ g \circ \left(\pi_\beta^{X,\theta}\right)^{-1} \subseteq h$. Fix some $g \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)})$; and, working inside M, express $X = \bigcup_{i<\lambda^+} X_i$ as the union of a smooth strictly increasing chain such that $|X_{i+1} \smallsetminus X_i| = \lambda^+$ for all $i < \lambda^+$.

Claim 9.3.5. *There exists an $i < \lambda^+$ such that $g[G_0] \leqslant \pi_\beta^{X_i,X}[G_\beta^{M(X_i)}]$ and $k = (\pi_\beta^{X_i,X})^{-1} \circ (g \restriction G_0) \in M(X_i)$.*

Proof of Claim 9.3.5. By statement $(d)_\beta$, we have that

$$\pi_\beta^{X,\theta}[G_\beta^{M(X)}] = \bigcup_{i<\lambda^+} \pi_\beta^{X_i,\theta}[G_\beta^{M(X_i)}].$$

So applying $(\pi_\beta^{X,\theta})^{-1}$ to both sides of this equality, we obtain that

$$G_\beta^{M(X)} = \bigcup_{i<\lambda^+} \pi_\beta^{X_i,X}[G_\beta^{M(X_i)}].$$

Since $|g[G_0]| = \lambda$, there exists an $j < \lambda^+$ such that $g[G_0] \leqslant \pi_\beta^{X_j,X}[G_\beta^{M(X_j)}]$. Since $k' = (\pi_\beta^{X_j,X})^{-1} \circ (g \restriction G_0) \in M(X)$ and $|k'| = \lambda$, there exists i such that $j < i < \lambda^+$ and $k' \in M(X_i)$. We claim that i satisfies our requirements. To see this, note that $\pi_\beta^{X_j,X_i} \in M(X_i)$ and that the following diagram commutes.

$$\begin{tikzcd}
& G_\beta^{M(X)} \\
G_\alpha^{M(X_j)} \arrow[ur, "\pi_\beta^{X_j,X}"] \arrow[r, "\pi_\beta^{X_j,X_i}"] & G_\alpha^{M(X_i)} \arrow[u, "\pi_\beta^{X_i,X}"']
\end{tikzcd}$$

It follows that

$$g[G_0] \leqslant \pi_\beta^{X_j,X}[G_\beta^{M(X_j)}] \leqslant \pi_\beta^{X_i,X}[G_\beta^{M(X_i)}]$$

and that

$$k = (\pi_\beta^{X_i,X})^{-1} \circ (g \restriction G_0) = \pi_\beta^{X_j,X_i} \circ k' \in M(X_i).$$

□

In particular, the following statement with parameters G, $k \in M(X_i)$ is true in $M(X)$.

$(9.3.5)_X$ *There exists an automorphism $g \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)})$ such that*

$$\pi_\beta^{X_i,X} \circ k \subseteq g.$$

By Lemma 9.3.2, statement $(9.3.5)_Y$ must be true for all $Y \in [\theta]^{\lambda^+}$ such that $X_i \subset Y$ and $|Y \smallsetminus X_i| = \lambda^+$; say, statement $(9.3.5)_Y$ is witnessed by the automorphism $g^Y \in \mathrm{Aut}^{M(Y)}(G_\beta^{M(Y)})$. We can now define a suitable automorphism $h \in \mathrm{Aut}^{M(\theta)}(G_\beta^{M(\theta)})$ as follows. Let $a \in G_\beta^{M(\theta)}$ be any element. By statements $(c)_\beta$ and $(b)_\beta$, there exists a set $Y \in [\theta]^{\lambda^+}$ such that $X_i \subset Y$, $|Y \smallsetminus X_i| = \lambda^+$ and $a \in \Gamma^Y$; and so we can define

$$h(a) = \pi_\beta^{Y,\theta} \circ g^Y \circ (\pi_\beta^{Y,\theta})^{-1}(a).$$

Of course, we must check that h is well-defined. To see this, it is enough to consider the case when $X_i \subset Y_1 \subset Y_2$, $|Y_1 \smallsetminus X_i| = |Y_2 \smallsetminus Y_1| = \lambda^+$ and $a \in \Gamma^{Y_1} \leqslant \Gamma^{Y_2}$. Notice that

$$(\pi_\beta^{Y_1,Y_2})^{-1} \circ (g^{Y_2} \restriction G_0) = (\pi_\beta^{Y_1,Y_2})^{-1} \circ \pi_\beta^{X_i,Y_2} \circ k = \pi_\beta^{X_i,Y_1} \circ k \subseteq g^{Y_1}.$$

So by Lemma 9.2.4, we have that

$$(\pi_\beta^{Y_1,Y_2}) \circ g^{Y_1} \circ (\pi_\beta^{Y_1,Y_2})^{-1} \subseteq g^{Y_2}.$$

It follows that h is well-defined. Also notice that

$$(\pi_\beta^{X,\theta})^{-1} \circ (h \restriction G_0) \subseteq g$$

and so by Lemma 9.2.4,

$$(\pi_\beta^{X,\theta}) \circ g \circ (\pi_\beta^{X,\theta})^{-1} \subseteq h,$$

as required. This completes the proof that statement $(a)_\alpha$ holds.

Next we shall prove that statement $(b)_\alpha$ holds. Fix some $X \in [\theta]^{\lambda^+}$. We shall first show that if $Y \in [\theta]^{\lambda^+}$ is such that $X \subset Y$ and $|Y \smallsetminus X| = \lambda^+$, then there exists a canonical embedding $\pi_\alpha^{X,Y} : G_\alpha^{M(X)} \to G_\alpha^{M(Y)}$. By Lemma 9.3.2, it is enough to find a single set $Y \in [\theta]^{\lambda^+}$ with these properties. We shall use the following observation to find such a set Y.

Claim 9.3.6. *Suppose that $S \in [\theta]^{\lambda^+}$ and that $X \subseteq S$. Then there exists $T \in [\theta]^{\lambda^+}$ such that $S \subseteq T$ and*

$$1_{\mathbb{P}} \Vdash_{\mathbb{P}} \pi_\alpha^{X,\theta}(g)[\Gamma^S] \leqslant \Gamma^T \textit{ for all } g \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)}).$$

Proof of Claim 9.3.6. Suppose that $S \in [\theta]^{\lambda^+}$ and that $X \subseteq S$. Initially we shall work within $M(\theta)$. By statement $(c)_\beta$, for each $g \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)})$ and $a \in \Gamma^S$, there exists a set $Y_{g,a} \in [\theta]^{\lambda^+}$ such that $\pi_\alpha^{X,\theta}(g)(a) \in \Gamma^{Y_{g,a}}$. Let

$$Y = S \cup \bigcup \{Y_{g,a} \mid g \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)}) \text{ and } a \in \Gamma^S\}.$$

Recall that G is an ascendent subgroup of both $\mathrm{Aut}^{M(X)}(G_\beta^{M(X)})$ and Γ^S. Consequently, since $M(X)$ and $M(S)$ both satisfy $2^\lambda = \lambda^+$, Theorem 3.3.1 implies that $|\mathrm{Aut}^{M(X)}(G_\beta^{M(X)})| \leq \lambda^+$ and $|\Gamma^S| \leq \lambda^+$. Thus $|Y| = \lambda^+$.

Now let $\widetilde{Y}$ be a $\mathrm{Fn}(\theta, 2, \kappa)$-name for Y. Since $\mathrm{Fn}(\theta, 2, \kappa)$ has the κ^+-*c.c.*, it follows that there exists a set $T \in [\theta]^{\lambda^+}$ such that

$$1_{\mathbb{P}} \Vdash_{\mathbb{P}} \widetilde{Y} \subseteq T.$$

Clearly T satisfies our requirements. $\square$

Thus, working within M, we can inductively define a smooth strictly increasing sequence $\langle X_i \mid i < \lambda^+\rangle$ such that $X_0 = X$, $|X_{i+1} \smallsetminus X_i| = \lambda^+$ and

$$1_{\mathbb{P}} \Vdash_{\mathbb{P}} \pi_\alpha^{X,\theta}(g)[\Gamma^{X_i}] \leqslant \Gamma^{X_{i+1}} \text{ for all } g \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)}).$$

Let $Y = \bigcup_{i<\lambda^+} X_i$. By statement $(d)_\beta$, we have that $\Gamma^Y = \bigcup_{i<\lambda^+} \Gamma^{X_i}$ and hence $\pi_\alpha^{X,\theta}(g)[\Gamma^Y] = \Gamma^Y$ for all $g \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)})$. For each $g \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)})$, define the automorphism $\pi_\alpha^{X,Y}(g) \in \mathrm{Aut}^{M(\theta)}(G_\beta^{M(Y)})$ by

$$\pi_\alpha^{X,Y}(g) = (\pi_\beta^{Y,\theta})^{-1} \circ \pi_\alpha^{X,\theta}(g) \circ \pi_\beta^{Y,\theta}.$$

Notice that

$$\pi_\alpha^{X,Y}(g) \restriction \pi_\beta^{X,Y}[G_\beta^{M(X)}] = (\pi_\beta^{X,Y})^{-1} \circ g \circ \pi_\beta^{X,Y}.$$

Since $\pi_\beta^{X,Y}$, $g \in M(Y)$, it follows that $\pi_\alpha^{X,Y}(g) \restriction G_0 \in M(Y)$. By Lemma 6.7.2, $\pi_\alpha^{X,Y}(g) \in M(Y)$. Thus $\pi_\alpha^{X,Y}$ is a canonical embedding of $\mathrm{Aut}^{M(X)}(G_\beta^{M(X)})$ into $\mathrm{Aut}^{M(Y)}(G_\beta^{M(Y)})$. It is now easily checked that the following diagram commutes.

$$\begin{tikzcd}
 & \mathrm{Aut}^{M(\theta)}(G_\beta^{M(\theta)}) \\
\mathrm{Aut}^{M(X)}(G_\beta^{M(X)}) \arrow[ur, "\pi_\alpha^{X,\theta}"] \arrow[r, "\pi_\alpha^{X,Y}"] & \mathrm{Aut}^{M(Y)}(G_\beta^{M(Y)}) \arrow[u, "\pi_\alpha^{Y,\theta}"']
\end{tikzcd}$$

Hence X,Y satisfy the conclusion of statement $(b)_\alpha$.

Next we shall prove that statement $(c)_\alpha$ holds. Let $g \in \mathrm{Aut}^{M(\theta)}(G_\beta^{M(\theta)})$. Then there exists $X_0 \in [\theta]^{\lambda^+}$ such that $g[G_0] \leqslant \Gamma^{X_0}$. Let $k = (\pi_\beta^{X_0,\theta})^{-1} \circ (g \restriction G_0)$. Then there exists $X_1 \supset X_0$ such that $|X_1 \smallsetminus X_0| = \lambda^+$ and $k \in M(X_1)$. Thus

$$(\pi_\beta^{X_1,\theta})^{-1} \circ (g \restriction G_0) = \pi_\beta^{X_0,X_1} \circ k \in M(X_1).$$

Arguing as in the previous paragraph, we can define a smooth strictly increasing sequence $\langle X_i \mid i < \lambda^+ \rangle \in M$ of elements of $[\theta]^{\lambda^+}$ such that if $Y = \bigcup_{i<\lambda^+} X_i$, then $g[\Gamma^Y] = \Gamma^Y$. Thus $h = (\pi_\beta^{Y,\theta})^{-1} \circ g \circ \pi_\beta^{Y,\theta} \in \mathrm{Aut}^{M(\theta)}(G_\beta^{M(Y)})$ and it suffices to show that $h \in M(Y)$. To see this, note that

$$h \restriction G_0 = \pi_\beta^{X_1,Y} \circ \Big((\pi_\beta^{X_1,\theta})^{-1} \circ (g \restriction G_0)\Big),$$

and so $h \restriction G_0 \in M(Y)$. By Lemma 6.7.2, $h \in M(Y)$.

Finally we shall prove statement $(d)_\alpha$. So suppose that $X = \bigcup_{i<\lambda^+} X_i$, where $\langle X_i \mid i < \lambda^+ \rangle \in M$ is a smooth strictly increasing sequence of elements of $[\theta]^{\lambda^+}$ such that $|X_{i+1} \smallsetminus X_i| = \lambda^+$ for each $i < \lambda^+$. It is clear that

$$\bigcup_{i<\lambda^+} \pi_\alpha^{X_i,\theta}[G_\alpha^{M(X_i)}] \leqslant \pi_\alpha^{X,\theta}[G_\alpha^{M(X)}].$$

Conversely suppose that $g \in \pi_\alpha^{X,\theta}[G_\alpha^{M(X)}]$ and let

$$h = (\pi_\alpha^{X,\theta})^{-1}(g) \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)}).$$

Arguing as above, there exists $i < \lambda^+$ such that $h[G_0] \leqslant \pi_\beta^{X_i,X}[G_\beta^{M(X_i)}]$ and $k = (\pi^{X_i,X})^{-1} \circ (h \restriction G_0) \in M(X_i)$. Notice that $|X \smallsetminus X_i| = |X_{i+1} \smallsetminus X_i| = \lambda^+$ and that $h \in \mathrm{Aut}^{M(X)}(G_\beta^{M(X)})$ satisfies $\pi_\beta^{X_i,X} \circ k \subseteq h$. So Lemma 9.3.2 implies that there exists $f \in \mathrm{Aut}^{M(X_{i+1})}(G_\beta^{M(X_{i+1})})$ such that $\pi_\beta^{X_i,X_{i+1}} \circ k \subseteq f$. Let $h' = \pi_\alpha^{X_{i+1},X}(f)$. Then $h' \restriction G_0 = h \restriction G_0$ and so $h' = h$. It follows that $g \in \pi_\alpha^{X_{i+1},\theta}[G_\alpha^{M(X_{i+1})}]$. This completes the proof of Lemma 9.3.4. □

We now easily obtain the following result.

THEOREM 9.3.7. *Let M be a c.t.m. and let $\kappa \leq \lambda < \lambda$ be regular cardinals which satisfy the following conditions.*

(a) $\kappa^{<\kappa} = \kappa$.
(b) $2^\lambda = \lambda^+$.
(c) $\theta \geq \lambda^{++}$ *and* $\theta^\lambda = \theta$.

Then $\mathbb{P} = \mathrm{Fn}(\theta, 2, \kappa)$ *is* κ*-closed and has the* κ^+*-c.c.. Hence* $\mathbb{P}$ *preserves cofinalities and cardinals. Furthermore, if* H *is a* $\mathbb{P}$*-generic filter over* M*, then the following statements are true in* $M[H]$*.*

(1) $2^\lambda = \theta$.

(2) *If* G *is a centreless group of cardinality* λ*, then* $\tau(G) < \lambda^{++}$*.*

Proof. Except for statement 9.3.7(2), the conclusion follows from Theorem 6.4.4. Let $G \in M[H]$ be a centreless group of cardinality λ. By Lemma 9.1.6, there exists a subset $X \subset \theta$ of cardinality λ such that $G \in M_1 = M[H \cap \mathrm{Fn}(X, 2, \kappa)]$. Let $Y = \theta \setminus X$. By Theorem 9.1.5, $H_Y = H \cap \mathrm{Fn}(Y, 2, \kappa)$ is a $\mathrm{Fn}(Y, 2, \kappa)$-generic filter over M_1 and $M[H] = M_1[H_Y]$. Working inside M_1, we have that $\mathrm{Fn}(Y, 2, \kappa)^M = \mathrm{Fn}(Y, 2, \kappa)^{M_1} \simeq \mathbb{P}^{M_1}$ and the cardinals κ, λ and θ continue to satisfy our original hypotheses. Hence Theorem 9.3.1 implies that $\tau(G) < \lambda^{++}$ in $M_1[H_1] = M[H]$. $\square$

Finally we will use a suitable reverse Easton iteration to complete the proof of Theorem 9.1.1. Let $V \vDash GCH$ and let $\{\theta_\beta \mid \beta \in On\}$ be the increasing enumeration of the class of limit cardinals. For each $\beta \in On$, define

$$\kappa_\beta = \begin{cases} \theta_\beta, & \text{if } \theta_\beta \text{ is regular;} \\ \theta_\beta^+, & \text{if } \theta_\beta \text{ is singular.} \end{cases}$$

Thus if $\beta > 0$, then $\kappa_\beta = \theta_\beta$ iff β is an inaccessible cardinal. We shall define a sequence $\langle \mathbb{P}_\beta \mid \beta \in On \rangle$ of forcing notions by induction on β.

Case 1. If $\beta = 0$, then $\mathbb{P}_0 = \{\emptyset\}$ is the trivial notion of forcing.

Case 2. If β is a limit ordinal which is not inaccessible, then $\mathbb{P}_\beta$ is the inverse limit of $\langle \mathbb{P}_\gamma \mid \gamma < \beta \rangle$.

Case 3. If β is inaccessible, then $\mathbb{P}_\beta$ is the direct limit of $\langle \mathbb{P}_\gamma \mid \gamma < \beta \rangle$.

Case 4. Finally if $\beta = \gamma + 1$ is a successor ordinal, then $\mathbb{P}_{\gamma+1} = \mathbb{P}_\gamma * \widetilde{\mathbb{Q}}_\gamma$, where $\widetilde{\mathbb{Q}}_\gamma$ is a $\mathbb{P}_\gamma$-name for the notion of forcing $\mathrm{Fn}(\theta_{\gamma+1}^+, 2, \kappa_\gamma) \in V^{\mathbb{P}_\gamma}$.

Let $\mathbb{P}_\infty$ be the direct limit of $\langle \mathbb{P}_\beta \mid \beta \in On \rangle$; and for each $\beta \in On$, let $\tilde{\mathbb{P}}_{\beta\infty}$ be the canonically chosen $\mathbb{P}_\beta$-name for a proper class notion of forcing such that $\mathbb{P}_\infty$ is isomorphic to a dense sub-order of $\mathbb{P}_\beta * \tilde{\mathbb{P}}_{\beta\infty}$. Let F be a $\mathbb{P}_\infty$-generic filter over V and let $V = V[F]$ be the corresponding generic extension. For each $\beta \in On$, let $F_\beta = F \cap \mathbb{P}_\beta$ and let $\mathbb{P}_{\beta\infty} = (\tilde{\mathbb{P}}_{\beta\infty})_{F_\beta}$.

Lemma 9.3.8. (a) $\mathbb{P}_\infty$ *preserves cofinalities and cardinals.*

(b) *For each ordinal β, the following statements are true in $V[F_\beta]$:*

(i) $2^\mu = \mu^+$ *for all cardinals* $\mu \geq \theta_\beta$.

(ii) $\kappa_\beta^{<\kappa_\beta} = \kappa_\beta$.

(iii) $\mathbb{P}_{\beta\infty}$ *is* κ_β*-closed.*

(c) $V[F]$ *is a model of* ZFC.

(d) *Let λ be regular and let α be the least ordinal such that $\lambda < \theta_\alpha$. Then* $V[F] \vDash 2^\lambda = \theta_\alpha^+ > \lambda^{++}$.

PROOF. This follows easily from Menas [**32**]. $\square$

Thus to complete the proof of Theorem 9.1.1, it is enough to show that if λ is a regular cardinal and $G \in V[F]$ is a centreless group of cardinality λ, then $V[F] \vDash \tau(G) < \lambda^{++}$. Clearly we can suppose that G has underlying set λ. Let β be the least ordinal such that $\lambda < \theta_\beta$. Clearly β is a successor ordinal, say $\beta = \gamma + 1$, and $\theta_\gamma \leq \kappa_\gamma \leq \lambda < \theta_{\gamma+1} < \kappa_{\gamma+1}$. Since $V[F_\beta] \vDash \mathbb{P}_{\beta\infty}$ is κ_β-closed, it follows that $G \in V[F_\beta]$, and Theorem 6.7.1 implies that $\tau^{V[F]}(G) = \tau^{V[F_\beta]}(G)$. Thus it suffices to prove that $V[F_\beta] \vDash \tau(G) < \lambda^{++}$. By Lemma 9.3.8(b), the following statements are true in $V[F_\gamma]$:

(a) $\kappa_\gamma^{<\kappa_\gamma} = \kappa_\gamma$.

(b) $2^\lambda = \lambda^+$.

(c) $\lambda^{++} < \theta_{\gamma+1}^+$ and $(\theta_{\gamma+1}^+)^\lambda = \theta_{\gamma+1}^+$.

Since $(\widetilde{\mathbb{Q}}_\gamma)_{F_\gamma} = \mathrm{Fn}(\theta_{\gamma+1}^+, 2, \kappa_\gamma)$, Theorem 9.3.7 implies that $V[F_{\gamma+1}] \vDash \tau(G) < \lambda^{++}$.

QUESTION 9.3.9. Is it consistent with ZFC that $\tau_\lambda < 2^\lambda$ for all infinite cardinals λ?

A positive answer to Question 9.3.9 would necessarily involve the use of large cardinals. To see this, recall that Corollary 4.1.14 says that $\tau_\lambda \geq \lambda^+$ for every infinite cardinal λ. Hence if $\tau_\lambda < 2^\lambda$ for every infinite cardinal λ, then $2^\lambda > \lambda^+$ for every infinite cardinal λ. In particular, $2^\lambda > \lambda^+$ when λ is a singular strong limit cardinal; and by Gitik [**11**], the consistency strength of this latter statement is strictly stronger than the consistency strength of a measurable cardinal. On the other hand, Woodin has shown that if there exists a supercompact cardinal, then there exists a model of ZFC in which $2^\lambda = \lambda^{++}$ for every infinite cardinal λ. (See Foreman-Woodin [**9**].) Thus it seems reasonable to conjecture that a positive answer to Question 9.3.9 can be obtained under the assumption that a supercompact cardinal exists.

9.4. Notes

The material in this chapter first appeared in Thomas [**50**].

Index

Bibliography

[1] E. Artin, *Algebraic numbers and algebraic functions*, Gordon and Breach, New York, 1967.

[2] J. E. Baumgartner, *Iterated forcing*, in Surveys in Set Theory (ed. A. R. D. Mathias), London Math. Soc. Lecture Notes Series **87**, Cambridge University Press, 1983, pp. 1–59.

[3] N. G. de Bruijn, *Embedding theorems for infinite groups*, Indag. Math. **19** (1957), 560–569.

[4] F. Clare, *Embedding theorems for infinite symmetric groups*, Indag. Math. **37** (1975), 155–163.

[5] R. M. Dudley, *Continuity of homomorphisms*, Duke Math. J. **28** (1961), 587–594.

[6] P. C. Eklof and A. H. Mekler, *Almost Free Modules: Set-Theoretic Methods*, North-Holland, Amsterdam, 1990.

[7] V. Faber, *Large abelian subgroups of some infinite groups II*, Rocky Mountain Journal of Mathematics **8** (1978), 481–490.

[8] U. Felgner and F. Haug, *The homomorphic images of infinite symmetric groups*, Forum Math. **5** (1993), 505–520.

[9] M. Foreman and W. H. Woodin, *The generalised continuum hypothesis can fail everywhere*, Annals of mathematics **133** (1991), 1–36.

[10] E. Fried and J. Kollár, *Automorphism groups of fields*, in Universal Algebra (E. T. Schmidt et al., eds.), Colloq. Math. Soc. János Bolyai, vol **24**, 1981, pp. 293–304.

[11] M. Gitik, *The strength of the failure of the singular cardinal hypothesis*, Ann. Pure Appl. Logic **51** (1991), 215–240.

[12] M. Gitik and M. Magidor, *The singular cardinals hypothesis revisited*, in Set Theory of the Continuum (H. Judah et al., eds.), Berkeley, CA, 1989, pp. 243–279.

[13] P. Hall, *On non-strictly simple groups*, Proc. Cambridge Philos. Soc. **59** (1963), 531–553.

[14] J. D. Hamkins, *Every group has a terminating transfinite automorphism tower*, Proc. Amer. Math. Soc. **126** (1998), 3223–3226.

[15] J. D. Hamkins and S. Thomas, *Changing the heights of automorphism towers*, Annals Pure Appl. Logic. **102** (2000), 139–157.

[16] W. A. Hodges, *Model Theory*, Cambridge University Press, Cambridge, 1993.

[17] J. A. Hulse, *Automorphism towers of polycyclic groups*, J. Algebra **16** (1970), 347–398.

[18] T. Jech, *Automorphisms of ω_1-trees*, Trans. Amer. Math. Soc. **173** (1972), 57–70.

[19] T. Jech, *Set Theory*, Academic Press, New York, 1978.

[20] T. Jech, *Multiple Forcing*, Cambridge University Press, Cambridge, 1986.

[21] T. Jech, *Singular cardinals and the PCF theory*, Bull. Symbolic Logic **1** (1995), 408–428.

[22] W. Just, S. Shelah and S. Thomas, *The automorphism tower problem revisited*, Advances in Mathematics **148** (1999), 243–265.

[23] G. Karpilovsky, *Field Theory*, Marcel Dekker, New York, 1988.

[24] D. W. Kueker, *Definability, automorphisms and infinitary languages*, in The Syntax and Semantics of Infinitary Languages (ed. J. Barwise), Lecture Notes in Math. **72**, Springer, Berlin, 1968, pp. 152–165.

[25] K. Kunen, *Combinatorics*, in Handbook of Mathematical Logic (ed. J. Barwise), North-Holland, Amsterdam, 1977, pp. 371–401.

[26] K. Kunen, *Set Theory*, North-Holland, Amsterdam, 1980.

[27] S. Lang, *Algebra*, Addison-Wesley, Menlo Park, 1984.

[28] H. D. Macpherson and P. M. Neumann, *Subgroups of infinite symmetric groups*, J. London Math. Soc. (2) **42** (1990), 64–84.

[29] M. Magidor, *On the singular cardinals problem I*, Israel J. Math **28** (1977), 1–31.

[30] R. McKenzie, *A note on subgroups of infinite symmetric groups*, Indag. Math. **33** (1971), 53–58.

[31] D. H. McLain, *Remarks on the upper central series of a group*, Proc. Glasgow Math. Assoc. **3** (1956), 38–44.

[32] T. K. Menas, *Consistency results concerning supercompactness*, Trans. Amer. Math. Soc. **223** (1976), 61–91.

[33] B. H. Neumann, *Some remarks on infinite groups*, J. London Math. Soc. **12** (1937), 120–127.

[34] P. M. Neumann, *On the structure of standard wreath products of groups*, Math. Z. **84** (1964), 343–373.

[35] M. R. Pettet, *A note on the automorphism tower for finite groups*, Proc. Amer. Math. Soc. **89** (1983), 182–183.

[36] B.I. Plotkin, *Groups of automorphisms of algebraic systems*, Wolters-Noordhoff Publishing, Groningen, 1972.

[37] E. B. Rabinovič, *Embedding theorems and De Bruijn's problem for bounded symmetric groups* (Russian), Dokl. Akad. Nauk BSSR **21** (1977), 784-787, 859.

[38] A. Rae and J. E. Roseblade, *Automorphism towers of extremal groups*, Math. Z. **117** (1970), 70–75.

[39] D. J. S. Robinson, *A course in the theory of groups*, Springer-Verlag, New York, 1995.

[40] E. Schenkman, *The tower theorem for finite groups*, Proc. Amer. Math. Soc. **22** (1969), 458–459.

[41] O. Schreier and B. L. van der Waerden, *Die Automorphismen der projektiven Gruppen*, Abh. Math. Sem. Univ. Hamburg **6** (1928), 303–322.

[42] W. R. Scott, *Group Theory*, Prentice-Hall, Englewood Cliffs, N. J., 1964.

[43] S. Shelah, *First order theory of permutation groups*, Israel J. Math. **14** (1973), 149–162.

[44] S. Shelah, *Cardinal Arithmetic*, Oxford Logic Guides **29**, Oxford, 1994.

[45] S. Shelah, *A countable structure does not have a free uncountable automorphism group*, Bull. London Math. Soc. **1** (2003), 1–7.

[46] J. H. Silver, *On the singular cardinals problem*, Proc. Int. Congr. Math. Vancouver 1974, pp. 265–268.

[47] S. Solecki, *Polish group topologies*, in Sets and Proofs (B. Cooper and J. Truss, eds.), London Math. Soc. Lecture Notes Series **258**, Cambridge University Press, 1999, pp. 339–364.

[48] M. Suzuki, *Group Theory I*, Grundlehren der mathematischen Wissenschaften **247**, Springer-Verlag, 1982.

[49] S. Thomas, *The automorphism tower problem*, Proc. Amer. Math. Soc. **95** (1985), 166–168.

[50] S. Thomas, *The automorphism tower problem II*, Israel J. Math. **103** (1998), 93–109.

[51] J. K. Truss, *Infinite permutation groups II: subgroups of small index*, J. Algebra **120** (1989), 494–515.

[52] H. Wielandt, *Eine Verallgemeinerung der invarienten Untergruppen*, Math. Z **45** (1939), 209–244.

[53] H. Zassenhaus, *Lehrbuch der Gruppentheorie*, B. G. Teubner, Liepzig, 1937.