Computing automorphisms of finite soluble groups

Michael J. Smith

August 1994

A thesis submitted for the degree of Doctor of Philosophy
of the Australian National University

Declaration

The work in this thesis is my own unless otherwise stated.

Michael J. Smith

Acknowledgements

My supervisor, Dr M.F. Newman, has been a constant and immeasurable source of help and guidance throughout my last three and a half years of study. I am very grateful for all that he has taught me during this time. I thank him for his never-ending enthusiasm, encouragement and interest in my work.

Along with Dr M.F. Newman, my advisors Dr E.A. O'Brien and Dr L.G. Kovács have spent many hours discussing ideas arising from my work and many enlightening hours discussing topics outside it. I thank them all for the time they devoted to these discussions and for their invaluable help in reading and refining parts of this thesis.

I thank Werner Nickel and Alice Niemeyer for their friendship, advice and encouragement. I also thank Alice Niemeyer for her assistance with some German translation, and for the many mathematical discussions we have had over the last year that have been such a boost to my confidence. I thank Dr C.R. Leedham-Green and Werner Nickel for some of the initial ideas for this work.

Philippa Harvey and Annette Hughes have helped greatly in administrative matters. I thank them both for their friendship and concern. The constant repartee from Philippa Harvey often supplied much appreciated light relief.

I thank my family and friends for their concern and understanding. I especially thank my wife Maree for her endless love and support. I greatly appreciate the large amount of time she spent enthusiastically helping me with the proof-reading of this thesis.

I thank the Centre for Mathematics and its Applications and the former Mathematics Research Section at the Australian National University for providing a stimulating environment for my studies. The constant stream of visitors from the international mathematical community is something I shall miss.

Finally, I would like to thank the Defence Science and Technology Organisation for providing me with the opportunity to study at the Australian National University. I also thank my Defence mentors Prof. K.J. Horadam and Dr I. Doherty for their support and interest in my work.

Abstract

There is a large collection of effective algorithms for computing information about finite soluble groups. The success in computation with these groups is primarily due to a computationally convenient representation of them by means of (special forms of) power conjugate presentations. A notable omission from this collection of algorithms is an effective algorithm for computing the automorphism group of a finite soluble group. An algorithm designed for finite groups in general provides only a partial answer to this deficiency.

In this thesis an effective algorithm for computing the automorphism group of a finite soluble group is described. An implementation of this algorithm has proved to be a substantial improvement over existing techniques available for finite soluble groups.

Table of contents

Chapter 1

Introduction

The exploitation of symmetry is a standard technique employed in many areas of mathematics. In algebraic and combinatorial contexts the symmetries of an object are usually called automorphisms. The group of all automorphisms of an object is important in many contexts. It may, for example, be useful as an invariant for distinguishing non-isomorphic objects. In other applications it may be used to improve the performance of algorithms that compute information about the object.

For example, the program of McKay (1978) for computing a canonical labelling of a graph uses the automorphism group of the graph to reduce substantially the time required for the computation. Canonical labellings provide a solution to the isomorphism problem for graphs. The algorithm is based on a backtrack search across an extremely large search tree associated with the graph, and the automorphisms are used to remove large sections of this tree from consideration. This extends the range of application of the canonical labelling program and therefore provides a solution of the isomorphism problem for much larger graphs than would otherwise be possible.

A recent application of automorphisms in computational group theory is to the study of Burnside groups. Investigation into the Burnside groups has proved to be fertile ground for the development of many successful programs for computing with finite p-groups. Proving that a finite presentation actually defines a Burnside group involves checking that an exponent law holds in the group. This in principle involves checking the order of every element of the group. Commutator identities and other considerations reduce the amount of work somewhat. Using the action of automorphisms of the group Newman & O'Brien (in preparation) have dramatically reduced the number of elements that must be checked to verify the exponent law holds. When a violation of the required exponent condition is discovered the presentation can be adjusted by adding relations that remove the violation, resulting in a presentation that is "closer" to defining the Burnside group. This adjustment process allows the algorithm to be used not only for verification purposes, but also as part of an algorithm for computing a presentation for a Burnside group which is substantially faster than previous methods.

The p-group generation algorithm, as described by Newman (1977), computes a list of representatives for the isomorphism types of p-groups of a given order. It has been successfully used to compute lists of all the groups of order 128 (James *et al.* 1990) and 256 (O'Brien 1991). Automorphism groups are of fundamental importance in this algorithm, since the calculation relies on identifying the isomorphism classes of certain extensions of a p-group with orbits of subgroups of a covering group under the action of automorphisms of the p-group. By choosing a canonical representative from each of these orbits, a related algorithm computes a canonical presentation for a p-group (O'Brien 1993), thereby solving the isomorphism problem for p-groups.

Part of the success in computation with p-groups is due to the use of a special form of finite presentation called a power conjugate presentation. For a group given by such a presentation the answers to many fundamental problems about the group become readily computable. In particular an algorithm called collection produces unique normal forms for elements of the group, thereby solving the word problem for the group.

An algorithm for generating finite soluble groups is described by Niemeyer (1993). In many ways it is a generalisation to finite soluble groups of the p-group generation algorithm. The automorphisms of a finite soluble group are as fundamental to the operation of this algorithm as the automorphisms of a p-group are to the p-group generation algorithm. As before, the orbits of subgroups of a covering group under the action of automorphisms determine isomorphism classes of certain extensions of a group. A number of complications arise when moving from the p-group context to that of soluble groups. In particular, the size of the covering group of a finite soluble group is generally much larger than the covering group of a comparable p-group, and problems arise in determining how the automorphisms of a group act on its covering group. Currently there is no implementation of this algorithm.

The computation of the automorphism group of a finite group is a difficult task. The earliest successful work on this problem was carried out by Felsch & Neubüser (1968, 1970). Their algorithm used information about the lattice of subgroups to choose a particular generating set for the group as well as a list of maps defined on this generating set. An exhaustive search of this list yielded all the automorphisms of the group. Gerhards & Altmann (1970) described an algorithm for computing the automorphism group of a finite soluble group whose order involves at least two distinct primes. It also relied on knowing the lattice of subgroups of the group. No implementation of their algorithm was developed.

An implementation of a general purpose algorithm for computing automorphism groups of finite groups is available as a standard function in CAYLEY (Cannon 1984). A description of this algorithm can be found in Robertz (1976). The algorithm only requires knowledge of the conjugacy classes of the group, and it contructs the automorphism group as a permutation group acting on a set related to the conjugacy classes. The performance of this algorithm depends critically on the structure of the conjugacy classes of the group.

The only other widely available implementation of an algorithm for computing automorphism groups of finite groups is an implementation of an algorithm for computing the automorphism group of a finite p-group (O'Brien 1994). This is closely related to the standard presentation algorithm mentioned earlier, since the computation of a canonical presentation for a p-group involves the computation of a large amount of its automorphism group. The implementation of this algorithm performs extremely well and it is widely available as a standard package in both GAP (Schönert *et al.* 1993) and MAGMA (Bosma & Cannon 1993).

In this thesis we describe an algorithm for computing the automorphism group of a finite soluble group. Since the existence of such an algorithm is not in question, the aim here is to describe an effective algorithm. The criterion for effectiveness is that it leads to an implementation that works well in practice. The algorithm presented here computes a generating set for the automorphism group of a finite soluble group that is given by a special form of power conjugate presentation. These presentations, which were developed by Leedham-Green, are called *special power conjugate presentations*. An implementation of an algorithm for computing such presentations has recently become available in GAP (see Eick 1993).

The p-group generation algorithm and the finite soluble group generation algorithm both produce descriptions of the automorphism group of each group that they generate. Both of these algorithms rely on computations that take place in covering groups. The algorithm described here performs all computations in quotients of the given finite soluble group, avoiding covering groups entirely. This is an important distinction, since in the soluble context the covering groups can be extremely large relative to the original group. It is likely that avoiding covering group computations entirely will allow the calculation of automorphism groups of much larger soluble groups than would be possible otherwise. A prerequisite of group generation is a description of the automorphism group of the starting group. Where the covering groups of quotients of a finite soluble group are too large to be handled in soluble group generation, this algorithm could be employed to compute the automorphism group of the starting group, allowing soluble group generation to proceed from there.

The algorithm described in this thesis has a prototype implementation in GAP. Even though this prototype does not incorporate all the ideas discussed in the thesis, it already performs substantially better than the current alternative (see Chapter 6 for a comparison). This prototype is available by contacting the author via electronic mail at the address `Michael.Smith@maths.anu.edu.au`. A full implementation of the algorithm is forthcoming. By incorporating all the features of the algorithm described here, the newer implementation should provide a substantial improvement in performance and range of application.

Chapter 2 sets up the notation and basic results that are required throughout the remainder of the thesis.

In Chapter 3 a more general problem than that required for computing automorphism groups of finite soluble groups is defined. This problem involves computing a generating set for the automorphism group of an extension of an elementary abelian p-group by a finitely presented group. The solution of this problem divides naturally into a number of separate computations. The simplest of these is shown to correspond to computing a basis for a vector space of functions from the finitely presented group into the elementary abelian group. This part of the computation is described in the remainder of the chapter.

The remaining parts of the computation are more easily performed in the case of a split extension. Chapter 4 describes the solution in this case. This calculation involves deciding isomorphism of modules for the finitely presented group, and much of the chapter is devoted to an algorithm for computing information about modules. An important element of the problem here is to compute a generating set for the group of module automorphisms. This is equivalent to finding a generating set for the centraliser of a matrix group in the general linear group, and an effective solution to this problem is covered at the end of the chapter.

In Chapter 5 we consider the general or nonsplit case. This relies on much of the information computed about the corresponding split extension, and involves some additional computations. The specialisation of this step to the types of extension required for soluble groups is also described.

Chapter 6 describes the algorithm for computing the automorphism group of a finite soluble group. We start the chapter by defining special power conjugate presentations, following the description given in Eick (1993). These presentations exhibit precisely the structure required to assemble the algorithm for finite soluble groups from the algorithms described in earlier chapters. We then show how some of the information exhibited by such a presentation can be used both to reduce the amount of work required to compute the automorphism group, and also to compute a generating set that exhibits more structural information for the automorphism group.

Presented in Chapter 7 is a hand calculation following the steps of the automorphism group algorithm. This calculation plays two roles. The first is as an extended example of the steps involved in computing the automorphism group of a split extension. More importantly, it shows that the algorithm may be applied to compute the automorphism groups of an infinite number of groups in a single calculation. An infinite family of groups is constructed, indexed by the set of all odd primes. The automorphism group of an arbitrary member of the family is computed. The result of this calculation is a single parameterised generating set for all of the automorphism groups as well as a function giving the order of each automorphism group. In principle it should be possible to incorporate algorithms for working with variable parameters into an implementation of the automorphism group algorithm, or indeed of other group theoretic algorithms, and have similar computations performed automatically.

Chapter 2

Preliminaries

The notation used in the thesis and some basic results are presented in this chapter. Robinson (1982) should be consulted for standard group theoretic notation that is not defined here.

Let p be a prime, k a positive integer, F a field, and R a ring. We adopt the following basic notation:

1) $\mathbb{F}_p$ is a finite field with p elements.
2) $\mathbb{M}(k, R)$ is the full matrix ring of dimension k over R.
3) $\mathrm{GL}(k, F)$ is the general linear group of dimension k over F.
4) $\mathrm{GL}(k, p)$ is $\mathrm{GL}(k, \mathbb{F}_p)$.
5) $F^\times$ is the multiplicative group of non-zero elements of F.

2.1 Presentations

Let $X = \{x_1, \ldots, x_n\}$ be a set and denote by X^- the set of formal inverses $\{x_1^{-1}, \ldots, x_n^{-1}\}$. Let X^* be the free monoid on $X^\pm = X \cup X^-$. An element w of X^* is called a *word over* X, and where necessary we shall write w as either $w(x_1, \ldots, x_n)$ or $w(X)$ to indicate that $w \in X^*$. For e a positive integer we define w^e to be the concatenation of e copies of w. If $w = x_{i_1}^{e_1} \ldots x_{i_t}^{e_t}$ then w^{-1} is the word $x_{i_t}^{-e_t} \ldots x_{i_1}^{-e_1}$. This allows us to define w^e in the obvious way for all integers e and words w. When w only involves a proper subset of the elements of X and it is important to indicate this, we write $w = w(x_{i_1}, \ldots, x_{i_t})$. A word $w \in X^*$ is *freely reduced* if it does not contain xx^{-1} or $x^{-1}x$ for any $x \in X$.

Let Y be another set, and $v_1, \ldots, v_n$ a sequence of words over Y. Given $w(x_1, \ldots, x_n) \in X^*$ and considering it as a polynomial function in n variables, we obtain a word over Y by substituting v_i for each x_i in w; we denote the resulting word by $w(v_1, \ldots, v_n) \in Y^*$.

Let X be a set and $\mathcal{R}$ a subset of X^*. If $F = F(X)$ is the free group on X, then $\mathcal{R}$ generates a unique normal subgroup R of F. The pair $\{X \mid \mathcal{R}\}$ is called a *presentation* for the group $G = F/R$. The set X is the generating set of G and $\mathcal{R}$ is the defining set of relators of G. We can also consider relations as well as relators in $\mathcal{R}$. A relation has the form $w = v$ for $w, v \in X^*$ and it corresponds to the relator wv^{-1}. We take advantage of this correspondence between relations and relators to use whichever form is convenient in a particular context. The presentation $\{X \mid \mathcal{R}\}$ is a finite presentation if there are only a finite number of elements in both X and $\mathcal{R}$.

We shall also use an internal form of presentation for a group. Define the homomorphism $\sigma\colon F \to G$ by $x \mapsto xR$ for all $x \in F$, and let $g_i = x_i^\sigma$ for each $x_i \in X$. Let $\mathcal{A}_G = \{g_1, \ldots, g_n\}$ and let $\mathcal{R}_G$ be the set of words in $\mathcal{A}_G^*$ obtained as the images of elements of $\mathcal{R}$ under σ. A word in $\mathcal{A}_G^*$ has a natural interpretation as an element of G. We call $\{\mathcal{A}_G \mid \mathcal{R}_G\}$ a presentation for G. The elements of the generating set are elements of G. The elements of the relator set are words over $\mathcal{A}_G$, each of which corresponds to the identity element of G. We define an equivalence relation on $\mathcal{A}_G^*$ such that two words $u, v \in \mathcal{A}_G^*$ are equivalent if and only if they correspond to the same element of the group.

Let $\{\mathcal{A}_G \mid \mathcal{R}_G\}$ be a presentation for a group G. A subgroup K of G is *exhibited* by the presentation if there exists subsets $\mathcal{A}_K \subseteq \mathcal{A}_G$ and $\mathcal{R}_K \subseteq \mathcal{R}_G$ such that $\{\mathcal{A}_K \mid \mathcal{R}_K\}$ is a presentation for K.

Let $w = w(g_1, \ldots, g_n)$ be a word in $\mathcal{A}_G^*$. Let $v_1, \ldots, v_n$ be a sequence of words in $\mathcal{A}_G^*$ and define a map γ from $\mathcal{A}_G$ into $\mathcal{A}_G^*$ by $\gamma\colon g_i \mapsto v_i$ for each $i \in \{1, \ldots, n\}$. Denote by w^γ the word $w(v_1, \ldots, v_n)$. The word w is *invariant under* γ if the words w and w^γ are equivalent; that is, $w = w^\gamma$ as elements of the group.

2.2 Groups

We give a brief description of polycyclic groups and polycyclic presentations; for a more complete description see Chapter 9 of Sims (1994). Let G be a group. If G has a subnormal series

$$G = G_1 > G_2 > \cdots > G_n > G_{n+1} = \langle 1 \rangle, \tag{2.1}$$

such that G_i/G_{i+1} is a cyclic group for all $i \in \{1, \dots, n\}$, then G is a *polycyclic group*. Such a series is a *polycyclic series* for G. For each $i \in \{1, \dots, n\}$ choose an element g_i of G such that $G_i/G_{i+1} = \langle g_i G_{i+1} \rangle$. The sequence $g_1, \dots, g_n$ is a *polycyclic generating sequence* for G based on the polycyclic series (2.1). Note that for each i, the sequence $g_i, \dots, g_n$ is a polycyclic series for G_i.

If G is a finite soluble group, then each factor in a polycyclic series for G is a finite cyclic group. We may refine the series to the point where each factor is a cyclic group of prime order. Let $g_1, \dots, g_n$ be a polycyclic generating sequence with respect to such a refined series. For each $i \in \{1, \dots, n\}$, $g_i G_{i+1}$ generates a cyclic group of order p_i for some prime p_i. Therefore $g_i^{p_i}$ is an element of G_{i+1} and there exists a word w_{ii} over $\{g_{i+1}, \dots, g_n\}$ such that $g_i^{p_i} = w_{ii}$ in G. Similarly, since G_{i+1} is a normal subgroup of G_i, for each $j \in \{i+1, \dots, n\}$ there exists a word w_{ij} over $\{g_{i+1}, \dots, g_n\}$ such that $g_i^{-1} g_j g_i = w_{ij}$ in G. Let $\mathcal{A}_G = \{g_1, \dots, g_n\}$ and let $\mathcal{R}_G$ be the set of relations $g_i^{p_i} = w_{ii}$ and $g_i^{-1} g_j g_i = w_{ij}$ for all $i \in \{1, \dots, n\}$ and $j \in \{i+1, \dots, n\}$. Then $\{\mathcal{A}_G \mid \mathcal{R}_G\}$ is a *power conjugate presentation* for G.

A power conjugate presentation for a group provides a lot of information about the group. In particular, it provides a solution to the word problem, since every element of the group has a unique normal form as a word over $\mathcal{A}_G$. This normal form can be computed by an algorithm called *collection* (see for example Sims 1994). Collection of a word involves a sequence of basic collection steps. A basic collection step replaces either a subword $g_i^{p_i}$ by w_{ii} or a subword $g_j g_i$ by $g_i w_{ij}$ for $j > i$. The first step is applied to the initial word, and each subsequent basic collection step is applied to the result of the previous one. Because the replacement words w_{ii} and w_{ij} do not involve elements of $\{g_1, \dots, g_i\}$, collection of an arbitrary word terminates after a finite number of basic collection steps in a word of the form $g_1^{e_1} g_2^{e_2} \dots g_n^{e_n}$ where $e_i \in \{0, \dots, p_i - 1\}$. A word of this form is called a *normal word*.

2.2.1 Group extensions

Let G and N be groups. A group H is an *extension* of G by N if there exists a normal subgroup L of H such that $L \cong N$ and $H/L \cong G$. Note that this definition of extension is the opposite to that used in many references, and it is sometimes referred to as a *downward extension*.

A *group extension* is an exact sequence of groups and homomorphisms

$$\mathcal{E}\colon\ 1 \longrightarrow M \xrightarrow{\rho} H \xrightarrow{\sigma} G \longrightarrow 1.$$

Clearly H is an extension of G by M. Typically M will be identified with its image M^ρ in H. Where M is abelian, the extension determines a unique map ξ from G into the automorphism group of M, inducing a G-module structure on M. Given $m \in M$ and $g \in G$, we write m^g for the image of m under the automorphism g^ξ. Given an element $\nu \in \operatorname{Aut} M$ we write $m\nu$ for the image of m under ν, which is the natural notation to use when M is written additively.

A map τ from G into H satisfying $g^{\tau\sigma} = g$ for all $g \in G$ is called a *transversal function*. The *factor set* associated with τ is the 2-cocycle $\varphi\colon G \times G \to M$ defined by the equation

$$g^\tau h^\tau = (gh)^\tau (g,h)\varphi$$

for all $g, h \in G$. The factor set satisfies the equation

$$(a, bc)\varphi + (b,c)\varphi = (ab, c)\varphi + ((a,b)\varphi)^c \tag{2.2}$$

for all $a, b, c \in G$. Note that every element of H has a unique expression as a product $g^\tau m$ with $g \in G$ and $m \in M$.

2.2.2 Group algebras

Let G be a finite group and p a prime. The *group algebra* of G over $\mathbb{F}_p$, denoted by $\mathbb{F}_pG$, is the set of all formal sums $\sum_{g\in G} a_g g$ where $a_g \in \mathbb{F}_p$ for all $g \in G$. It is an $\mathbb{F}_p$-algebra with addition and multiplication defined by

$$\left(\sum_{g\in G} a_g g\right) + \left(\sum_{g\in G} b_g g\right) = \sum_{g\in G} (a_g + b_g) g$$
$$\left(\sum_{g\in G} a_g g\right)\left(\sum_{g\in G} b_g g\right) = \sum_{g\in G}\left(\sum_{hk=g} a_h b_k\right) g,$$

and the obvious linear action of the field.

Let $X = \{x_1, \ldots, x_n\}$ and let F be the free group on X. Let $\sigma : x_i \mapsto g_i$ define an epimorphism of F onto a finitely presented group G. The relators of

G correspond to elements of F which generate a subgroup whose normal closure is the kernel of σ. Note that σ extends linearly to an algebra homomorphism $\mathbb{F}_p F \to \mathbb{F}_p G$.

The free differential calculus was introduced by Fox (1953). The following definition of Fox derivatives is from Johnson (1980).

Definition 2.1 Let $w = y_1 \dots y_t$ be an element of F, where $y_i \in X^{\pm}$ for $i \in \{1, \dots, t\}$. The *Fox derivative* of w with respect to $x_i \in X$ is an element of $\mathbb{F}_p F$ given by

$$\frac{\partial w}{\partial x_i} = \sum_{j=1}^{t} a_j, \qquad a_j = \begin{cases} y_{i+1} \cdots y_t, & \text{if } y_j = x_i, \\ -y_i \cdots y_t, & \text{if } y_j = x_i^{-1}, \\ 0, & \text{otherwise.} \end{cases} \tag{2.3}$$

Note that the above definition is for a *right* Fox derivative. There is a corresponding definition for left Fox derivatives, but they are not needed here. Let $w(g_1, \dots, g_n)$ be a word in $\mathcal{A}_G^*$. The Fox derivative of $w(x_1, \dots, x_n)$ with respect to x_i is an element of $\mathbb{F}_p F$, and is mapped under σ to an element of $\mathbb{F}_p G$. We call this element of $\mathbb{F}_p G$ the Fox derivative of $w(g_1, \dots, g_n)$ with respect to g_i, and denote it by $\partial w / \partial g_i$.

2.2.3 Modules

Let G be a group. By an $\mathbb{F}_p G$-module of dimension k we shall mean an elementary abelian p-group M of order p^k together with a homomorphism ξ mapping G into $\mathrm{Aut}\, M$. With respect to a minimal generating set for M, $\mathrm{Aut}\, M$ has a natural identification as the general linear group $\mathrm{GL}(k, p)$, in which case the $\mathbb{F}_p G$-module structure of M is defined by a homomorphism $\xi \colon G \to \mathrm{GL}(k, p)$.

Let M be a $\mathbb{F}_p G$-module of dimension k with associated representation ξ. The endomorphism algebra of M, denoted by $\mathrm{End}_G(M)$, is the $\mathbb{F}_p$-algebra of all abelian group endomorphisms of M that commute with the action of G on M. Let N be a $\mathbb{F}_p G$-module of dimension l with associated representation η. Denote by $\mathrm{Hom}_G(M, N)$ the abelian group of all $\mathbb{F}_p G$-module homomorphisms between M and N. Fix generating sets for M and N as abelian groups. The endomorphism algebra $\mathrm{End}_G(M)$ has a natural interpretation as a subalgebra of $\mathbb{M}(k, \mathbb{F}_p)$ consisting of those matrices A satisfying $g^\xi A = A g^\xi$ for all $g \in G$. The groups $\mathrm{Hom}_G(M, N)$ has a natural interpretation as the abelian group of all $k \times l$ matrices over $\mathbb{F}_p$ satisfying $g^\xi A = A g^\eta$ for all $g \in G$.

2.3 Schreier generators

Let F be the free group on the set $X = \{x_1, \ldots, x_n\}$. Let H be a subgroup of F. The (right) cosets of H partition F and by choosing a single element from each coset we obtain a *(right) transversal* for H in F. If T is a transversal for H in F, then for $w \in F$ we denote by $\overline{w}$ the unique element of $Hw \cap T$.

A transversal T is a *Schreier transversal* if for any word uv in T, the word u is also in T. That is, T is a Schreier transversal if it contains all initial subwords of each of its elements. The following theorem defines the set of *Schreier generators* associated with a transversal. Note that the transversal need not be a Schreier transversal.

Theorem 2.2 The elements of the set

$$\{tx(\overline{tx})^{-1} \mid t \in T, x \in X\}$$

generate the subgroup H.

Proof See for example Johnson (1980), Section 2.5. ■

This theorem is the basis for a fundamental algorithm that computes a generating set for a subgroup of a group. The *orbit-stabiliser* algorithm takes as input a generating set $\mathcal{A}$ for a group G, a set Z that G acts on and an initial point $z_0 \in Z$. It computes a transversal T for the stabiliser S in G of z_0. It achieves this by computing the orbit of z_0 under the action of G. Each point in the orbit is the image of z_0 under the action of some element of G, and the set T of all these elements is a transversal for S in G. The Schreier generators associated with T form a generating set for S. An alternative form of the algorithm relies on simply having a test for membership of the subgroup S.

This algorithm is used in a number of calculations described later in the thesis. Note that the method of constructing the orbit has not been specified, and different methods for filling out the orbit lead to different transversals. A pseudo-code description of the generic algorithm follows.

Algorithm 2.3 ORBIT-STABILISER

Input: A set $\mathcal{A}$ generating a group G and a membership test for when $g \in G$ lies in a subgroup S of G.

Output: A transversal for S in G and a generating set for S.

Set T to $\{1\}$
repeat
    Note the size of T
    for each $t \in T$ and $\alpha \in \mathcal{A}$ do
        Set u to $t\alpha$
        Add u to T if $Su \cap T = \{\}$
    end do
until the size of T has not increased
Return T and the Schreier generators associated with T

The check for $Su \cap T = \{\}$ is where either the partial orbit or the membership test is required. When the orbit of some point z_0 is being computed, this check is for whether the new point is already in the partial orbit constructed so far. When the algorithm is relying on a test for membership for S, this check is for whether $t'u^{-1}$ lies in S for each t' in the partial transversal, since $t'u^{-1} \in S$ implies that u is in the same coset as t'.

Chapter 3

Lifting automorphisms

We consider the problem of computing the automorphism group of a finite soluble group G that is given by a power conjugate presentation $\{\mathcal{A}_G \mid \mathcal{R}_G\}$. The description of the automorphism group that we will obtain is a generating set of automorphisms, where each automorphism is represented by its action on $\mathcal{A}_G$ as words in $\mathcal{A}_G^*$. That is, if $\mathcal{A}_G = \{g_1, \ldots, g_n\}$, then an automorphism α is represented by a sequence of words $w_1, \ldots, w_n \in \mathcal{A}_G^*$ such that w_i is a word corresponding to the group element g_i^α for $i \in \{1, \ldots, n\}$. In addition the order of the automorphism is computed and some of the elements of the generating set will be identified as inner automorphisms.

The automorphism group calculation relies on an algorithm of Leedham-Green (described in Eick 1993) which takes the given power conjugate presentation and computes a *special power conjugate presentation* for G. We describe these special power conjugate presentations in Chapter 6, but here we need only note that such a presentation is a consistent power conjugate presentation based on a particular polycyclic series for the group. The polycyclic series is a refinement of a series

$$G = G_0 > G_1 > G_2 > \ldots > G_l > G_{l+1} = 1,$$

where, for all $i \in \{1, \ldots, l\}$, the quotient G_i/G_{i+1} is an elementary abelian p_i-group for some prime p_i, and G_i/G_{i+1} is characteristic in G/G_{i+1}. We describe an algorithm that takes a generating set for the automorphism group of G/G_i and computes from this a generating set for the automorphism group of G/G_{i+1}. The automorphism group of G is computed by applying this basic step recursively.

We shall describe sections of the algorithm as they apply to a more general problem than that required for the basic step of the soluble group automorphism algorithm. Let G be a finitely presented group and M an elementary abelian

p-group for some prime p. Let H be an extension of G by M, so that we have a group extension

$$\mathcal{E}\colon\ 1 \longrightarrow M \overset{\rho}{\longrightarrow} H \overset{\sigma}{\longrightarrow} G \longrightarrow 1.$$

Given a suitable presentation for H together with a generating set for $\operatorname{Aut} G$, the algorithm will compute a generating set for the subgroup of $\operatorname{Aut} H$ consisting of all the automorphisms of H that map M^ρ to itself. We shall return to the specific context of soluble groups in Chapter 6.

3.1 Scenario

We begin by specifying the conditions on the presentation for H that we require. We then define a collection process for words in the generating set of H and extend it to a collection process for a larger class of words related to the extension. We use this extended collection process to compute systems of linear equations related to group theoretic questions, in such a way that a question has an affirmative answer precisely when the system of linear equations has a solution (such techniques are described by Plesken 1987).

Definition 3.1 Let $\mathcal{E}$ be a group extension

$$\mathcal{E}\colon\ 1 \longrightarrow M \overset{\rho}{\longrightarrow} H \overset{\sigma}{\longrightarrow} G \longrightarrow 1, \tag{3.1}$$

where G is a finitely presented group with presentation $\{\mathcal{A}_G \,|\, \mathcal{R}_G\}$ and M is an elementary abelian p-group of order p^k for some prime p. Let $\mathcal{A}_G = \{g_1, \ldots, g_n\}$. A presentation $\{\mathcal{A}_H \,|\, \mathcal{R}_H\}$ for H *exhibits* the extension $\mathcal{E}$ if the following conditions hold:

1) $\mathcal{A}_H$ can be partitioned into two subsets $\{h_1, \ldots, h_n\}$ and $\{m_1, \ldots, m_k\}$ such that $M^\rho = \langle m_1, \ldots, m_k \rangle$ and $\sigma\colon h_i \mapsto g_i$ for $i \in \{1, \ldots, n\}$.
2) The relation set $\mathcal{R}_H$ is the disjoint union of $\mathcal{R}_H^{(1)}$, $\mathcal{R}_H^{(2)}$ and $\mathcal{R}_H^{(3)}$, where:
   a) $\mathcal{R}_H^{(1)}$ contains relations $m_i^p = 1$ and $m_j^{m_i} = m_j$ for all $i \in \{1, \ldots, k\}$ and $j \in \{i+1, \ldots, k\}$.
   b) $\mathcal{R}_H^{(2)}$ contains relations $m_j^{h_i} = w_{ij}(m_1, \ldots, m_k)$ for all $i \in \{1, \ldots, n\}$ and $j \in \{1, \ldots, k\}$.

c) $\mathcal{R}_H^{(3)}$ contains a relation $r(h_1, \ldots, h_n) = s(m_1, \ldots, m_k)$ for each relator $r(g_1, \ldots, g_n)$ in $\mathcal{R}_G$.

Clearly a presentation of this form defines a group H which is an extension of G by M. While these conditions look quite restrictive, such presentations arise naturally in a number of different ways:

1) When computing the automorphism group of a soluble group S, which is our primary motivation. Here G and H are both quotients of S by successive terms of a subnormal series for S. They both inherit power conjugate presentations from that of S.
2) G is a finitely presented group, M is an elementary abelian p-group, and ξ is a homomorphism from G into $\operatorname{Aut} M$ inducing a G-module structure on M. A presentation for the semi-direct product of G by M is easily constructed from the presentation for G and the action of G on M via ξ.
3) G is a finitely presented group, M is an elementary abelian p-group, ξ is a homomorphism from G into $\operatorname{Aut} M$, and $\varphi\colon G \times G \to M$ a 2-cocycle. As in the semi-direct product construction above, we can obtain a presentation for an arbitrary extension of G by M using this information (see for example Sims 1994, Section 11.3).

Let τ be a transversal function from G into H and φ the factor set associated with τ. Let ξ be the homomorphism from G into $\operatorname{Aut} M$. With respect to the generating set $\{m_1, \ldots, m_k\}$ of M, the representation ξ maps G into $\mathrm{GL}(k, p)$, where the j-th row of the matrix g_i^ξ is given by the exponent vector of $w_{ij}(m_1, \ldots, m_k)$ in M. The representation ξ extends linearly to a map $\xi\colon \mathbb{F}_p G \to \mathbb{M}(k, \mathbb{F}_p)$ which defines an $\mathbb{F}_p G$-module structure for M (see for example Sims 1994, Proposition 10.8.1).

Note that words in $\mathcal{A}_H^*$ that involve only generators from M inherit the natural normal form for elements of M, and that, with respect to $\{m_1, \ldots, m_k\}$, the elements of M may be represented as vectors of length k over $\mathbb{F}_p$. We may assume that the presentation $\{\mathcal{A}_H \mid \mathcal{R}_H\}$ has all words $w_{ij} \in M$ in normal form.

A word $w \in \mathcal{A}_H^*$ is in *normal form* if it does not contain a subword of the form $m_j h_i$ and all subwords $v(m_1, \ldots, m_k)$ of w are in normal form for M. If G is given by a power conjugate presentation, then the presentation for H is also a power conjugate presentation and it is natural to also require that there are no subwords $h_j h_i$ with $i < j$. However, here we only assume that G is a finitely presented group, and as such it may not have a collection process defined for it.

Definition 3.2 Let $\{\mathcal{A}_H \mid \mathcal{R}_H\}$ be a presentation for a group H that exhibits the extension $\mathcal{E}$, and let w be a word over $\mathcal{A}_H$. A *basic collection step* applied to w is one of the following:

1) choose a non-normal subword $m_j h_i$ in w and replace it by the equivalent normal word $h_i w_{ij}(m_1, \ldots, m_k)$;
2) choose a non-normal subword $v(m_1, \ldots, m_k)$ in w and replace it by its equivalent normal form in M.

The word obtained as the result of a basic collection step is equivalent to w. A *collection* of a word $w \in \mathcal{A}_H^*$ consists of a sequence of basic collection steps, each one applied to the result of the previous step. When no more basic steps can be performed the process terminates and returns a word in normal form that is equivalent to w.

If a collection terminates after a finite number of steps, the end result is a word of the form $u(h_1, \ldots, h_n)\, v(m_1, \ldots, m_k)$ with v in normal form. Note that the definition of collection does not specify how a non-normal subword is chosen at each basic collection step. Recall that a collection with respect to a power conjugate presentation always terminates after a finite number of steps, irrespective of the particular choice of non-normal subword at each step. The next result proves that the same is also true here.

Lemma 3.3 A collection of a word in $\mathcal{A}_H^*$ terminates after a finite number of basic collection steps and returns a normal word that is equivalent to w.

Proof Let w be a word in $\mathcal{A}_H^*$, and let $\mathcal{A} = \{h_1, \ldots, h_n\}$ be the set of generators of H that lie outside M. We argue by induction on the number of occurrences of elements of $\mathcal{A}$ in w. Suppose that the result is true for all words with at most $k-1$ occurrences of generators from $\mathcal{A}$, and let w be a word containing exactly k occurrences of generators from $\mathcal{A}$. Write w as a product $u_1 h v_1$ where u_1 lies in $\mathcal{A}_M^*$, h is an element of $\mathcal{A}$, and v_1 is a word in $\mathcal{A}_H^*$ containing $k-1$ generators from $\mathcal{A}$. A collection step on w either replaces a subword of $u_1 h$ or a subword of v_1, since no collection step replaces subwords starting with a generator from $\mathcal{A}$. Take a sequence of basic collection steps on w and consider the initial sequence of steps that do not replace a subword containing the generator h. This initial sequence can be partitioned into two subsequences, the first containing all the replacements that take place to the left of h, the second containing all the replacements that take place to the right of h. Since u_1 is a word from M it will reach normal form after a finite number of collection steps. By the inductive hypothesis, v_1 collects to a normal word after a finite number of collection steps. Therefore the initial sequence of collection steps applied to w is a finite sequence that results in a word of the form $u_1'hv_1'$. If u_1' is the empty word then we are finished, otherwise the next collection step involves h. After this collection step we have a word u_2hv_2 where the length of u_2 is strictly less than the length of u_1. Therefore, after a finite number of collection steps we obtain a word hv_t where v_t is again a word with at most $k-1$ occurrences of generators from $\mathcal{A}$. ■

Collecting two different words in $\mathcal{A}_H^*$ that represent the same element of H may result in different normal words. Let w be a word over $\mathcal{A}_H$. The positions of the h_i in w determine their positions in the normal form, since they are never interchanged by a collection step. A consequence of this fact is that collection of a single word w will always result in the same normal word irrespective of the particular sequence of basic steps performed. If uv and uv' are the collected forms for w obtained by different sequences of collection steps, then v and v' represent the same element of M and, since they are both in normal form for M, they must be equal as words.

We now define a more general class of words related to our extension $\mathcal{E}$. The definition is motivated by the fact that if u is some element of M, and e is an element of the group algebra $\mathbb{F}_pG$, then e acts on u to produce another element u^e of M. We append a set of "indeterminate" elements of M to our generating set $\mathcal{A}_H$ and construct generalised words over this extended set, allowing exponents in $\mathbb{F}_pG$ for the new generators.

Definition 3.4 For a set $U = \{u_1, \ldots, u_t\}$ we define $(\mathcal{A}_H, U)^*$ to be the set of all generalised words of the form

$$w = a_1^{e_1} \ldots a_s^{e_s},$$

where $a_i \in \mathcal{A}_H \cup U$ and the exponents e_i are integers when $a_i \in \mathcal{A}_H$ and elements of $\mathbb{F}_pG$ when $a_i \in U$.

We say that a word w in $(\mathcal{A}_H, U)^*$ is in normal form if it contains no subwords of the following forms: $m_j h_i$, $u_k^e h_i$, $u_k^e m_j$, or any non-normal $v(m_1, \ldots, m_k)$. We extend the collection process for words in $\mathcal{A}_H^*$ to a collection process for words in $(\mathcal{A}_H, U)^*$. The additional basic collection steps express the relationships that must hold when the elements of U are thought of as elements of M.

Definition 3.5 Let $U = \{u_1, \ldots, u_t\}$ be a set, and let w be a generalised word in $(\mathcal{A}_H, U)^*$. A basic collection step applied to w is one of the following:

1) replace a subword $m_j h_i$ by $h_i w_{ij}(m_1, \ldots, m_k)$.
2) replace a non-normal subword $v(m_1, \ldots, m_k)$ by its equivalent normal form.
3) replace $u_j^e h_i$ by $h_i u_j^{eg_i}$ where $g_i = h_i^\sigma \in \mathcal{A}_G$.
4) replace $u_j m_i$ by $m_i u_j$.
5) replace $u_j^f u_i^e$ by $u_i^e u_j^f$ for $i < j$.
6) replace $u_i^e u_i^f$ by u_i^{e+f} if $e + f \neq 0$ and by the empty word otherwise.

A *collection* of a word $w \in (\mathcal{A}_H, U)^*$ consists of a sequence of basic collection steps, each one applied to the result of the previous step. When no more basic steps can be performed the process terminates and returns a word in normal form.

There is also a natural notion of equivalence for words in $(\mathcal{A}_H, U)^*$. If $w(\mathcal{A}_H, u_1, \ldots, u_t)$ and $v(\mathcal{A}_H, u_1, \ldots, u_t)$ are elements of $(\mathcal{A}_H, U)^*$, then we say that w and v are equivalent if for all maps μ from U into M the words $w(\mathcal{A}_H, u_1^\mu, \ldots, u_t^\mu)$ and $v(\mathcal{A}_H, u_1^\mu, \ldots, u_t^\mu)$ are equivalent in the usual sense. This definition of equivalence for generalised words is consistent with the interpretation of U as a set of elements of M.

Lemma 3.6 Let w be an element of $(\mathcal{A}_H, U)^*$, and let v be the word obtained after a single basic collection step has been applied to w. Then w and v are equivalent.

Proof It suffices to show that the non-normal subword is equivalent to its replacement for a basic collection step. For collection steps 4) and 5) it follows from the fact that M is abelian. For collection steps 3) and 6) it follows from the fact that M is an $\mathbb{F}_p G$-module. ■

Clearly equivalence of generalised words is an equivalence relation. An immediate consequence of Lemma 3.6 is that if the collection terminates and returns a word in normal form, then this word is equivalent to the initial word.

Lemma 3.7 Collection of a word in $(\mathcal{A}_H, U)^*$ terminates after only finitely many basic collection steps and returns a normal word that is equivalent to w.

Proof The proof is a simple modification of the proof of Lemma 3.3. Rather than an induction on the number of occurrences of $h_1, \ldots, h_n$ in w, perform induction on the number of occurrences of elements of U and apply Lemma 3.3 to subwords that do not contain elements of U. ∎

The result of a collection of a word $w \in (\mathcal{A}_H, U)^*$ is a word of the form

$$v_1(h_1, \ldots, h_n)\, v_2(m_1, \ldots, m_k)\, u_1^{e_1} \ldots u_t^{e_t},$$

where the e_i are elements of $\mathbb{F}_p G$ and v_1, v_2 are elements of $\mathcal{A}_H^*$. As before, equivalent words may collect to different normal words, but a single initial word collects to a unique normal form irrespective of the particular sequence of basic collection steps.

3.2 Automorphisms of an extension

An automorphism γ of H normalises M if $M^\gamma = M$. If γ normalises M, then it induces an automorphism of G which we denote by γ/M. Given an automorphism α of G we say that α *lifts to* H if and only if there exists an automorphism γ of H such that $\gamma/M = \alpha$. Such a γ is called *a lifting of* α to H, and when α is the identity automorphism of G we call γ a lifting of the identity. We shall use lowercase Greek letters for automorphisms, typically choosing γ for an automorphism of H, α and β for automorphisms of G, ι for the identity automorphism of G, and ν for an automorphism of M.

Definition 3.8 Let $\mathcal{E}$ be the group extension given by the short exact sequence (3.1). The automorphism group of $\mathcal{E}$, denoted by $\operatorname{Aut}\mathcal{E}$, is the subgroup of all automorphisms of H that normalise M.

Note that when M is a characteristic subgroup of H the automorphism group of $\mathcal{E}$ is all of $\operatorname{Aut} H$.

Let A be a generating set for $\operatorname{Aut} G$. We aim to compute a generating set for $\operatorname{Aut} \mathcal{E}$. The representation that we use for automorphisms is as follows: an automorphism α of G is specified by a sequence $\{w_1, \ldots, w_n\}$, with $w_i \in \mathcal{A}_G^*$ for $i \in \{1, \ldots, n\}$, such that if $g_i \in \mathcal{A}_G$ then w_i is a word representing the element g_i^α of G. Given two such automorphisms, multiplying them is simply a matter of substitution: if α and β are automorphisms and $g_i^\alpha = w(g_1, \ldots, g_n)$, then the automorphism $\alpha\beta$ maps g_i to $w(g_1^\beta, \ldots, g_n^\beta)$. The computation of products is the only requirement for the automorphism group algorithm when $\mathcal{E}$ is split. In the general case, however, we also require a means of computing the inverse of an automorphism.

Let γ be an automorphism of $\mathcal{E}$. Then γ/M is an automorphism of G and $\gamma|_M$ is an automorphism M. Therefore we have a natural homomorphism

$$\vartheta\colon \operatorname{Aut} \mathcal{E} \longrightarrow \operatorname{Aut} G \times \operatorname{Aut} M, \tag{3.2}$$

where the image of γ under ϑ is $(\gamma/M, \gamma|_M)$. We shall divide the computation of $\operatorname{Aut} \mathcal{E}$ into two parts based on the short exact sequence

$$1 \longrightarrow \operatorname{Ker} \vartheta \longrightarrow \operatorname{Aut} \mathcal{E} \longrightarrow \operatorname{Im} \vartheta \longrightarrow 1 \tag{3.3}$$

associated with (3.2). We will describe effective algorithms for computing generating sets for the image and kernel of ϑ. Let $\mathcal{A}_C$ be a generating set for $\operatorname{Im} \vartheta$, and $\mathcal{A}_K$ a generating set for $\operatorname{Ker} \vartheta$. Let $\mathcal{A}'_C$ be a set containing at least one preimage of each element of $\mathcal{A}_C$ under ϑ. The proof of the next result is trivial and omitted.

Lemma 3.9 $\mathcal{A}'_C \cup \mathcal{A}_K$ is a generating set for $\operatorname{Aut} \mathcal{E}$.

Therefore the problem of determining a generating set of $\operatorname{Aut} \mathcal{E}$ can be divided into two sub-problems. The first is computing a generating set for the image of ϑ, and taking a set of preimages of its elements. This step is handled in the next two chapters, first for split extensions and then for non-split extensions. The second step is to compute a generating set for the kernel of ϑ, and this is considered in the remainder of this chapter.

3.3 Derivations

The kernel of ϑ has a natural interpretation as the subgroup of automorphisms of H that correspond to derivations from G to M.

Definition 3.10 A derivation from a group G to a G-module M is a map $\delta\colon G \to M$ satisfying

$$(ab)^{\delta} = (a^{\delta})^{b} + b^{\delta} \tag{3.4}$$

for all $a, b \in G$. The derivations from G to M form an abelian group, denoted $\mathrm{Der}(G, M)$, where addition is defined by $a^{\delta_1+\delta_2} = a^{\delta_1} + a^{\delta_2}$ for all $a \in G$ and $\delta_1, \delta_2 \in \mathrm{Der}(G, M)$.

It is well known (see for example Robinson 1982, Chapter 11) that there is a bijection from $\mathrm{Der}(G, M)$ to the set of complements of M in the semi-direct product of G by M. There is also a bijection between the subgroup of *inner derivations* in $\mathrm{Der}(G, M)$, denoted by $\mathrm{Inn}(G, M)$, and the conjugacy class of complements of M in the semi-direct product of G by M.

Lemma 3.11 The kernel of ϑ is isomorphic to $\mathrm{Der}(G, M)$.

Proof (see also Robinson 1981) Let $\delta \in \mathrm{Der}(G, M)$ and define $\gamma_{\delta}\colon H \to H$ by

$$(g^{\tau} m)^{\gamma_{\delta}} = g^{\tau} g^{\delta} m \tag{3.5}$$

for all $g^{\tau} m \in H$. That γ_{δ} is an endomorphism of H follows from (3.4). Noting that $\gamma_{\delta_1+\delta_2} = \gamma_{\delta_1}\gamma_{\delta_2}$, we have γ_{δ} an automorphism of H since $\gamma_{-\delta}$ is a two-sided inverse for it. Clearly γ_{δ} lies in $\mathrm{Ker}\,\vartheta$, and hence the map $\eta\colon \delta \mapsto \gamma_{\delta}$ is a group homomorphism from $\mathrm{Der}(G, M)$ to $\mathrm{Ker}\,\vartheta$. If γ_{δ} is the identity on H, then $g^{\delta} = 1$ for all $g \in G$, implying that η is injective. Now suppose γ lies in the kernel of ϑ. So $(g^{\tau} m)^{\gamma} = g^{\tau} m_g m$ for some $m_g \in M$. Therefore

$$\begin{aligned}
(gh)^{\tau}\, m_{gh} &= (gh)^{\tau\gamma} \\
&= (g^{\tau} h^{\tau}\, (g, h)\varphi^{-1})^{\gamma} \\
&= g^{\tau} m_g\, h^{\tau}\, m_h\, (g, h)\varphi^{-1} \\
&= g^{\tau} h^{\tau}\, {m_g}^{h^{\tau}}\, m_h\, (g, h)\varphi^{-1} \\
&= (gh)^{\tau}\, {m_g}^{h^{\tau}}\, m_h.
\end{aligned}$$

Hence $m_{gh} = m_g^h m_h$ so $\delta\colon g \mapsto m_g$ is a derivation, and γ is clearly the corresponding automorphism γ_δ. ■

We calculate a generating set for the abelian group of derivations by converting the problem to that of finding a basis for the solution space of a homogeneous system of linear equations over $\mathbb{F}_p$. A major step in this direction is the projective sequence for a finitely presented group given below in Theorem 3.13.

Definition 3.12 The *augmentation mapping* $\epsilon\colon \mathbb{F}_pG \to \mathbb{F}_p$ is defined by

$$\epsilon\colon \sum_{g\in G} f_g g \longmapsto \sum_{g\in G} f_g.$$

The *augmentation ideal* of $\mathbb{F}_pG$ is the kernel of ϵ and is denoted by $\mathrm{Aug}(\mathbb{F}_pG)$. The augmentation ideal is generated as an $\mathbb{F}_pG$-module by the set $\{g_i - 1 \mid g_i \in \mathcal{A}_G\}$ (Johnson 1990, Proposition 11.2.1).

Theorem 3.13 Let $\{g_1, \ldots, g_n \mid r_1, \ldots, r_s\}$ be a finite presentation for a group G. Define maps $\beta\colon (\mathbb{F}_pG)^n \to \mathbb{F}_pG$ and $\alpha\colon (\mathbb{F}_pG)^s \to (\mathbb{F}_pG)^n$ by

$$\beta\colon (y_1, \ldots, y_n) \longmapsto \sum_{i=1}^{n} (g_i - 1)y_i,$$

$$\alpha\colon (y_1, \ldots, y_s) \longmapsto (v_1, \ldots, v_n), \quad v_j = \sum_{i=1}^{s} y_i \frac{\partial r_j}{\partial g_i}.$$

Then the sequence

$$(\mathbb{F}_pG)^s \xrightarrow{\alpha} (\mathbb{F}_pG)^n \xrightarrow{\beta} \mathbb{F}_pG \xrightarrow{\epsilon} \mathbb{F}_p \longrightarrow 0 \tag{3.6}$$

is a projective resolution of the trivial $\mathbb{F}_pG$-module.

Proof See Johnson (1990), Chapter 11. ■

We obtain from (3.6) the exact sequence

$$(\mathbb{F}_pG)^s \xrightarrow{\alpha} (\mathbb{F}_pG)^n \xrightarrow{\beta} \mathrm{Aug}(\mathbb{F}_pG) \longrightarrow 0,$$

which induces an exact sequence of homomorphism groups,

$$\mathrm{Hom}_G((\mathbb{F}_pG)^s, M) \xleftarrow{\alpha^*} \mathrm{Hom}_G((\mathbb{F}_pG)^n, M) \xleftarrow{\beta^*} \mathrm{Hom}_G(\mathrm{Aug}(\mathbb{F}_pG), M) \longleftarrow 0.$$

The abelian group $\mathrm{Hom}_G((\mathbb{F}_p G)^t, M)$ is isomorphic to M^t for each positive integer t, so the above sequence corresponds to

$$M^s \xleftarrow{\alpha^*} M^n \xleftarrow{\beta^*} \mathrm{Hom}_G(\mathrm{Aug}(\mathbb{F}_p G), M) \longleftarrow 0, \tag{3.7}$$

where β^* is the homomorphism

$$\beta^* \colon f \longmapsto ((g_1 - 1)f, \ldots, (g_n - 1)f)$$

for all $f \in \mathrm{Hom}_G(\mathrm{Aug}(\mathbb{F}_p G), M)$, and α^* is

$$\alpha^* \colon (y_1, \ldots, y_n) \longmapsto (v_1, \ldots, v_s), \quad v_j = \sum_{i=1}^{n} y_i \left(\frac{\partial r_j}{\partial g_i}\right)^{\xi}.$$

The kernel of α^* consists of those vectors $y = (y_1, \ldots, y_n) \in M^n$ for which

$$\sum_{i=1}^{n} y_i \left(\frac{\partial r_j}{\partial g_i}\right)^{\xi} = 0$$

for all $r_j \in \mathcal{R}_G$. Such a vector corresponds to a homomorphism f_y from $\mathrm{Aug}(\mathbb{F}_p G)$ to M defined by

$$f_y \colon (g_i - 1) \longmapsto y_i.$$

The following theorem identifies the correspondence between the kernel of α^* and the group of derivations (see for example Robinson 1982, Proposition 11.4.5).

Lemma 3.14 $\mathrm{Hom}_G(\mathrm{Aug}(\mathbb{F}_p G), M)$ is isomorphic to $\mathrm{Der}(G, M)$, with the isomorphism defined by $f \mapsto \delta_f$ where $g^{\delta_f} = (g-1)f$ for all $g \in G$.

We can now define a system of linear equations over $\mathbb{F}_p$ whose solution space is isomorphic to the group of derivations. If $y_i \in M$ is the vector $(y_{i1}, \ldots, y_{ik}) \in \mathbb{F}_p^k$, then $(y_1, \ldots, y_n) \in M^n$ corresponds to a derivation if and only if

$$\sum_{i=1}^{n} (y_{i1}, \ldots, y_{ik}) \left(\frac{\partial r_j}{\partial g_i}\right)^{\xi} = 0 \tag{3.8}$$

for all $j \in \{1, \ldots, s\}$. The image of a Fox derivative under ξ is a $k \times k$ matrix over $\mathbb{F}_p$, and so the s matrix equations (3.8) correspond to the following system of linear equations in the variables y_{ij}:

$$(y_{11},\ldots,y_{1k},y_{21},\ldots,y_{nk})\begin{pmatrix} A_{11} & A_{12} & \ldots & A_{1s} \\ A_{21} & A_{22} & \ldots & A_{2s} \\ \vdots & \vdots & \ddots & \vdots \\ A_{n1} & A_{n2} & \ldots & A_{ns} \end{pmatrix} = (0,\ldots,0), \qquad (3.9)$$

where A_{ij} is the submatrix $\left(\partial r_j/\partial g_i\right)^{\xi}$ for $i \in \{1,\ldots,n\}$ and $j \in \{1,\ldots,s\}$. This provides a convenient means to compute an $\mathbb{F}_p$-basis for the vector space of derivations $\mathrm{Der}(G,M)$. It reduces the problem to computing Fox derivatives of the relations, and then solving a system of ks equations in nk unknowns.

Applying the isomorphism of Lemma 3.11 to the results of the derivation computation, we obtain an algorithm with the following description:

Algorithm 3.15 KERNEL-THETA

Input: The finite presentation for G and the representation ξ.

Output: A generating set $\{\gamma_1,\ldots,\gamma_t\}$ for $\mathrm{Ker}\,\vartheta$.

The key to the algorithm is the computation of the Fox derivatives of a word $w \in \mathcal{A}_G^*$ with respect to each of the elements of $\mathcal{A}_G$. The following pseudo-code computes the matrix images of the Fox derivatives of a word with respect to each of the generators from $\mathcal{A}_G$ in a single pass through the word. The word used as input has the form $w = y_1 \ldots y_t$, where $y_i \in \mathcal{A}_G^{\pm}$. The derivatives are computed by iterating a basic step across the word from right to left. The matrix u is the image under ξ of the subword of w to the right of working position.

```
Set L_1, ..., L_t to the k × k zero matrix over F_p
Set u to the k × k identity matrix over F_p
for i ∈ {t, ..., 1} do
    y_i is g_k^e for some k ∈ {1, ..., n} and e ∈ {±1}
    if e > 0 then
        Set L_k to L_k + u
    else
        Set L_k to L_k − y_i^ξ u
    end if
    Set u to y_i^ξ u
end do
Return L_1, ..., L_t
```

After computing the Fox derivatives of each relator, the system of linear equations is solved and a basis for the vector space of derivations has been found. From this basis a set of automorphisms is easily constructed.

Chapter 4

Split case

We consider the lifting step for $\mathcal{E}$ a split extension. We define an action of the automorphism group of G on a set of $\mathbb{F}_p G$-modules related to M. The action has the property that an automorphism of G lifts to H precisely when it maps M to a module isomorphic to M. An algorithm for deciding when two modules are isomorphic is then described. It relies on computing a decomposition of a module into indecomposable summands. Some of the algorithms developed by Schneider (1990) play a crucial role in this computation. Liftings of the identity that do not correspond to derivations are shown to correspond to $\mathbb{F}_p G$-module automorphisms of M, and the information used to find a decomposition of M is reused to produce a generating set for the subgroup of these liftings.

4.1 Scenario

We use the notation of Chapter 3. We assume throughout this chapter that $\mathcal{E}$ is a split extension, and that $\{h_1, \ldots, h_n\}$ generates a complement, Q, for M in H. Hence $Q \cong G$ and we can choose a transversal function $\tau\colon G \to H$ that maps G into Q. Note that τ is a group homomorphism. Let $A = \{\alpha_1, \ldots, \alpha_t\}$ be the given generating set for $\operatorname{Aut} G$.

We begin by defining an action of the automorphism group of G on $\mathbb{F}_p G$-modules. We shall see that this action is closely related to the automorphism lifting problem.

4.2 An action on modules

With respect to the fixed generating set $\{m_1, \ldots, m_k\}$, the subgroup M has a natural interpretation as the $\mathbb{F}_p G$-module associated with the representation $\xi\colon G \to \mathrm{GL}(k,p)$.

Definition 4.1 Let N be an $\mathbb{F}_p G$-module associated with a representation $\eta:$ $G \to \mathrm{GL}(k,p)$. For $\alpha \in \mathrm{Aut}\, G$, we define N^α to be the module associated with the representation $\eta' = \alpha^{-1}\eta$.

Let $\mathcal{M} = \{M^\alpha \mid \alpha \in \mathrm{Aut}\, G\}$, a finite set of $\mathbb{F}_p G$-modules.

Lemma 4.2 The action of $\mathrm{Aut}\, G$ on $\mathcal{M}$ given by Definition 4.1 is a permutation representation of $\mathrm{Aut}\, G$ on $\mathcal{M}$.

Proof $(M^\alpha)^\beta = M^{\alpha\beta}$ is a direct consequence of $\beta^{-1}\alpha^{-1}\xi = (\alpha\beta)^{-1}\xi$. Clearly $M^\iota = M$. ■

Although the representations associated with M^α and M^β for $\alpha, \beta \in \mathrm{Aut}\, G$ have identical images in $\mathrm{GL}(k,p)$, the two modules are not necessarily isomorphic to each other. Recall that two modules in $\mathcal{M}$, associated with representations ξ and ξ', are isomorphic as $\mathbb{F}_p G$-modules if and only if there exists an element of $\mathrm{GL}(k,p)$ that conjugates g^ξ to $g^{\xi'}$ for each $g \in G$.

I am indebted to W. Nickel and Dr. C.R. Leedham-Green for the following theorem.

Theorem 4.3 An automorphism α of G lifts to H if and only if M and $M^{\alpha^{-1}}$ are isomorphic as $\mathbb{F}_p G$-modules.

Proof Suppose that M and $M^{\alpha^{-1}}$ are isomorphic. Let ν be an element of $\mathrm{GL}(k,p)$ satisfying $\nu^{-1} g^\xi \nu = g^{\alpha\xi}$ for all $g \in G$. Define the map γ on H by $(g^\tau m)^\gamma = g^{\alpha\tau} m\nu$. This map is well defined, since every element of H has a unique expression as a product $g^\tau m$ with $g \in G$ and $m \in M$. Furthermore, because τ is a group homomorphism, γ is an endomorphism of H, since

$$\begin{aligned} g^{\alpha\tau}\, m\nu\, h^{\alpha\tau}\, n\nu &= g^{\alpha\tau}\, h^{\alpha\tau}\, (m\nu)^{h^{\alpha}}\, n\nu \\ &= (g^{\alpha} h^{\alpha})^{\tau}\, m(\nu h^{\alpha\xi})\, n\nu \\ &= (gh)^{\alpha\tau}\, m(h^{\xi}\nu)\, n\nu \\ &= (gh)^{\alpha\tau}\, (m^{h} n)\nu. \end{aligned}$$

The restriction of γ to M is ν, and γ induces α on G. Therefore γ is an automorphism of H, and α lifts to γ.

Now suppose γ is a lifting of α to H. Denote by ν the restriction of γ to M. For $g \in G$, $g^{\tau\gamma}$ is some element of the coset $g^{\alpha\tau}M$, say $g^{\alpha\tau}m_g$ with $m_g \in M$. Consider

$$g^{\alpha\tau}m_g\, m^g\nu = (g^{\tau}m^{g})^{\gamma} = (mg^{\tau})^{\gamma} = m\nu\, g^{\alpha\tau}m_g = g^{\alpha\tau}m_g(m\nu)^{g^{\alpha}},$$

which implies that $mg^{\xi}\nu = m\nu g^{\alpha\xi}$ for all $g \in G$ and $m \in M$. Hence $g^{\xi}\nu = \nu g^{\alpha\xi}$, and so M and $M^{\alpha^{-1}}$ are isomorphic. ∎

4.3 Compatible pairs

Recall from Chapter 3 the homomorphism ϑ mapping $\operatorname{Aut}\mathcal{E}$ into the direct product $\operatorname{Aut} G \times \operatorname{Aut} M$. We shall call a pair (α, ν) of automorphisms in $\operatorname{Aut} G \times \operatorname{Aut} M$ a *compatible pair* (Robinson 1981) if

$$\nu g^{\alpha\xi} = g^{\xi}\nu \tag{4.1}$$

for all $g \in G$. This is precisely the condition that ν is an $\mathbb{F}_p G$-module isomorphism from M to $M^{\alpha^{-1}}$. The following is a direct consequence of Theorem 4.3.

Corollary 4.4 The group of compatible pairs is precisely the image of ϑ.

Since the compatible pairs coincide with the image of ϑ the introduction of the term "compatible pair" may appear redundant. In the non-split case, however, the situation is more complicated and the two no longer coincide.

Let $C(\mathcal{E}) \leq \operatorname{Aut} G \times \operatorname{Aut} M$ be the group of compatible pairs for the extension $\mathcal{E}$. Let π be the restriction to $C(\mathcal{E})$ of the projection map from $\operatorname{Aut} G \times \operatorname{Aut} M$ onto $\operatorname{Aut} G$. Then we have a natural composition of homomorphisms,

$$\operatorname{Aut} \mathcal{E} \xrightarrow{\vartheta} C(\mathcal{E}) \xrightarrow{\pi} \operatorname{Aut} G.$$

The image of π is the subgroup of $\operatorname{Aut} G$ consisting of all the automorphisms of G that lift to H. Therefore, if $\{\beta_1, \ldots, \beta_t\}$ is a generating set for this image, and $B = \{(\beta_1, \nu_1), \ldots, (\beta_t, \nu_t)\}$ is a preimage of the generating set under π, then $B \cup \operatorname{Ker} \pi$ generates the group of compatible pairs. We begin by computing a generating set for the image of π.

We assume that we have an effective algorithm for deciding when two modules are isomorphic, and that an explicit isomorphism is returned when the answer is affirmative (such an algorithm is described in Section 4.4). Theorem 4.3 provides a test for determining whether an automorphism α of G lifts to H. Since we have the generating set A for the automorphism group of G, we can perform an orbit-stabiliser calculation to find a generating set for $\operatorname{Im} \pi$.

Define an equivalence relation on $\mathcal{M}$ such that N_1 is equivalent to N_2 if and only if N_1 and N_2 are isomorphic as $\mathbb{F}_p G$-modules. Let $\mathcal{M}'$ be the set of equivalence classes in $\mathcal{M}$, and denote its elements by $[M_i]$ for $i \in \{1, \ldots, s\}$ and $M_i \in \mathcal{M}$. Let $M_1 = M$.

Lemma 4.5 The action of $\operatorname{Aut} G$ on $\mathcal{M}$ induces a permutation representation of $\operatorname{Aut} G$ on $\mathcal{M}'$.

Proof Let $\alpha \in \operatorname{Aut} G$, and $i \in \{1, \ldots, s\}$. If $N \in [M_i]$, then $N \cong M_i$, and therefore there exists $\nu \in \operatorname{GL}(k, p)$ such that $g^\eta = \nu^{-1} g^{\eta_i} \nu$ for all $g \in G$, where η and η_i are the representations associated with N and M_i respectively. In particular, $(g^{\alpha^{-1}})^\eta = \nu^{-1} (g^{\alpha^{-1}})^{\eta_i} \nu$ for all $g \in G$, which implies that $N^\alpha \cong M_i^\alpha$. Therefore we may define $[M_i]^\alpha$ to be $[M_i^\alpha]$. Clearly $[M_i]^{\alpha\beta} = [M_i^\alpha]^\beta$ and $[M_i]^\iota = [M_i]$ for all $i \in \{1, \ldots, s\}$ and $\alpha, \beta \in \operatorname{Aut} G$. ■

Since we are assuming the existence of an algorithm for determining module isomorphism, we can decide when $[M_i]^\alpha = [M_j]$ for all $i, j \in \{1, \ldots, s\}$ and $\alpha \in \operatorname{Aut} G$. Since we have a generating set A for $\operatorname{Aut} G$, we can calculate the stabiliser of $[M]$ under $\operatorname{Aut} G$ using the orbit-stabiliser algorithm. Let $\{\kappa_1, \ldots, \kappa_r\}$ be the set of Schreier generators returned for the stabiliser of $[M]$. For each $j \in \{1, \ldots, r\}$ let

ν_j be an $\mathbb{F}_p G$-module isomorphism from M to $M^{\kappa_j^{-1}}$. Then (κ_j, ν_j) is a compatible pair for $\mathcal{E}$. The set $\{(\kappa_1, \nu_1), \ldots, (\kappa_r, \nu_r)\}$ generates a subgroup of the group of compatible pairs $C(\mathcal{E})$, which may be a proper subgroup. The set must be supplemented with generators of the kernel of $\pi : \operatorname{Aut} G \times \operatorname{Aut} M \to \operatorname{Aut} G$ in order to obtain a generating set for $C(\mathcal{E})$. Assume for the moment that we have a generating set for the kernel, say $\{(\kappa_{r+1}, \nu_{r+1}), \ldots, (\kappa_{r+l}, \nu_{r+l})\}$. Then

$$\{\gamma_i \mid \gamma_i \colon g^\tau m \longmapsto g^{\kappa_i \tau} m\nu_i,\ 1 \leq i \leq r+l\}$$

is a set of automorphisms of H that generates a supplement for $\operatorname{Ker} \vartheta$ in $\operatorname{Aut} \mathcal{E}$.

Assuming that we have computed a generating set for the kernel of the projection π, then we have computed the automorphism group of the extension in the split case. In Section 4.6 we compute a generating set for $\operatorname{Ker} \pi$. This calculation relies on an algorithm for determining a decomposition of M into a sum of indecomposable submodules. Such an algorithm is described as part of the module isomorphism algorithm.

4.4 Determining module isomorphism

Many of the results stated and used here are standard results about modules and endomorphism algebras, and can be found, for example, in Feit (1982) and Pierce (1982). The algorithms for finding decompositions and for deciding indecomposability are due to Schneider (1990). An $\mathbb{F}_p$-basis for the endomorphism algebra of a module, or for the homomorphism group between two modules is a prerequisite for many of these algorithms, and we begin by describing a method for computing such bases.

4.4.1 Computing module homomorphisms

Let M and N be two $\mathbb{F}_p G$-modules of dimensions k and l respectively. Let ξ and η be the associated representations. We wish to construct an $\mathbb{F}_p$-basis for $\mathrm{Hom}_G(M, N)$.

An element X of $\mathrm{Hom}_G(M, N)$ corresponds to a $k \times l$ matrix (x_{ij}) with entries in $\mathbb{F}_p$ satisfying

$$g^\xi X = X g^\eta \tag{4.2}$$

for all $g \in G$. Let $(g^\xi)_{ij}$ denote the (i, j)-th entry of g^ξ, and similarly for $(g^\eta)_{ij}$. Then (4.2) is equivalent to the equations

$$\sum_{r=1}^{k} (g^\xi)_{ir} x_{rj} - \sum_{s=1}^{l} x_{is} (g^\eta)_{sj} = 0 \tag{4.3}$$

for all $i \in \{1, \ldots, k\}$ and $j \in \{1, \ldots, l\}$. It suffices to check that X satisfies this equation for all g in a generating set of G. Therefore, as g_i runs over the elements of $\mathcal{A}_G$, the equations (4.3) form a homogeneous system of linear equations in the x_{ij} such that an $\mathbb{F}_p$-linear basis of the solution space of the system corresponds to an $\mathbb{F}_p$-linear basis for $\mathrm{Hom}_G(M, N)$. The number of indeterminates is kl and the number of equations is nkl. An implementation of this simple algorithm for computing module homomorphism group bases is quite effective for modules of dimension less than 30.

When $M = N$ this algorithm computes a basis for the endomorphism algebra of M. When the dimension of M is large, the calculations described above become quite time consuming and better techniques are required for an effective implementation. We consider two cases. First suppose that the order of G is coprime to p. Note that when G is given by a consistent power conjugate presentation, we can easily identify whether the order of G is coprime to p. In this case the representation is completely reducible, and therefore every submodule of M is also a direct summand of M. The MEAT-AXE algorithm (Parker 1984) can be applied to compute a proper submodule of M. It is an effective algorithm which works on the principal that a null vector of an element of the group algebra that has nullity 1 is likely to lie a proper invariant subspace. Using MEAT-AXE we compute a submodule M_1 of M, and write $M = M_1 \oplus M_2$. Then

$$\mathrm{End}_G(M) \cong \mathrm{End}_G(M_1) \oplus \mathrm{End}_G(M_2) \oplus \mathrm{Hom}_G(M_1, M_2) \oplus \mathrm{Hom}_G(M_2, M_1).$$

Therefore to calculate $\mathrm{End}_G(M)$ we calculate each of the summands on the right-hand side separately, which involves significantly less work than computing the left-hand side directly. By recursively applying MEAT-AXE to the summands M_1 and M_2, we eventually get a decomposition of M into irreducible summands.

If the order of G is not coprime to p, then there exists an element of p-power order in G. Note that when G is given by a consistent power conjugate presentation, such an element is easily identified. Schneider (1990) describes an algorithm for computing the endomorphism algebra of modules in the case where a generator of a nontrivial cyclic p-subgroup P of G is given. A reduction in the number of unknowns that need to be considered is obtained by representing an arbitrary element of $\mathrm{End}_G(M)$ by an $\mathbb{F}_p$-linear combination of the basis elements for $\mathrm{End}_P(M)$. The basis elements for $\mathrm{End}_P(M)$ are easily constructed, and implementations of this algorithm, which are widely available as intrinsic functions of both CAYLEY and MAGMA, perform well on module dimensions as large as several hundred.

We have established algorithms with the following descriptions:

Algorithm 4.6 ENDO

Input: An $\mathbb{F}_p G$-module M.

Output: An $\mathbb{F}_p$-basis for the endomorphism algebra of M.

Algorithm 4.7 HOM

Input: Two $\mathbb{F}_p G$-modules M and N.

Output: An $\mathbb{F}_p$-basis for the vector space of homomorphisms from M to N.

4.4.2 Indecomposable modules

Let M and N be two $\mathbb{F}_p G$-modules associated with representations ξ and η mapping G into $\mathrm{GL}(k, p)$. Verifying that the two modules are isomorphic can be as simple as computing a basis for $\mathrm{Hom}_G(M, N)$ and taking a few random elements in search of an isomorphism. Consider the case of $M \cong N$ and $M = M_1 \oplus \ldots \oplus M_t$ with the M_i irreducible and pairwise non-isomorphic. The probability of choosing an isomorphism in $\mathrm{Hom}_G(M, N)$ is at least $(1 - \frac{1}{p})^t$. Particularly when the number of irreducible summands is low this probability is quite high, and a limited random search for an isomorphism is likely to be successful. For example, if the prime is 3 and there are 4 irreducible summands, then with probability 0.9 a set of 10 randomly chosen elements of $\mathrm{Hom}_G(M, N)$ will contain at least one isomorphism.

However, when it is not known in advance whether or not the two modules are isomorphic, the failure to produce an isomorphism after such a random search does not imply that the two modules are not isomorphic. We require a deterministic algorithm for deciding the issue, and moreover one which admits an effective implementation.

For indecomposable modules the situation is much simpler. To see this we need to examine the structure of the endomorphism algebra of an indecomposable module. The modules we are considering are all finite dimensional over a finite field, so their endomorphism algebras are finite dimensional algebras over the same finite field. The following results are well known (see for example Pierce 1982). Let A be a finite dimensional algebra over a field. Denote by $\mathrm{Rad}(A)$ the radical of A, which is a nilpotent ideal in A.

Definition 4.8 A finite dimensional algebra A is a *local algebra* if $A/\mathrm{Rad}(A)$ is a field. If A is a local algebra, then $A/\mathrm{Rad}(A)$ is called the *residue field* of A.

Lemma 4.9 For an algebra A, the following are equivalent:

1) A is a local algebra.
2) $\mathrm{Rad}(A)$ is the set of all non-units of A.
3) The unique maximal left ideal of A is $\mathrm{Rad}(A)$.

Theorem 4.10 If an $\mathbb{F}_p G$-module M is indecomposable, then $\mathrm{End}_G(M)$ is a local algebra.

This last theorem is the major step towards an effective module isomorphism algorithm, since as the following lemma shows, it provides an effective deterministic algorithm for deciding whether two indecomposable modules are isomorphic. It simply requires a basis for $\mathrm{Hom}_G(M, N)$.

Lemma 4.11 Let M and N be indecomposable $\mathbb{F}_p G$-modules, and $\{f_1, \ldots, f_t\}$ an $\mathbb{F}_p$-basis for $\mathrm{Hom}_G(M, N)$. Then $M \cong N$ if and only if there exists $i \in \{1, \ldots, t\}$ such that f_i is an isomorphism.

Proof We need only prove that $M \cong N$ implies that one of the f_i is an isomorphism. Let $\alpha \colon N \to M$ be an $\mathbb{F}_p G$-module isomorphism. For each i, $f_i\alpha$ is an endomorphism of M, and so the $\mathbb{F}_p$-span of $\{f_1\alpha, \ldots, f_t\alpha\}$ is a subgroup of $\mathrm{End}_G(M)$. Take $\beta \in \mathrm{End}_G(M)$. Consider $\beta\alpha^{-1} \colon M \to N$. This is an $\mathbb{F}_p G$-module homomorphism, and hence there exist coefficients $a_i \in \mathbb{F}_p$ such that $\sum_{i=1}^{t} a_i f_i = \beta\alpha^{-1}$, whence $\sum a_i(f_i\alpha) = \beta$. Therefore $\{f_1\alpha, \ldots, f_t\alpha\}$ spans $\mathrm{End}_G(M)$. But $\mathrm{End}_G(M)$ is a local algebra (by Theorem 4.10), so the singular elements form an ideal in $\mathrm{End}_G(M)$ equal to $\mathrm{Rad}(\mathrm{End}_G(M))$ (by Lemma 4.9). Since there exists at least one non-singular endomorphism of M, at least one of the $f_i\alpha$ must be non-singular. ■

As a consequence, if M and N are both indecomposable, then we need only calculate the rank of each element of a basis for $\mathrm{Hom}_G(M, N)$ to check whether the two modules are isomorphic. Therefore we have an algorithm with the following description:

Algorithm 4.12 INDEC-ISOMORPHIC

Input: Two indecomposable $\mathbb{F}_p G$-modules M and N.

Output: An isomorphism $M \to N$ if $M \cong N$, or "False" otherwise.

Since we are working with finite dimensional modules, we have at our disposal the fundamental Krull-Schmidt Theorem.

Theorem 4.13 (Krull-Schmidt) Let A be a finite dimensional algebra over a field, and $V \neq 0$ a finite dimensional A-module. Then:

1) V is a finite direct sum of indecomposable modules.
2) If $V = \bigoplus_{i=1}^{m} V_i = \bigoplus_{j=1}^{n} W_j$, where each V_i, W_j is indecomposable, then $n = m$ and, after possibly reordering the W_j, $V_i \cong W_i$ for all $i \in \{1, \ldots, m\}$.

An immediate corollary of this theorem suits our purpose.

Corollary 4.14 Let $M = \bigoplus_{i=1}^{m} M_i$ and $N = \bigoplus_{j=1}^{n} N_j$ be $\mathbb{F}_p G$-modules with M_i, N_j indecomposable. Then $M \cong N$ if and only if $n = m$ and, after possibly reordering the N_j, $M_i \cong N_i$ for all $i \in \{1, \ldots, m\}$.

Assuming that a full decomposition of each module can be calculated, Corollary 4.14 says that if we find a bijection between isomorphic summands of each module, then the two modules are isomorphic. The isomorphism test for the summands is solved by Lemma 4.11. We now consider the problem of calculating such decompositions.

4.5 Finding module decompositions

An element e of $\mathrm{End}_G(M)$ is called an *idempotent* if $e^2 = e$. Associated with a decomposition $M_1 \oplus M_2$ of a module M are two idempotents e and $1-e$ in $\mathrm{End}_G(M)$, where e is the projection of M onto M_1 relative to $M_1 \oplus M_2$. Conversely, an arbitrary idempotent e leads to the decomposition $M = Me \oplus M(1-e)$. A simple algorithm for computing a decomposition of M searches the endomorphism algebra of M for idempotents. However, a larger class of elements of $\mathrm{End}_G(M)$ which also lead to decompositions may be used to increase the probability of success in a (random) search. These are the so-called Fitting elements in $\mathrm{End}_G(M)$ (Schneider 1990).

Definition 4.15 An endomorphism f of M is a *Fitting element* if it is singular but not nilpotent.

Clearly idempotents are also Fitting elements. An endomorphism f is a Fitting element if and only if there exists some integer $n \geq 1$ such that $0 < \mathrm{rank}(f^n) = \mathrm{rank}(f^{n+1}) < k$, and as a consequence, $\mathrm{rank}(f^n) = \mathrm{rank}(f^m)$ for all $m \geq n$. The fact that they lead to a direct decomposition of the module follows from Fitting's Lemma.

Lemma 4.16 (Fitting's Lemma) Let n be the length of a composition series of an A-module V, and $f \in \mathrm{End}_A(V)$. Then $V = \mathrm{Im}\, f^n \oplus \mathrm{Ker}\, f^n$.

Fitting's Lemma also demonstrates that we need only examine powers of an endomorphism up to the composition length of the module in order to check whether it is a Fitting element. How common are Fitting elements in $\mathrm{End}_G(M)$? Schneider gives a simple analysis of the case $M = M_1 \oplus M_2$, with both summands absolutely simple; that is, all submodules are irreducible. When M_1 and M_2 are not isomorphic, the number of endomorphisms is p^2, of which 2 are idempotents and $2(p-1)$ are Fitting elements. When M_1 and M_2 are isomorphic, there are

p^4 endomorphisms, of which $p(p+1)$ are idempotents and $p(p+1)(p-1)$ are Fitting elements. When $p > 2$ there are at least twice as many Fitting elements as idempotents irrespective of whether or not M_1 and M_2 are isomorphic. For larger p it is clear that the factor rapidly becomes much larger than two. More generally, the factor depends on the structure of the summands, but for many modules, even of characteristic 2, it will be much larger than 2.

4.5.1 Verifying indecomposability

When a Fitting element is not found after a number of random samples of the endomorphism algebra, indecomposability is only one of two possible explanations. We require a deterministic algorithm that decides whether the module is indecomposable. The endomorphism algebra of an indecomposable module has some special properties that may be exploited as the basis of such an algorithm. It is still worthwhile doing a brief random search for a Fitting element, since if one is found it avoids the overhead associated with the deterministic algorithm.

Following Lemma 4.9, to verify that a module is indecomposable it suffices to check that the singular elements of $\mathrm{End}_G(M)$ are all nilpotent. Alternatively, if a basis of nilpotent elements is found, such that the ideal generated by them is nilpotent and $\mathrm{End}_G(M)$ modulo the ideal is a field, then the ideal is the radical of $\mathrm{End}_G(M)$ and M is indecomposable. A detailed description of a suitable algorithm appears in Schneider (1990). Basically, the algorithm begins by constructing an approximation, J_i, to the radical of $\mathrm{End}_G(M)$. We may take J_0 to be $\{0\}$. For each approximation J_i, if an attempt to verify that $\mathrm{End}_G(M)/J_i$ is a field succeeds, then the module is indecomposable. If $\mathrm{End}_G(M)/J_i$ is a field, say $\mathbb{F}_{p^r}$ for some $r > 0$, then there exists a nonsingular element in $\mathrm{End}_G(M) - J_i$ which has multiplicative order a multiple of $p^r - 1$. If the field verification step fails, it is for one of two reasons:

1) The discovery of a nilpotent element, f, lying outside of J_i. In this case the algorithm calculates the ideal J_{i+1} generated by f and J_i, and restarts the field verification step with this better approximation to the radical.
2) The discovery of a Fitting element, f, in which case the module is not indecomposable and a decomposition is given by the image and kernel of f^n for some n.

This procedure is guaranteed to terminate after a finite number of steps since the ascending series of approximations J_i to $\mathrm{Rad}(\mathrm{End}_G(M))$ must eventually become constant, and we can ensure that the field checking step eventually checks enough elements of $\mathrm{End}_G(M)/J_i$; note, moreover, that the proportion of primitive elements is high so that this search is not time consuming.

The existence of an algorithm for computing a direct decomposition of a module with indecomposable summands is a consequence of the preceding discussion. It takes as input a representation $\xi\colon G \to \mathrm{GL}(k,p)$ defining an $\mathbb{F}_pG$-module M, and a basis for the endomorphism algebra of M. It consists of the following basic step. Search for a Fitting element of $\mathrm{End}_G(M)$. If no such element is found, apply an indecomposability test, leading either to the conclusion that M is indecomposable, or to a Fitting element in $\mathrm{End}_G(M)$. When a Fitting element f has been found, it corresponds to a decomposition $M_1 \oplus M_2$ of M. Bases for the endomorphism algebras of M_1 and M_2 are computed and the algorithm recurses on these summands.

Algorithm 4.17 MODULE-DECOMPOSE

Input: An $\mathbb{F}_pG$-module M and an $\mathbb{F}_p$-basis for $\mathrm{End}_G(M)$.

Output: A list $M_1, \ldots, M_t$ of indecomposable submodules of M such that $M = \bigoplus_{i=1}^{t} M_i$. In addition, for each M_i, a basis for $\mathrm{Rad}(\mathrm{End}_G(M_i))$ and a non-singular endomorphism w_i that generates the residue field of $\mathrm{End}_G(M_i)$.

Note that only the endomorphism algebra of M needs to be computed by the linear equations method described in Section 4.4.1. The endomorphism algebras of submodules required for the recursive step may be obtained directly from the endomorphism algebra of M as the following result indicates.

Lemma 4.18 Let $\{f_1, \ldots, f_t\}$ be an $\mathbb{F}_p$-basis for $\mathrm{End}_G(M)$, and let $M = M_1 \oplus M_2$. If $\pi\colon M \to M_1$ is the projection map onto M_1 relative to $M_1 \oplus M_2$, then $\{f_1|_{M_1}\pi, \ldots, f_t|_{M_1}\pi\}$ contains an $\mathbb{F}_p$-basis for $\mathrm{End}_G(M_1)$.

Proof The map $f \mapsto f|_{M_1}\pi$ is clearly onto, since each $f_1 \in \mathrm{End}_G(M_1)$ has $f_1 \oplus 0$ as a preimage. ■

4.5.2 Exploiting the decomposition

The algorithm that computes a decomposition of a module into indecomposable summands was motivated by the question of deciding when two modules are isomorphic. However, in the context of lifting automorphisms, the decomposition algorithm plays a more important role. Since we are typically interested in deciding whether M^α is isomorphic to M^β for $\alpha, \beta \in \operatorname{Aut} G$, we can use the fact that the action of an automorphism of G preserves the decomposition of M, and hence avoid computing decompositions of modules obtained during the orbit-stabiliser calculation.

Lemma 4.19 Let $M = \bigoplus_{i=1}^{t} M_i$. If $\alpha \in \operatorname{Aut} G$, then $M^\alpha = \bigoplus_{i=1}^{t} M_i^\alpha$.

Proof Let $\xi : G \to \mathrm{GL}(k, p)$ be the representation associated with M, and let $\alpha \in \operatorname{Aut} G$. Let $M = M_1 \oplus M_2$, and $\xi = \xi_1 \oplus \xi_2$ the associated decomposition of ξ, where ξ_i is the representation associated with module M_i. Hence $\alpha\xi = \alpha\xi_1 \oplus \alpha\xi_2$ and M^α has a decomposition into two submodules associated with representations $\alpha\xi_1$ and $\alpha\xi_2$; namely M_1^α and M_2^α respectively. The result follows by induction on the number of summands. ∎

Therefore, before commencing the orbit-stabiliser calculation in Section 4.2, we calculate a decomposition of M into indecomposable summands. When a new module M^α is generated by the orbit-stabiliser algorithm its decomposition is obtained by applying α to each of the submodules in the decomposition of M.

4.6 Module automorphisms

Recall the composition of mappings,

$$\operatorname{Aut} \mathcal{E} \xrightarrow{\vartheta} \operatorname{Aut} G \times \operatorname{Aut} M \xrightarrow{\pi} \operatorname{Aut} G.$$

In Section 4.2 we computed a generating set $\{\beta_1, \ldots, \beta_t\}$ for the image of $\vartheta\pi$, and obtained from this a set B of elements in $C(\mathcal{E})$ generating a subgroup of $C(\mathcal{E})$. This subgroup is a supplement in $\operatorname{Im} \vartheta$ of the kernel of

$$\pi\colon \operatorname{Im} \vartheta \longrightarrow \operatorname{Aut} G.$$

The elements of this kernel correspond to liftings of the identity which, apart from the trivial lifting, act non-trivially on M. Such automorphisms correspond to $\mathbb{F}_p G$-module automorphisms of M, since the compatibility condition for a pair (ι, ν) is

$$g^{\xi}\nu = \nu g^{\xi}, \tag{4.4}$$

which is precisely the condition for the group automorphism ν of M to be a module endomorphism. We denote by $\mathrm{Aut}_G(M)$ the group of $\mathbb{F}_p G$-module automorphisms of M, which is simply the multiplicative group of invertible elements in $\mathrm{End}_G(M)$. Therefore, a generating set for the group of compatible pairs may be obtained from B and a generating set of the module automorphism group of M.

It is clear from (4.4) that the automorphisms ν of M that may appear in pairs (ι, ν) are precisely the elements of the centraliser of G^{ξ} in $\mathrm{GL}(k,p)$. Matrix group algorithms have not yet been developed to a stage that would allow computing centralisers of matrix groups effectively. A technique used in some computer algebra packages is to compute a base and strong generating set for a permutation representation of the matrix group, and then apply standard permutation group algorithms for finding the centraliser. However, even for modules of relatively small dimension this calculation can be time consuming. An alternative method that makes use of the decomposition of the module is described next. This method uses information already obtained when computing the decomposition of the module; namely, a basis for the radical of the endomorphism algebra of each indecomposable summand, and the module automorphism that was found to verify indecomposability of the summand.

4.6.1 Homogeneous modules

We first consider the case of $N = \bigoplus_{i=1}^{n} M_i$, with $M_i \cong M$ for all $i \in \{1, \ldots, n\}$, where M is an indecomposable $\mathbb{F}_p G$-module. By Krull-Schmidt, every indecomposable direct summand of N is isomorphic to M, and we shall call such a module *homogeneous*. Note that this is a slight departure from the usual definition which is for direct sums of semisimple modules. We shall use I for the identity element of $\mathrm{Aut}_G(N)$, and 1 for the identity element of $\mathrm{Aut}_G(M)$. Let J denote the radical of $\mathrm{End}_G(M)$.

Lemma 4.20 (see for example Pierce 1982, Corollary 3.4a) $\mathrm{End}_G(N)$ is isomorphic to the full matrix algebra $\mathbb{M}(n, \mathrm{End}_G(M))$.

In the light of this lemma, we shall identify $\mathrm{End}_G(N)$ with the set of $n \times n$ matrices over $\mathrm{End}_G(M)$.

Let the order of $\mathrm{End}_G(M)$ be p^r, and the order of J be p^s. By Theorem 4.10, $\mathrm{End}_G(M)$ is a local algebra, and therefore $F = \mathrm{End}_G(M)/J$ is the residue field of order $q = p^{r-s}$. By Lemma 4.9, the radical of $\mathrm{End}_G(M)$ is the set of all singular elements in $\mathrm{End}_G(M)$, so $\mathrm{Aut}_G(M) = \mathrm{End}_G(M) - J$ and its order is precisely $p^s(q-1)$.

Theorem 4.21 (Wedderburn Principal Theorem) (see for example Pierce 1982, Corollary 11.6) If F is a perfect field and A is a finite dimensional F-algebra, then there is a subalgebra B of A such that $A = B \oplus \mathrm{Rad}(A)$.

The group $\mathrm{Aut}_G(M)$ has a normal subgroup $P = \{1 + y \mid y \in J\}$ of order p^s. By the Wedderburn Principal Theorem, there exists a subalgebra B of $\mathrm{End}_G(M)$ that is isomorphic to the residue field F. Since $B \cap J = \{0\}$ the multiplicative group $B^\times$ of nonzero elements of B is a subgroup of $\mathrm{Aut}_G(M)$, and since its order is $q-1$ it is a complement for P in $\mathrm{Aut}_G(M)$. Let w be an element of $\mathrm{Aut}_G(M)$ such that $w + J$ is a primitive element of F. Hence $q-1$ divides the order of w, and since p and $q-1$ are coprime, we may choose w such that it has order exactly $q-1$. The cyclic group W generated by w has order $q-1$ and is a complement of P in $\mathrm{Aut}_G(M)$. By the Schur–Zassenhaus Theorem, W is conjugate to $B^\times$, and therefore $W \cup \{0\}$ is a subalgebra of $\mathrm{End}_G(M)$ isomorphic to the residue field F, with primitive element w. We identify F with $W \cup \{0\}$.

The general linear group $\mathrm{GL}(n, F)$ has a 2 element generating set, where each element is a matrix whose non-zero entries are explicitly given as powers of a primitive element of F (Taylor 1987). Let $U(w)$ and $V(w)$ be two such matrices and denote by K the subgroup of $\mathrm{Aut}_G(N)$ that they generate. We have a 2 element generating set for a subgroup of $\mathrm{Aut}_G(N)$, and our aim is to find a larger set that generates all of $\mathrm{Aut}_G(N)$. We begin by proving an elementary but important lemma.

Lemma 4.22 Let $X, Y \in \mathbb{M}(n, J)$, and $C \in \mathbb{M}(n, \mathrm{End}_G(M))$. Then X is nilpotent, and $X + Y$, CY and YC are all elements of $\mathbb{M}(n, J)$.

Proof If m is a positive integer, then each entry of X^m is an $\mathbb{F}_p$-linear combination of products of length m in the entries of X. Hence X^m lies in $\mathbb{M}(n, J^m)$. Since J is a nilpotent ideal, there exists a positive integer m such that $J^m = 0$. Hence $X^m = 0$. Since J is closed under addition, $X+Y \in \mathbb{M}(n, J)$. The entries of YC and CY are $\mathbb{F}_p$-linear combinations of terms $C_{ik}Y_{kj}$ and $Y_{ik}C_{kj}$ respectively. Since the Y_{ij} all lie in the ideal J, both CY and YC are elements of $\mathbb{M}(n, J)$. ■

Define a subset L of $\mathbb{M}(n, \mathrm{End}_G(M))$ by

$$L = \{I + Y \mid Y \in \mathbb{M}(n, J)\}.$$

If $I+Y$ is an element of L, then Y is a nilpotent matrix by Lemma 4.22 and $Y^m = 0$ for some integer m. Therefore $(I+Y)(I-Y+Y^2-Y^3+\cdots+(-Y)^m) = I$ implies that $I+Y$ is an invertible matrix, and hence $L \subseteq \mathrm{Aut}_G(N)$. Moreover the inverse of $I+Y$ is also an element of L. If $I+X$ and $I+Y$ are elements of L, then $(I+X)(I+Y) = I + (X+Y+XY) \in L$. Hence L is a subgroup of $\mathrm{Aut}_G(N)$. Let C be an invertible element of $\mathbb{M}(n, \mathrm{End}_G(M))$. For $I + Y \in L$,

$$C^{-1}(I+Y)C = I + C^{-1}YC.$$

Lemma 4.22 implies $C^{-1}YC \in \mathbb{M}(n, J)$, and so $I + C^{-1}YC \in L$. Therefore L is a normal subgroup of $\mathrm{Aut}_G(N)$.

Definition 4.23 Let C be an endomorphism of N. Define $\bar{C}$ by $\bar{C}_{ij} = b$ where $C_{ij} = b + y$ for $b \in F$ and $y \in J$.

Observe that if C is an element of $\mathrm{End}_G(N)$, then $\bar{C}$ is in K, and $C - \bar{C}$ is in $\mathbb{M}(n, J)$.

Lemma 4.24 $KL = \mathrm{Aut}_G(N)$.

Proof For $C \in \mathrm{Aut}_G(N)$, let $C - \bar{C} = Y \in \mathbb{M}(n, J)$. Then $C = \bar{C} + Y = \bar{C}(I + \bar{C}^{-1}Y)$ where $\bar{C} \in K$ and, by Lemma 4.22, $I + \bar{C}^{-1}Y \in L$. ■

Lemma 4.25 Let $\mathcal{B} = \{b_1, \ldots, b_t\}$ be a basis for J such that $\mathcal{B}_i = \mathcal{B} \cap J^i$ is a basis for J^i. Then $\mathcal{D} = \{1 + b_j \mid b_j \in \mathcal{B}\}$ is a generating set for P.

Proof Let y be an element of J. We have to show that $1+y$ is in the subgroup generated by $\mathcal{D}$. We prove this by induction. Let $y_1 = y$. Clearly $1 = 1 + y - y_1$. Assume that $\prod_{j=1}^{s} x_j = 1 + y - y_i$ with $x_1, \ldots, x_s \in \mathcal{D}$ and $y_i \in J^i$ for some $i > 0$. Since $\mathcal{B}_i$ is a basis for J^i, there exist elements $f_i \in \mathbb{F}_p$ such that $y_i = \sum_{b_i \in \mathcal{B}_i} f_i b_i$. Therefore

$$\prod_{b_i \in \mathcal{B}_i} (1 + b_i)^{f_i} = 1 + y_i - y_{i+1}$$

for some $y_{i+1} \in J$. But y_{i+1} is a sum of products of length at least two in elements of $\mathcal{B}_i$, hence $y_{i+1} \in J^{i+1}$. Consider

$$\begin{aligned} \prod_{j=1}^{s} x_j \times \prod_{b_i \in \mathcal{B}_i} (1 + b_i)^{f_i} &= (1 + y - y_i)(1 + y_i - y_{i+1}) \\ &= 1 + y - y_{i+1} + y(y_i - y_{i+1}) - y_i(y_i - y_{i+1}). \end{aligned}$$

Clearly $y(y_i - y_{i+1}) - y_i(y_i - y_{i+1})$ lies in J^{i+1}. The result follows by a finite induction on i, since J is nilpotent. ∎

It is worth noting that we cannot relax the conditions of Lemma 4.25 and use an arbitrary basis of J. For a nilpotent $\mathbb{F}_p$-algebra A, the *circle product* on A is defined by $a \circ b = a + b + ab$ for all $a, b \in A$, and A forms a group under this operation. Lemma 4.25 is equivalent to the statement that a particular type of $\mathbb{F}_p$-basis for a nilpotent algebra A also generates A as a circle group. An arbitrary basis for A will not necessarily generate it as a circle group. The following counterexample is due to Dr. L.G. Kovács (personal communication). Let A be the nilpotent $\mathbb{F}_2$-algebra generated by a single element a and the single relation $a^4 = 0$. Consider the circle subgroup generated by a. The set of non-zero elements of this proper subgroup is $\{a, a^2, a + a^2 + a^3\}$. While this set is an $\mathbb{F}_2$-basis for A, it does not generate A as a circle group.

Now we can write down a simple generating set for the $\mathbb{F}_p G$-module automorphism group of a homogeneous module. First define $z(b)$ for $b \in J$ to be the matrix

$$z(b) = \begin{pmatrix} 1+b & & & & \\ & 1 & & & \\ & & 1 & & \\ & & & \ddots & \\ & & & & 1 \end{pmatrix}$$

in $\mathbb{M}(n, \mathrm{End}_G(M))$.

Theorem 4.26 Let $\mathcal{B}$ be a basis for J such that $\mathcal{B} \cap J^i$ is a basis for J^i, and let $Z = \{z(b) \mid b \in \mathcal{B}\}$. Then $\{U(w), V(w)\} \cup Z$ is a generating set for $\mathrm{Aut}_G(N)$.

Proof Let T be the group generated by $\{U(w), V(w)\} \cup Z$. We shall exhibit a set of matrices in T that correspond to elementary row operation matrices, and then show that every element of $\mathrm{Aut}_G(N)$ can be reduced to an element that is obviously in T.

First observe that the pair of elements $U(w), V(w)$ generate the group K which contains

$$d_1(w^i) = \begin{pmatrix} w^i & & & & \\ & 1 & & & \\ & & 1 & & \\ & & & \ddots & \\ & & & & 1 \end{pmatrix}$$

for all $w^i \in W$, and also all the permutation matrices in $\mathbb{M}(n, \mathrm{End}_G(M))$. Define d_j similarly for each position on the main diagonal. Since $z(b_i) \in T$ for all $b_i \in \mathcal{B}$, Lemma 4.25 implies that $z(b) \in T$ for all $b \in J$. Therefore, for an arbitrary non-singular element a of $\mathrm{End}_G(M)$, we have $a = w^i + b$ for some $b \in J$, and hence $d_1(a) = d_1(w^i)z(w^{-i}b)$ is in T.

Let $y(c)$ denote the matrix

$$\begin{pmatrix} 1 & c & & & \\ & 1 & & & \\ & & 1 & & \\ & & & \ddots & \\ & & & & 1 \end{pmatrix}$$

for $c \in \mathrm{End}_G(M)$. If $b \in J$, then $1 + b$ is invertible, of order l say. Restricting our attention to the top right-hand corner, we have

$$\begin{aligned} y(-1)z(b)y(1)z(b)^{l-1} &= \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1+b & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1+b & 0 \\ 0 & 1 \end{pmatrix}^{l-1} \\ &= \begin{pmatrix} 1+b & b \\ 0 & 1 \end{pmatrix} \begin{pmatrix} (1+b)^{l-1} & 0 \\ 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} \\ &= y(b). \end{aligned}$$

Therefore, T contains the matrices $d_i(a)$ for all invertible $a \in \mathrm{End}_G(M)$ and $r_{ij}(b)$ for all $b \in \mathrm{End}_G(M)$, where $i, j \in \{1, \ldots, n\}$, $i \neq j$, and $r_{ij}(b)$ is the sum of the identity matrix and a matrix whose single non-zero entry is b lying in the (i,j)-th position. These matrices correspond to elementary row and column operation matrices, and we can use them to perform Gaussian elimination on an arbitrary matrix over $\mathrm{End}_G(M)$.

Let C be any invertible element of $\mathbb{M}(n, \mathrm{End}_G(M))$. Let $j \in \{1, \ldots, n\}$ be such that columns $1, \ldots, j-1$ of C are zero below the main diagonal. Let $C = \bar{C} + Y$ with $Y \in \mathbb{M}(n, J)$. By Lemma 4.22, Y is nilpotent, and $C^m = \bar{C}^m$ for a sufficiently large integer m. So $\bar{C}$ is invertible and column j of $\bar{C}$ must contain a non-zero entry. Hence column j of C contains an invertible entry and after possibly multiplying C by a permutation matrix, we may assume that C_{jj} is invertible. Therefore

$$\left(\prod_{i>j} r_{ij}(-C_{jj}^{-1} C_{ij})\right) C$$

is a matrix for which all entries below the main diagonal in columns $1, \ldots, j$ are zero. By induction, there exists a matrix E such that EC is upper triangular, with invertible elements on the main diagonal. Since E is a product of permutation matrices and $r_{ij}(b)$ for $b \in \mathrm{End}_G(M)$, it is an element of T. A similar argument using column operations on EC leads to ECD a diagonal matrix with invertible entries on the main diagonal, and $E, D \in T$. But then ECD certainly lies in T, which implies that $C \in T$. ■

4.6.2 An example of the homogeneous case

Let $G = \langle a \rangle$ be a cyclic group of order 6, and $p = 3$. Let ξ be the representation mapping G into $\mathrm{GL}(2, p)$ defined by

$$a^{\xi} = \begin{pmatrix} 2 & 2 \\ 0 & 2 \end{pmatrix}.$$

Let M be the associated $\mathbb{F}_p G$-module. The endomorphism algebra of M is computed by solving the homogeneous system of linear equations in variables x_{ij} arising from the matrix equation

$$\begin{pmatrix} 2 & 2 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \end{pmatrix} = \begin{pmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \end{pmatrix} \begin{pmatrix} 2 & 2 \\ 0 & 2 \end{pmatrix}.$$

A basis for the solution space of these equations leads to a basis

$$\left\{ b_1 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, b_2 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \right\}$$

for the endomorphism algebra of M. The element b_2 is nilpotent and obviously generates a 1-dimensional ideal J in $\mathrm{End}_G(M)$. The quotient $\mathrm{End}_G(M)/J$ is isomorphic to the scalar matrices and hence is a field. Therefore J is the radical of $\mathrm{End}_G(M)$ and M is indecomposable. The residue field $F = \mathrm{End}_G(M)/J$ has 3 elements and $w = 2b_1$ is a primitive element of the residue field embedded in $\mathrm{End}_G(M)$.

Let $N = M \oplus M \oplus M$, an $\mathbb{F}_p G$-module of dimension 6. A generating set for $\mathrm{GL}(3,3)$ is

$$\left\{ \begin{pmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 2 & 0 & 1 \\ 2 & 0 & 0 \\ 0 & 2 & 0 \end{pmatrix} \right\}$$

and following the recipe given in Theorem 4.26, a generating set for $\mathrm{Aut}_G(N)$ is the union of

$$\left\{ \begin{pmatrix} 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 2 & 0 & 0 & 0 & 1 & 0 \\ 0 & 2 & 0 & 0 & 0 & 1 \\ 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 & 0 \end{pmatrix} \right\}$$

and

$$\left\{ \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \right\}.$$

4.6.3 Inhomogeneous modules

Next we consider the general case,

$$N = \bigoplus_{i=1}^{t} N_i, \qquad \text{where } N_i = \bigoplus_{j=1}^{n_i} M_{ij}$$

with the M_{ij} indecomposable submodules of N such that $M_{ij} \cong M_{kl}$ if and only if $i = k$. The submodules N_i are the homogeneous components of N. Let $M_i = M_{i1}$ so that N_i is a direct sum of n_i copies of M_i. Let J_i be the radical of $\mathrm{End}_G(M_i)$. Note that a generating set for $\mathrm{Aut}_G(N_i)$ is given by Theorem 4.26. We assume that $t > 1$.

Choose a basis for N that is the union of bases of the submodules M_{ij}. With respect to this basis, an endomorphism A of N is equivalent to a matrix over $\mathbb{F}_p$ of the form

$$\begin{pmatrix} A_{11} & A_{12} & \dots & A_{1t} \\ A_{21} & A_{22} & \dots & A_{2t} \\ \vdots & \vdots & \ddots & \vdots \\ A_{t1} & A_{t2} & \dots & A_{tt} \end{pmatrix}$$

(see for example McDonald 1974, Theorem VII.1). The A_{ii} are blocks of size $n_i \times n_i$ with entries in $\mathrm{End}_G(M_i)$. For $i \neq j$, the A_{ij} are blocks of size $n_i \times n_j$ with entries in $\mathrm{Hom}_G(M_i, M_j)$.

Let K be the subgroup of $\mathrm{Aut}_G(N)$ defined by

$$K = \{A \in \mathrm{Aut}_G(N) \mid A_{ij} = 0, i \neq j\}.$$

Lemma 4.27 $K \cong \mathrm{Aut}_G(N_1) \times \cdots \times \mathrm{Aut}_G(N_t)$.

Proof This is clear from the block diagonal form of elements of K. ■

We identify $\mathrm{Aut}_G(N_i)$ with the corresponding subgroup of $\mathrm{Aut}_G(N)$ given in Lemma 4.27. We now show that an automorphism of N restricts to an automorphism on each homogeneous component N_i of N. For this purpose we need another standard result (see for example Karpilovsky 1992, Corollary 1.10.8).

Lemma 4.28 Let U, V, W and $M_1, \dots, M_n$ be R-modules such that

$$U = V \oplus W = \bigoplus_{i=1}^{n} M_i.$$

If V is non-zero and V, $M_1, \dots, M_n$ are indecomposable, then there exists $j \in \{1, \dots, n\}$ such that

$$U = W \oplus M_j = V \oplus \bigoplus_{\substack{i=1 \\ i \neq j}}^{n} M_i.$$

Lemma 4.29 If $A \in \mathrm{Aut}_G(N)$, then $A_{ii} \in \mathrm{Aut}_G(N_i)$.

Proof Let A be an automorphism of N, and let V be an indecomposable summand of N_1. Define $\pi_1 \colon N \to N_1$ to be the projection associated with $N = \bigoplus N_i$. It is enough to show that the image of V under $A\pi_1$ is isomorphic to V. Certainly $N = V \oplus W$ for some submodule W. Since A is an automorphism of N, we also have $N = VA \oplus WA$. By Lemma 4.28, there exists an $l \in \{1, \ldots, n_1\}$ such that

$$N = WA \oplus M_{1l} = VA \oplus \bigoplus_{i,j \neq 1,l} M_{ij}.$$

Therefore, if π_{1l} is the projection onto M_{1l} associated with $N_1 = \bigoplus M_{1j}$, then $VA\pi_{1l} = M_{1l}$. Hence $V \mapsto VA\pi_{1l}$ has trivial kernel, and so $V \mapsto VA\pi_1$ also has trivial kernel. ∎

Definition 4.30 Let $i, j \in \{1, \ldots, t\}$ with $i \neq j$. Let b be an element of $\mathrm{Hom}_G(M_i, M_j)$. Define the matrix $Y^{(ij)}(b)$ in $\mathrm{End}_G(N)$ by

$$Y^{(ij)}(b) = \begin{pmatrix} I_{11} & & & & & & \\ & \ddots & & & & & \\ & & I_{ii} & & & & \\ & & & \ddots & & & \\ & & H_{ij} & & I_{jj} & & \\ & & & & & \ddots & \\ & & & & & & I_{tt} \end{pmatrix}$$

where I_{ii} is the identity in $\mathrm{Aut}_G(N_i)$, and H_{ij} is an $n_i \times n_j$ block with entries in $\mathrm{Hom}_G(M_i, M_j)$ given by

$$H_{ij} = \begin{pmatrix} b & 0 & \ldots & 0 \\ 0 & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 0 \end{pmatrix}.$$

Clearly $Y^{(ij)}(b)$ is an automorphism of N with order p.

If b is an element of $\mathrm{Hom}_G(M_i, M_j)$, then $Y^{(ij)}(b)$ is the automorphism of N that acts as the identity on components M_{kj} for $(k, j) \neq (i, 1)$, and by $m \mapsto m \oplus mb \in M_{i1} \oplus M_{j1}$ for $m \in M_{i1}$.

Now we construct a generating set for the automorphism group of an inhomogeneous module.

Theorem 4.31 Let $\mathcal{B}^{(ij)}$ be a basis for $\mathrm{Hom}_G(M_i, M_j)$, and let T_i be a generating set for $\mathrm{Aut}_G(N_i)$ for $i, j \in \{1, \ldots, t\}$, $i \neq j$. Then

$$\left(\bigcup_{i=1}^{t} T_i\right) \cup \left(\bigcup_{\substack{i,j=1 \\ i\neq j}}^{t} \left\{Y^{(ij)}(b) \mid b \in \mathcal{B}^{(ij)}\right\}\right)$$

is a generating set for $\mathrm{Aut}_G(N)$.

Proof Let S be the group generated by the set. The proof shall mimic that of Theorem 4.26. We shall exhibit a set of elements in S that correspond to elementary row and column operations for $\mathrm{End}_G(N)$, and show that every automorphism of N reduces to a matrix that obviously lies in S.

First observe that for $b, c \in \mathrm{Hom}_G(M_i, M_j)$, we have $Y^{(ij)}(b)Y^{(ij)}(c) = Y^{(ij)}(b + c)$. Since S contains $Y^{(ij)}(b)$ for all b in the basis $\mathcal{B}^{(ij)}$, S contains $Y^{(ij)}(b)$ for all b in $\mathrm{Hom}_G(M_i, M_j)$. For each $k \in \{1, \ldots, n_i\}$ there exists a permutation matrix $P_{1k}^{(i)}$ in $\mathrm{Aut}_G(N_i)$ that interchanges rows 1 and k. Similarly for $l \in \{1, \ldots, n_j\}$ there exists $P_{1l}^{(j)} \in \mathrm{Aut}_G(N_j)$ interchanging rows 1 and l. Moreover, both of these permutation matrices lie in S. Therefore, the matrix

$$Y_{kl}^{(ij)}(b) = (P_{1l}^{(j)})^{-1}(P_{1k}^{(i)})^{-1}Y^{(ij)}(b)P_{1k}^{(i)}P_{1l}^{(j)}$$

also lies in S, and this is the matrix that differs from the identity matrix by the single entry b in the position corresponding to $\mathrm{Hom}_G(M_{ik}, M_{jl})$.

Let A be an automorphism of N. By Lemma 4.29, each A_{ii} is an automorphism of N_i. By multiplying A by the diagonal matrix whose i-th diagonal entry is $(A_{ii})^{-1}$, we obtain an automorphism A' whose diagonal entries A'_{ii} are the identity automorphisms of the homogeneous components N_i, $i \in \{1, \ldots, t\}$. Since this diagonal matrix lies in S, A will be an element of S if and only if A' is an element of S. So we need only consider those A whose diagonal blocks are the identity.

Let j be such that $A_{ik} = 0$ for all $i \in \{1, \ldots, t\}$ and $k \in \{1, \ldots, j-1\}$. Let $i \in \{j+1, \ldots, t\}$ be minimal such that A_{ij} is non-zero. Then A_{ij} is a block of size $n_i \times n_j$ over $\mathrm{Hom}_G(M_i, M_j)$. Denote by b_{kl} the entries of A_{ij} for $k \in \{1, \ldots, n_i\}$

and $l \in \{1, \ldots, n_j\}$. If columns $1, \ldots, l-1$ are zero in A_{ij}, then the matrix

$$\left(\prod_{k=1}^{n_i} Y_{kl}^{(ij)}(-b_{kl}) \right) A$$

is an automorphism of N whose corresponding (i,j)-th block has columns $0, \ldots, l$ all zero. By induction, CA is an upper triangular element of $\mathrm{Aut}_G(N)$, for some $C \in S$. A similar argument using column operations on CA leads to CAD a block diagonal matrix, with $C, D \in S$. But CAD is clearly an element of S, therefore A is also an element of S. ∎

Chapter 5

General case

In this section we consider the general case of a non-split extension $\mathcal{E}$. We start with a generating set for $\operatorname{Aut} G$ and compute a generating set for $\operatorname{Aut} \mathcal{E}$. As in the split case, we consider the map ϑ from $\operatorname{Aut} \mathcal{E}$ into the direct product of $\operatorname{Aut} G$ by $\operatorname{Aut} M$. We show that a necessary condition for a pair of automorphisms to lie in the image of ϑ is that it is a compatible pair. For split extensions this was also a sufficient condition. For non-split extensions it is no longer sufficient. Adapting some ideas from Robinson (1981), we determine a criterion for deciding whether a compatible pair lies in the image of ϑ, and when it does, a method for finding a preimage of it under ϑ. This solves the most general non-split case. We then examine a special situation that arises when the automorphism group of a soluble group is being computed; in this case the action of an automorphism on G uniquely determines the action of a lifting on M.

5.1 Scenario

Let $\mathcal{E}$ be a non-split extension. We have a transversal function $\tau : G \to H$ satisfying $g_i^\tau = h_i$ for all $g_i \in \mathcal{A}_G$, and $1^\tau = 1$. As before, let $\xi \colon G \to \mathrm{GL}(k, p)$ be the representation defining the $\mathbb{F}_p G$-module structure of M.

Let $\varphi \colon G \times G \to M$ be the factor set associated with τ. For this computation we require the factor set to be computable; that is, given (u, v) with $u, v \in \mathcal{A}_G^*$ we can evaluate $(u, v)\varphi$ as an element of M. In the context of soluble groups, where both G and H correspond to quotients of a group defined by a consistent power conjugate presentation, there is a natural choice for the transversal function, and relative to this transversal function the factor set is easily computed. If $w(g_1, \ldots, g_n)$ is a normal word representing an element g of G, then define g^τ

to be the element of H whose normal word is $w(h_1, \ldots, h_n)$. Computing the factor set with respect to this transversal function reduces to collecting a word: if $u(g_1, \ldots, g_n)$ and $v(g_1, \ldots, g_n)$ are normal words, and $w(g_1, \ldots, g_n)$ is the normal word after collection of their product, then the result after collection of

$$w(h_1, \ldots, h_n)^{-1} u(h_1, \ldots, h_n) v(h_1, \ldots, h_n)$$

is the normal word representing the value of $(u, v)\varphi$.

Recall that a pair (α, ν) in $\operatorname{Aut} G \times \operatorname{Aut} M$ is a compatible pair if it satisfies

$$\nu g^{\alpha\xi} = g^{\xi}\nu \tag{5.1}$$

for all $g \in G$. Chapter 4 detailed how to calculate a generating set for the group $C(\mathcal{E})$ of compatible pairs. In the split case $C(\mathcal{E})$ is precisely the image of ϑ. This is no longer true when the extension does not split.

Definition 5.1 A pair (α, ν) is an *inducible pair* if $(\alpha, \nu) \in \operatorname{Im} \vartheta$.

As we shall see shortly, the image of ϑ is always a subgroup $C(\mathcal{E})$. Therefore, if we have an effective criterion for deciding when a compatible pair is inducible, then we can perform an orbit-stabiliser calculation to find a generating set for $\operatorname{Im} \vartheta$. We begin by associating with each inducible pair a function from G into M which will provide the basis for a suitable criterion.

Lemma 5.2 Let $\gamma \in \operatorname{Aut} \mathcal{E}$, and $\gamma^{\vartheta} = (\alpha, \nu)$. Then

$$\gamma\colon g^{\tau} m \longmapsto g^{\alpha\tau}\, (g)\psi\, m\nu, \tag{5.2}$$

for some function $\psi\colon G \to M$.

Proof Since γ/M is α, we have

$$(g^{\tau} m)^{\gamma} = g^{\alpha\tau} n_{g,m}$$

for some element $n_{g,m}$ of M. Also $\gamma|_M$ is ν, so $m^{\gamma} = m\nu$. Therefore

$$g^{\alpha\tau} n_{g,m} = g^{\alpha\tau}\, n_{g,1}\, m\nu$$

for all $g \in G$ and $m \in M$. Hence $n_{g,m} = n_{g,1}\, m\nu$, and the function $\psi\colon g \mapsto n_{g,1}$ is the required function from G into M. ∎

For $\gamma \in \operatorname{Aut}\mathcal{E}$ we have a unique function $\psi\colon G \to M$ associated with γ. For an inducible pair (α, ν) we have at least one function ψ associated with (α, ν), since there is at least one and usually more than one preimage of (α, ν) under ϑ.

We now prove the relationship mentioned earlier between compatible pairs and the image of ϑ for a non-split extension.

Lemma 5.3 $\operatorname{Im}\vartheta \leq C(\mathcal{E})$.

Proof Let $\gamma \in \operatorname{Aut}\mathcal{E}$, and let $\gamma^{\vartheta} = (\alpha, \nu)$. Let $\psi : G \to M$ be the function associated with γ. Therefore

$$\begin{aligned} g^{\alpha\tau}\,(g)\psi\,(m\nu)^{g^{\alpha}} &= m\nu\, g^{\alpha\tau}\,(g)\psi \\ &= (mg^{\tau})^{\gamma} \\ &= (g^{\tau}m^{g})^{\gamma} \\ &= g^{\alpha\tau}\,(g)\psi\, m^{g}\nu \end{aligned}$$

for all $g \in G$ and $m \in M$. Hence $\nu g^{\alpha\xi} = g^{\xi}\nu$ for all $g \in G$, which is precisely the compatibility condition (5.1) for the pair (α, ν). ■

We now prove that while ψ is neither a homomorphism or a derivation from G to M, it is characterised by its action on a generating set of G.

Lemma 5.4 Let $\gamma \in \operatorname{Aut}\mathcal{E}$, and let $\gamma^{\vartheta} = (\alpha, \nu)$. The function $\psi : G \to M$ associated with γ is determined by its action on $\mathcal{A}_G$.

Proof Consider the equation $g^{\tau}h^{\tau} = (gh)^{\tau}\,(g,h)\varphi$. Comparing the image of each side of this equation under γ, we find that

$$g^{\alpha\tau}\,(g)\psi\, h^{\alpha\tau}\,(h)\psi = (gh)^{\alpha\tau}\,(gh)\psi\,(g,h)\varphi\nu.$$

Reducing to an equation in M and writing it additively we have

$$(g^{\alpha}, h^{\alpha})\varphi + ((g)\psi)^{h^{\alpha}} + (h)\psi = (gh)\psi + (g,h)\varphi\nu.$$

Therefore, given a word w in $\mathcal{A}_G^*$, by recursively applying the relation

$$(gh)\psi = ((g)\psi)^{h^{\alpha}} + (h)\psi + (g^{\alpha}, h^{\alpha})\varphi - (g,h)\varphi\nu \tag{5.3}$$

we can express $(w)\psi$ in terms of $(g_i)\psi$ for $g_i \in \mathcal{A}_G$. ■

We use the existence of this function ψ as a criterion for whether a compatible pair is inducible. The following theorem is a slight modification of Proposition 4.1 in Robinson (1981).

Theorem 5.5 Let (α, ν) be a compatible pair for $\mathcal{E}$. There exists $\gamma \in \operatorname{Aut}\mathcal{E}$ with $\gamma^{\vartheta} = (\alpha, \nu)$ if and only if there exists a function $\psi\colon G \to M$ satisfying (5.3) for all $g, h \in G$.

Proof We need only show that if (α, ν) is a compatible pair for $\mathcal{E}$ and there exists a map ψ satisfying (5.3), then we can find a preimage $\gamma \in \operatorname{Aut}\mathcal{E}$ of (α, ν) under ϑ. Define a map γ on H by

$$(g^{\tau} m)^{\gamma} = g^{\alpha\tau}\ (g)\psi\ m\nu. \tag{5.4}$$

This map is well defined, since every element of H has a unique expression of the form $g^{\tau} m$. Furthermore, γ is a endomorphism of H, since

$$\begin{aligned}
(g^{\tau} m)^{\gamma}(h^{\tau} n)^{\gamma} &= g^{\alpha\tau}\ (g)\psi\ m\nu\ h^{\alpha\tau}\ (h)\psi\ n\nu \\
&= g^{\alpha\tau} h^{\alpha\tau}\ ((g)\psi)^{h^{\alpha}}\ (h)\psi\ (m\nu)^{h^{\alpha}}\ n\nu \\
&= (gh)^{\alpha\tau}\ (g^{\alpha}, h^{\alpha})\varphi\ ((g)\psi)^{h^{\alpha}}\ (h)\psi\ (m\nu)^{h^{\alpha}}\ n\nu \\
&= (gh)^{\alpha\tau}\ (gh)\psi\ (g,h)\varphi\nu\ (m\nu)^{h^{\alpha}}\ n\nu, && \text{using (5.3)} \\
&= (gh)^{\alpha\tau}\ (gh)\psi\ (g,h)\varphi\nu\ m^{h}\nu\ n\nu, && (\alpha,\nu) \text{ compatible} \\
&= (gh)^{\alpha\tau}\ (gh)\psi\ ((g,h)\varphi\ m^{h} n)\nu, && \text{since } \nu \in \operatorname{Aut} M \\
&= ((gh)^{\tau}(g,h)\varphi\ m^{h} n)^{\gamma}.
\end{aligned}$$

Note that $\gamma/M = \alpha$ and $\gamma|_M = \nu$. Therefore γ is an automorphism of H, and is such that $\gamma^{\vartheta} = (\alpha, \nu)$. ∎

It is not obvious how the criterion arising from the last theorem can be used in practice. By using Lemma 5.4, which says that we can represent ψ by its n images on $\mathcal{A}_G$, we convert the criterion into a form that is more amenable to computation.

Corollary 5.6 Let (α, ν) be a compatible pair for $\mathcal{E}$. There exists $\gamma \in \operatorname{Aut}\mathcal{E}$ with $\gamma^{\vartheta} = (\alpha, \nu)$ if and only if there exist elements $u_1, \ldots, u_t \in M$ such that

$$r(g_1^{\alpha\tau} u_1, \ldots, g_n^{\alpha\tau} u_n) = s(m_1, \ldots, m_k)\nu \tag{5.5}$$

as elements of H for each relation $r(h_1, \ldots, h_n) = s(m_1, \ldots, m_k)$ in $\mathcal{R}_H^{(3)}$.

Proof Recall that $\mathcal{R}_H^{(3)}$ consists of those relations of H which are preimages of relations of G. The result hinges on proving that if such elements u_i exist, then the map $\gamma\colon \mathcal{A}_H \to H$ defined by

$$\begin{aligned} h_i &\longmapsto g_i^{\alpha\tau} u_i & \qquad 1 \le i \le n \\ m_j &\longmapsto m_j \nu & \qquad 1 \le j \le k \end{aligned}$$

extends to an automorphism of H. In the light of Theorem 5.5, if (α, ν) is inducible, then the elements u_i of M are the values on $\mathcal{A}_G$ of a function $\psi\colon G \to M$ associated with (α, ν).

If all the relations of H are invariant under γ, then γ extends to an endomorphism of H. Consider the relations in $\mathcal{R}_H$ that do not lie in $\mathcal{R}_H^{(3)}$. There are 2 cases:

1) Relations in $\mathcal{R}_H^{(1)}$ which define the elementary abelian structure of M. The action of γ on $\{m_1, \ldots, m_k\}$ is via an automorphism of M, so these relations are invariant under γ.
2) Relations in $\mathcal{R}_H^{(2)}$ of the form $m_j^{h_i} = w_{ij}(m_1, \ldots, m_k)$ which define the module action on M. The image of the right-hand side under γ is a word representing the element $w_{ij}(m_1, \ldots, m_k)\nu$ of M. The image of the left-hand side is a word representing the element $(m_j\nu)^{g_i^{\alpha\tau} u_i}$ which, by the compatibility of (α, ν), is equal to $(m_j^{g_i}\nu)^{u_i}$. The term inside the parentheses is an element of M, so conjugation by u_i does nothing to it. Hence the image of the left-hand side under γ represents the element $m_j^{h_i}\nu$ of M, and therefore the relation is invariant under γ.

If $\{u_1, \ldots, u_n\}$ is a sequence of elements of M satisfying (5.5) for all relations in $\mathcal{R}_H^{(3)}$, then the map γ extends to an endomorphism of H. In this case γ is clearly an automorphism satisfying $\gamma^\vartheta = (\alpha, \nu)$. Conversely, if $\gamma^\vartheta = (\alpha, \nu)$ and $\psi\colon G \to M$ is the function associated with γ, then

$$\{u_1, \ldots, u_n \mid u_i = (g_i)\psi, i \in \{1, \ldots, n\}\}$$

is a sequence of elements of M satisfying (5.5) for all relations in $\mathcal{R}_H^{(3)}$. ∎

Let (α, ν) be a compatible pair. We construct an inhomogeneous system of linear equations over $\mathbb{F}_p$, such that the system has a solution if and only if (α, ν) is inducible. Moreover, a solution of the system explicitly describes a function ψ associated with (α, ν), and this information allows a preimage of (α, ν) under ϑ to be constructed.

Let $U = \{u_1, \ldots, u_n\}$ and let R be an element of $\mathcal{R}_H^{(3)}$ written as a relator $r(h_1, \ldots, h_n)s(m_1, \ldots, m_k)$. Consider the word

$$W = r(g_1^{\alpha\tau} u_1, \ldots, g_n^{\alpha\tau} u_n) s(m_1\nu, \ldots, m_k\nu) \tag{5.6}$$

in $(\mathcal{A}_H, U)^*$. This word collects to a word of the form

$$\overline{W} = r'(h_1, \ldots, h_n) s'(m_1, \ldots, m_k) u_1^{f_1} \ldots u_n^{f_n}$$

where the exponents f_j are elements of $\mathbb{F}_p G$ for $j \in \{1, \ldots, n\}$. The subword $r'(h_1, \ldots, h_n)$ is equivalent to an element of M. To see this consider

$$r(g_1^{\alpha\tau}, \ldots, g_n^{\alpha\tau}) = r(g_1^{\alpha}, \ldots, g_n^{\alpha})^{\tau} s''(m_1, \ldots, m_k)$$

for some $s''(m_1, \ldots, m_k)$. The right-hand side is equal to $s''(m_1, \ldots, m_k)$ since $r(g_1, \ldots, g_n)$ is a relator of G and α is an automorphism of G. We can compute $s''(m_1, \ldots, m_k)$ since it only involves a sequence of factor set computations. Collection of the left-hand side results in $r'(h_1, \ldots, h_n)s'(m_1, \ldots, m_k)$, and hence $r'(h_1, \ldots, h_n)$ is equivalent to the element $v = s''(m_1, \ldots, m_k)\left(s'(m_1, \ldots, m_k)\right)^{-1}$ of M. Therefore we have

$$\overline{W} = v(m_1, \ldots, m_k)\, u_1^{f_1} \ldots u_n^{f_n}. \tag{5.7}$$

Let $R_1, \ldots, R_t$ be the elements of $\mathcal{R}_H^{(3)}$, and define $\overline{W}_i$ as above for each R_i, where $\overline{W}_i$ is the product of $v_i \in M$ and elements of U with exponents $f_{i1}, \ldots, f_{in}$.

By Corollary 5.6, the pair (α, ν) lies in the image of ϑ if and only if there exists a map $U \to M$ such that the image of $\overline{W}_i$ under this map is the identity for each $i \in \{1, \ldots, t\}$. If we represent such a mapping by $u_j \mapsto (u_{j1}, \ldots, u_{jk})$ for each $j \in \{1, \ldots, n\}$, then (α, ν) lies in the image of ϑ if and only if the following inhomogeneous system of linear equations in the u_{ij} has a solution:

$$(u_1, \ldots, u_n) \begin{pmatrix} A_{11} & A_{12} & \ldots & A_{1t} \\ A_{21} & A_{22} & \ldots & A_{2t} \\ \vdots & \vdots & \ddots & \vdots \\ A_{n1} & A_{n2} & \ldots & A_{nt} \end{pmatrix} = -(v_1, \ldots, v_t), \tag{5.8}$$

where A_{ij} is the $k \times k$ matrix f_{ij}^{ξ}, and the v_i are the elements of M defined above in (5.7). If the system does have a solution, say $(y_1, \ldots, y_n)$, then the function $\psi\colon G \to M$ defined by $(g_i)\psi = y_i$ is a function from G to M associated with the preimage γ of (α, ν) given by (5.4).

We therefore have a membership test for the image of ϑ as a subgroup of $C(\mathcal{E})$. Given the generating set for $C(\mathcal{E})$ as input, an orbit-stabiliser calculation returns a generating set for $\operatorname{Im}\vartheta$. We have an algorithm with the following description:

Algorithm 5.7

Input: A generating set for $C(\mathcal{E})$.

Output: A set $\{\gamma_1, \ldots, \gamma_t\} \subseteq \operatorname{Aut}\mathcal{E}$ such that $\{\gamma_1^{\vartheta}, \ldots, \gamma_t^{\vartheta}\}$ generates $\operatorname{Im}\vartheta$.

During the calculation the values on $\mathcal{A}_G$ of a function ψ associated with each pair in the generating set is obtained as a by-product. Hence we can write down a preimage under ϑ of each element of the generating set for $\operatorname{Im}\vartheta$. Together with the set of liftings of the identity arising from derivations from G to M, the set obtained from this algorithm generates the automorphism group of $\mathcal{E}$.

5.2 Soluble group case

We now consider the lifting problem in a more restricted situation. This arises when computing automorphism groups of soluble groups (as described in the next chapter). Here the elements of $\mathcal{A}_M$ can be expressed as words in earlier generators, and from this information we can compute the possible action on M that matches a given automorphism of G. In this way the module isomorphism computations required in the general case are avoided.

Let P be a characteristic p-subgroup of H exhibited by the presentation for H, and such that M is the last term in the lower exponent-p central series of P (this series is defined in the next chapter). This implies that M is central in P and every generator of M has a definition in terms of previous generators for P. For simplicity assume that the generators of P lie at the end of the polycyclic generating sequence for H. Let $l < n$ be such that

$$\{h_l, \ldots, h_n, m_1, \ldots, m_k\}$$

generates P. A definition for $m_t \in \mathcal{A}_M$ is a word w_t in generators $h_l, \ldots, h_n$ such that $w_t = m_t$ in H. Since each element of $\mathcal{A}_M$ has a definition, the images of these earlier generators under an automorphism uniquely determines the action of the automorphism on M. This fact can be used to avoid the computation of a generating set for the group of compatible pairs.

A significant improvement is obtained when the presentation for P is a labelled power conjugate presentation (see for example Celler *et al.* 1993). This means that each generator $m_t \in \mathcal{A}_M$ occurs in a relation of one of the following forms:

1) $h_j^p = m_t$ for some $j \in \{l, \ldots, n\}$, or
2) $h_j^{h_i} = h_j m_t$ for some $i, j \in \{l, \ldots, n\}$, $i < j$.

Therefore each of the generators m_t can be expressed either as a power of an earlier generator, $m_t = h_j^p$, or as a commutator of earlier generators, $m_t = [h_j, h_i]$. The right-hand side of the appropriate equation is the definition of m_t.

We use the definitions of the elements of $\mathcal{A}_M$ to associate with an automorphism α of G a unique group endomorphism ν_α of M. When α lifts to H this endomorphism ν_α is actually an automorphism of M and it is equal to the restriction to M of a lifting of α. Hence, when α lifts to H, the pair (α, ν_α) is a compatible pair for $\mathcal{E}$.

Let α be an automorphism of G. For $m_t \in \mathcal{A}_M$ with definition $w_t(h_l, \ldots, h_n)$, define m_t' to be the element of M corresponding to the word $w_t(g_1^{\alpha\tau}, \ldots, g_n^{\alpha\tau})$. Define a map ν_α from M to M by $m_t \mapsto m_t'$ for all $i \in \{1, \ldots, k\}$. Since the m_t freely generate M as an abelian group, this uniquely defines a group endomorphism ν_α of M.

Lemma 5.8 If $\alpha \in \operatorname{Aut} G$ lifts to H, then (α, ν_α) is a compatible pair for $\mathcal{E}$.

Proof Let γ be a lifting of α. Let $u_1, \ldots, u_n$ be the elements of M satisfying $g_i^{\tau\gamma} = g_i^{\alpha\tau} u_i$ for each $i \in \{1, \ldots, n\}$. Let m_t be an element of the generating set of M with $w_t(h_l, \ldots, h_n)$ as its definition. Since γ is an automorphism of H, the equation $m_t = w_t$ is invariant under γ. Consider the two possible types of definition w_t for m_t. If $w_t = h_i^p$ for some h_i, then $w_t^\gamma = ((g_i^\tau)^p)^\gamma = (g_i^{\alpha\tau} u_i)^p$. But M is central in P, so $w_t^\gamma = (g_i^{\alpha\tau})^p$ which is m_t'. The other type of definition is $w_t = [h_j, h_i]$ for some h_i and h_j. Again using the fact that M is central, $w_t^\gamma = [g_j^{\alpha\tau}, g_i^{\alpha\tau}] = m_t'$. Therefore $\nu_\alpha = \gamma|_M$ and (α, ν_α) is a compatible pair. ■

We already have an algorithm for deciding whether a compatible pair is inducible. By combining this last lemma with Corollary 5.6, we obtain a criterion for whether an automorphism α of G lifts to H. Given $\alpha \in \mathrm{Aut}\, G$ we compute the endomorphism ν_α that is associated with it. If ν_α has full rank, then $\nu_\alpha \in \mathrm{Aut}\, M$. To check whether (α, ν_α) is a compatible pair we examine the relations $\nu_\alpha g^{\alpha\xi} = g^\xi \nu_\alpha$ with g running over a minimal generating set for G. If (α, ν_α) is a compatible pair, then we apply the membership test used in Algorithm 5.7 to check whether α lifts to H. Given a generating set A for $\mathrm{Aut}\, G$, we perform an orbit-stabiliser calculation to determine a set of elements of $\mathrm{Aut}\, G$, as words in A^*, that generates the subgroup of $\mathrm{Aut}\, G$ that lifts to H. As before, a by-product of the lifting criterion for an automorphism α is a map $\psi \colon G \to M$ associated with the pair (α, ν_α) from which a lifting of α may be constructed.

5.2.1 Implementing the collection

The key to the membership test is the construction of the system of linear equations (5.8) that determine when a pair is inducible. These equations arise from collecting the generalised words given in (5.6). We now give a description of the key elements of this collection algorithm in the case of a soluble group.

We have a compatible pair (α, ν) and a relator R in $\mathcal{R}_H^{(3)}$. Since H has a power conjugate presentation, we can write R in the form

$$h_i^{p_i} = h_{i+1}^{a_{i+1}} \dots h_n^{a_n} v_i$$

if it is a power relation, or

$$h_j h_i = h_{i+1}^{a_{i+1}} \dots h_n^{a_n} v_i$$

if it is a conjugation relation. In either case the exponents a_i are non-negative integers and v_i is an element of M.

Let w_0 be a word in $\mathcal{A}_H^*$ with non-negative exponents, and let γ be the mapping of $\mathcal{A}_H$ into $(\mathcal{A}_H, U)^*$ defined by $\gamma \colon h_i \mapsto g_i^{\alpha\tau} u_i$ and $\gamma \colon m_i \mapsto m_i \nu$. We compute the image of the word w_0 under the map γ by applying the generalised collection steps given in Definition 3.5. We represent an element of the group algebra $\mathbb{F}_p G$ by a list of pairs (g, a) with $g \in G$ and $a \in \mathbb{F}_p$, where the empty list represents the zero element. Elements of M are represented by their exponent vectors. At the end of each step in the collection we have a collected part $g_c^\tau m_c\, u_1^{f_1} \dots u_n^{f_n}$ and an uncollected part w such that w_0^γ and

$$g_c^\tau m_c\, u_1^{f_1} \dots u_n^{f_n}\, w^\gamma$$

are equivalent as elements of H. The initialisation step sets g_c to the identity, m_c to the zero vector, lists f_1 through f_n to the empty list, and w to w_0. It then performs a collection process as the following pseudo-code indicates.

```
while w is non-trivial  do
    w is yw' for some y ∈ 𝒜_H
    Set w to w'
    if y = h_i for some i ∈ {1,...,n}  then
        — collect the h_i past the elements of U
        for each list f_j  do
            for each (g, a) in f_j  do
                Replace (g, a) by (gg_i^α, a)
                if gg_i^α is the identity  then
                    Increment a by 1
                end if
            end do
            if none of the gg_i^α were the identity  then
                Append (1, 1) to f_j
            end if
        end do
        — the rest of the collection
        Set m_c to m_c g_i^{αξ} + (g_c, g_i^α)φ
        Set g_c to g_c g_i^α
    else
        — the generator y was from M
        Set m_c to m_c + yν
    end if
end do
```

Each side of the relation is collected using this algorithm and the difference between the results evaluated. The matrix $A_i = f_i^\xi$ can be computed by adding the terms ag^ξ together for each (g, a) in the list f_i. The vector $v = -m_c$ is a component of the right-hand side of the inhomogeneous system of linear equations.

Chapter 6

Finite soluble group case

In this section we describe an algorithm for computing the automorphism group of a finite soluble group. Leedham-Green has developed an algorithm for computing a particular type of presentation for a finite soluble group which has a number of computational advantages over the usual power conjugate presentation. We shall call such a presentation a *special power conjugate presentation*. Algorithms for computing a special power conjugate presentation from an arbitrary power conjugate presentation are described by Eick (1993). These special presentations exhibit precisely the structure required to assemble an automorphism group algorithm for finite soluble groups from the algorithms described in previous sections.

6.1 Special power conjugate presentations

We begin by choosing a particular characteristic series for a finite soluble group and then take a well chosen polycylic generating sequence based on a refinement of the series. Many of the definitions and results in this section are well known, and can be found, for example, in Huppert (1967) and Robinson (1982). Throughout this section, unless stated otherwise, G denotes a finite soluble group, N a finite nilpotent group, and P a finite group of prime power order.

Definition 6.1 The *lower central series* of G is the series of characteristic subgroups

$$\gamma_1(G) \geq \gamma_2(G) \geq \dots$$

where $\gamma_1(G) = G$ and $\gamma_{i+1}(G) = [\gamma_i(G), G]$ for $i \geq 1$.

A group N is *nilpotent* if $\gamma_{n+1}(N) = 1$ for some n. The following result about nilpotent groups is well known.

Lemma 6.2 The following are equivalent:

1) N is nilpotent.
2) N is the direct product of its Sylow subgroups.
3) $[a, b] = 1$ for all $a, b \in N$ of coprime order.

Each term of the lower central series of G is a characteristic subgroup of G. Let $G_1 = G$ and for $i > 1$ define $G_i = \bigcap_{j=1}^{\infty} \gamma_j(G_{i-1})$. In this way we obtain a series

$$G_1 \geq G_2 \geq G_3 \geq \dots \tag{6.1}$$

of characteristic subgroups of G. Since G is finite there exists an integer $n \geq 1$ such that $G_{n+1} = G_{n+2} = 1$. The minimal such n is called the *nilpotent length* of G. The characteristic series (6.1) is called the (*lower*) *nilpotent series* of G.

The *lower exponent-p central series* of a group is of fundamental importance in computation with p-groups. In particular, the version of the p-quotient algorithm described by Newman (1977) is based on it. The series is defined for groups in general as follows.

Definition 6.3 The *lower exponent-p central series* of a group G is

$$\mathcal{P}_0^p(G) \geq \mathcal{P}_1^p(G) \geq \mathcal{P}_2^p(G) \geq \dots$$

where $\mathcal{P}_0^p(G) = G$ and $\mathcal{P}_i^p(G) = [\mathcal{P}_{i-1}^p(G), G](\mathcal{P}_{i-1}^p(G))^p$ for $i \geq 1$.

Each of the quotients $\mathcal{P}_i^p(G)/\mathcal{P}_{i+1}^p(G)$ is elementary abelian of exponent p and central in $G/\mathcal{P}_{i+1}^p(G)$. If P is a p-group, then we write $\mathcal{P}_i(P)$ for $\mathcal{P}_i^p(P)$. The subgroup $\mathcal{P}_1(P)$ is equal to the Frattini subgroup of P, which is the intersection of all maximal subgroups of P. Note that $\mathcal{P}_c(P) = 1$ for some $c \geq 0$. The smallest such c is called the *exponent-p class* of P, or simply the *class* of P.

Let N be a nilpotent group, and let $P_1, \dots, P_t$ be the Sylow subgroups of N for primes $p_1, \dots, p_t$ respectively. We now define a characteristic series of subgroups of N. This series has the property that each quotient of successive terms of the series is a direct product of elementary abelian groups (hence the name "DEA").

Definition 6.4 The *DEA-series* of N is

$$\mathcal{P}_0^*(N) \geq \mathcal{P}_1^*(N) \geq \mathcal{P}_2^*(N) \geq \dots \tag{6.2}$$

where $\mathcal{P}_0^*(N) = N$ and

$$\mathcal{P}_j^*(N) = \prod_{i=1}^{t} \mathcal{P}_j^{p_i}(P_i) \tag{6.3}$$

for all $j \geq 1$.

The factor group $\mathcal{P}_j^*(N)/\mathcal{P}_{j+1}^*(N)$ is the direct product of the quotients $\mathcal{P}_j^{p_i}(P_i)/\mathcal{P}_{j+1}^{p_i}(P_i)$ for $i \in \{1, \dots, t\}$. The next result shows that the DEA-series of a nilpotent group N is simply the lower exponent-q central series of N, where q is the product of the distinct primes $p_1, \dots, p_t$ dividing the order of N.

Lemma 6.5 If $q = p_1 \dots p_t$, then

$$\mathcal{P}_j^*(N) = [\mathcal{P}_{j-1}^*(N), N](\mathcal{P}_{j-1}^*(N))^q \tag{6.4}$$

for all $j \geq 1$.

Proof Assume that the result holds for $\mathcal{P}_{j-1}^*(N)$. It follows from Lemma 6.2 point (3) that the commutator $[\mathcal{P}_{j-1}^*(N), N]$ is equal to the product of the commutators $[\mathcal{P}_{j-1}^{p_i}(P_i), P_i]$ for $i \in \{1, \dots, t\}$. Also, since the p_i are all distinct, $(\mathcal{P}_{j-1}^*(N))^{p_1 \cdots p_t}$ is the product of $(\mathcal{P}_{j-1}^{p_i}(P_i))^{p_i}$ for $i \in \{1, \dots, t\}$. The result follows by induction. ∎

Let c_i be the class of the Sylow p_i-subgroup of N for $i \in \{1, \dots, t\}$. Then $\mathcal{P}_c^*(N) = 1$ where c is the maximum of the c_i, and $\mathcal{P}_j^*(N) \neq 1$ for all $j < c$. The integer c is the *DEA-length* of N.

We now define a characteristic series for N that is a refinement of the DEA-series. This new series has the property that each quotient of successive terms is an elementary abelian group (hence the name "EA"). Where the DEA-series descends the lower exponent-p_i central series of each Sylow p_i-subgroup simultaneously, the new series descends each of them separately.

Definition 6.6 Let N be a nilpotent group, and let $p_1, \dots, p_t$ be the distinct prime divisors of the order of N, ordered so that $p_1 < p_2 < \dots < p_t$. Let the

DEA-length of N be c. Then the characteristic series

$$\mathcal{L}_{1,p_1}(N) \geq \mathcal{L}_{1,p_2}(N) \geq \ldots \geq \mathcal{L}_{1,p_t}(N) \geq \mathcal{L}_{2,p_1}(N) \geq \ldots \geq \mathcal{L}_{c,p_t}(N) \tag{6.5}$$

is the *EA-series* of N, where

$$\mathcal{L}_{j,p_k}(N) = \prod_{i=1}^{k-1} \mathcal{P}_{j+1}^{p_i}(P_i) \times \prod_{i=k}^{t} \mathcal{P}_j^{p_i}(P_i) \tag{6.6}$$

for all $j \in \{1, \ldots, c\}$ and $k \in \{2, \ldots, t\}$.

Since the definition of $\mathcal{L}_{k,p_1}$ coincides with the definition of $\mathcal{P}_j^*(N)$, the EA-series is a refinement of the DEA-series. The class of a particular Sylow subgroup P_i may be less than the DEA-length of N, and consequently the EA-series may contain repeated subgroups leaving trivial quotients. We define a new series (called "LG" for Leedham-Green) with these repeated terms removed. A term of the EA-series survives in the new series if and only if the class of the corresponding Sylow subgroup is large enough.

Definition 6.7 The *LG-series* of N is the subseries of the EA-series (6.5) containing the subgroups $\mathcal{L}_{i,p_k}(N)$ for which $c_k \geq i$.

While the term $\mathcal{L}_{i,p_1}(N)$ may have been excised when moving from the EA-series to the LG-series, the subgroup itself still survives in the LG-series, possibly with a different label. Therefore we can always refer to the subgroup $\mathcal{L}_{i,p_1}(N)$ in the LG-series.

Having defined the LG-series for nilpotent groups, we now extend the definition to finite soluble groups. Let G be a finite soluble group of nilpotent length n. For each $i \in \{1, \ldots, n\}$ the quotient group $N_i = G_i/G_{i+1}$ is a nilpotent group. For each subgroup $\mathcal{L}_{j,p_k}(N_i)$ in the LG-series of N_i, let G_{i,j,p_k} be the subgroup of G containing G_{i+1} and satisfying

$$G_{i,j,p_k}/G_{i+1} = \mathcal{L}_{j,p_k}(N_i). \tag{6.7}$$

The series of subgroups so obtained is the LG-series of G. Clearly the LG-series is a refinement of the nilpotent series. For $w = (i, j, p_k)$ denote by G_w the subgroup G_{i,j,p_k}. If G_w and $G_{w'}$ are successive terms of the LG-series of G, then $G_w/G_{w'}$ is an elementary abelian group of prime exponent, characteristic in $G/G_{w'}$.

For $w = (i, j, p_k)$ indexing a term G_w of the LG-series of G, denote by $\mathrm{succ}(w)$ the index of the following term of the LG-series. We introduce an ordering on these indices. If $w' = (i', j', p_{k'})$, then $w < w'$ if:

1) $i < i'$, or
2) $i = i'$ and $j < j'$, or
3) $i = i'$, $j = j'$ and $p_k < p_{k'}$.

The subscript i in G_{i,j,p_k} indicates depth in the nilpotent series of G. The subgroups G_{i,j,p_1} form the DEA-series of G. The subscript j indicates depth in the DEA-series of G_i/G_{i+1}. The j-th quotient of this DEA-series is $G_{i,j,p_1}/G_{i,j+1,p_1}$ and it is a direct product of elementary abelian groups with exponents $p_1, \ldots, p_t$. The subscript p_k indicates how many of these elementary abelian subgroups are included in $G_{i,j,p_k}/G_{i,j+1,p_1}$; specifically, it includes those for primes $p_k, \ldots, p_t$.

We refine the LG-series of G to a polycyclic series with prime order factors. The special power conjugate presentation for G will later be defined as a power conjugate presentation based on such a refinement. Let $\{g_1, \ldots, g_l\}$ be a polycyclic generating sequence arising from a refinement of the LG-series of G. We define four weight functions on an element g_m of the sequence. Let $w = (i, j, p_k)$ be such that $g_m \in G_w$ but $g_m \notin G_{w'}$ where $w' = \mathrm{succ}(w)$. Define $\mathrm{wt}(g_m) = w$, $\mathrm{wt}_1(g_m) = i$, $\mathrm{wt}_2(g_m) = j$, and $\mathrm{wt}_3(g_m) = p_k$. Note that $\mathrm{wt}_3(g_m)$ is the order of g_m modulo $\langle g_{m+1}, \ldots, g_l \rangle$.

Before we can define the special power conjugate presentation we need to define some subgroups of G that will be exhibited by it.

6.1.1 Hall subgroups

Definition 6.8 Let π be a set of prime numbers. An integer n is a π-*number* if all the primes dividing it lie in π.

We denote by π' the set of all primes that do not lie in π. Hence an integer n is a π'-number if none of the primes dividing it lie in π. A subgroup of G is called a π-subgroup if its order is a π-number.

Definition 6.9 A subgroup H of a group G is a *Hall* π-*subgroup* if $|H|$ is a π-number and $|G : H|$ is a π'-number.

If G is a finite soluble group and π is a set of primes, then G has Hall π-subgroups and they are all conjugate in G. The Schur–Zassenhaus theorem states that if N is a normal Hall π-subgroup of finite group G, then complements of N exist and they are all conjugate in G.

6.1.2 Nilpotent heads

Since G_i/G_{i+1} is the direct product of its Sylow subgroups $P_1, \ldots, P_t$, its Frattini subgroup is $\Phi(P_1) \times \ldots \times \Phi(P_t)$. Hence $\Phi(G_i/G_{i+1}) = G_{i,2,p_1}/G_{i+1}$. We denote by $G_{i,*}$ the subgroup $G_{i,2,p_1}$.

Definition 6.10 The factor groups $G_i/G_{i,*}$ for $i \in \{1, \ldots, n\}$ are called the *nilpotent heads* of G.

Each nilpotent head is the Frattini quotient of one of the nilpotent factors G_i/G_{i+1}. We shall refer to the nilpotent heads of G as simply the heads of G.

Definition 6.11 Let G be a group, L a normal subgroup of G, and K a subgroup of G containing L. A complement of K/L in G is a subgroup U of G satisfying $KU = G$ and $K \cap U \leq L$.

Theorem 6.12 Every nilpotent head of a finite soluble group has a complement.

Proof See Eick (1993), Satz II.95. ∎

The last theorem states that for each $i \in \{1, \ldots, n\}$ there exists a subgroup K_i of G such that $G = K_i G_i$ and $K_i \cap G_i \leq G_{i,*}$. A complement to a head in G is called a *head complement*.

6.1.3 Special power conjugate presentations

A special power conjugate presentation for a finite soluble group G is a power conjugate presentation based on a refinement of the LG-series of G with the additional property that it exhibits Hall subgroups and head complements of G.

Definition 6.13 Let G be a finite soluble group and let $\mathcal{A}_G = \{g_1, \ldots, g_l\}$ be a polycyclic generating sequence based on a refinement of the LG-series of G. A power conjugate presentation $\{\mathcal{A}_G \mid \mathcal{R}_G\}$ for G is a *special power conjugate presentation* if it satisfies the following conditions:

1) For each set π of prime numbers, the subgroup generated by

$$\{g_m \mid \mathrm{wt}_3(g_m) \in \pi\}$$

is a Hall π-subgroup of G.

2) For each $i \in \{1, \ldots, n\}$, the subgroup generated by

$$\{g_m \mid \mathrm{wt}_1(g_m) \neq i \text{ or } \mathrm{wt}_1(g_m) = i \text{ and } \mathrm{wt}_2(g_m) > 1\}$$

is a complement of the head $G_i/G_{i,*}$ in G.

Since a special power conjugate presentation is based on a refinement of the LG-series of G, we have, in addition to those listed above, the following property:

3) For $w = (i, j, p_k)$ the index of G_w in the LG-series, the set

$$\{g_m \mid \mathrm{wt}(g_m) \geq w\}$$

contains a polycyclic generating sequence for G_w.

Eick (1993) describes an algorithm for computing a special power conjugate presentation from an arbitrary power conjugate presentation for a finite soluble group. Also described are a number of algorithms for computing structural information about the group given such a presentation. These algorithms have recently become available as a standard package in GAP (Schönert *et al.* 1993).

6.2 Automorphisms of finite soluble groups

Let G be a finite soluble group given by a special power conjugate presentation. We compute the automorphism group of G by iterating a basic step for each quotient of G by a term of its LG-series. Let K and L be successive terms of the LG-series of G. The basic step takes a generating set for the automorphism group of G/K and computes a generating set for the automorphism group of G/L.

Note that K/L is an elementary abelian p-group for some prime p, and that it is characteristic in G/L. The group G/L is an extension of G/K by K/L, and we may apply the algorithms that were described in earlier chapters.

Whether or not the extension splits, the generating set returned is the union of a set of liftings of automorphisms of G/K and a set of liftings of the identity to G/L. Moreover the set of liftings of the identity generates the subgroup of $\operatorname{Aut} G/L$ that restricts to the identity on G/K. In each case, the set of liftings of automorphisms of G/K is computed by performing an orbit-stabiliser calculation; that is, by constructing a Schreier transversal for the subgroup S of $\operatorname{Aut} G/K$ that lifts to G/L and then returning the associated Schreier generators as the generating set for S. By choosing the transversal carefully we can construct a generating set that exhibits information about how the automorphism group restricts to subgroups and previous quotient groups of the LG-series of G.

Definition 6.14 Let $\mathcal{A} = \{\alpha_1, \ldots, \alpha_r\}$ be a generating set for the automorphism group A of a group G, and let

$$G = K_0 > K_1 > \ldots > K_t > 1 \tag{6.8}$$

be a series of normal subgroups of G. We say that $\mathcal{A}$ *exhibits restrictions* to the subnormal series (6.8) if there exist index sets I_i for $i \in \{0, \ldots, t\}$ such that the set $\mathcal{A}_i = \{\alpha_j \mid j \in I_i\}$ generates the subgroup of A consisting of all automorphisms of G that restrict to the identity on G/K_i.

A generating set $\mathcal{A}$ exhibits restrictions to a subgroup K of G if it exhibits restrictions to the series $G > K > 1$. If $\mathcal{A}$ is a generating set for A exhibiting restrictions to the subnormal series given in Definition 6.14, then the subsets of $\mathcal{A}$ satisfy $\mathcal{A}_0 = \mathcal{A}$ and $\mathcal{A}_i \supseteq \mathcal{A}_j$ for $i < j$.

Let G be a finite soluble group given by a special power conjugate presentation, and let the LG-series of G be

$$G = K_0 > K_1 > K_2 > \ldots > K_l > 1. \tag{6.9}$$

Let $\mathcal{A} = \{\alpha_1, \ldots, \alpha_t\}$ be a generating set for the automorphism group A of G/K_l that exhibits restrictions to the LG-series of G/K_l. Let S be the subgroup of A that lifts to $\operatorname{Aut} G$. Let A_i be the subgroup of A generated by $\mathcal{A}_i$ for $i \in \{0, \ldots, l-1\}$. Since $\mathcal{A}_i \supseteq \mathcal{A}_j$ for $i < j$, we have a chain

$$A = A_0 > A_1 > A_2 > \ldots > A_{l-1} \tag{6.10}$$

of subgroups of A. Let $T = \{t_1, \ldots, t_r\}$ be a transversal for S in A. If T is the result of an orbit-stabiliser calculation then it typically yields little information about the chain of subgroups (6.10). Let t_i be an element of T, and let $j \in \{0, \ldots, l-1\}$ be maximal subject to $St_i \cap A_j \neq \{\}$. Choose $t'_i \in St_i \cap A_j$. Then the set $T' = \{t'_1, \ldots, t'_r\}$ is also a transversal for S in A. We will now show that the set of Schreier generators arising from this new transversal yields a large amount of information about the chain of subgroups (6.10).

Let $T'_i = T' \cap A_i$ and $S_i = S \cap A_i$ for $i \in \{0, \ldots, l-1\}$. Observe that $T'_i \supseteq T'_j$ and $S_i \geq S_j$ for $i < j$.

Lemma 6.15 T'_i is a transversal for S_i in A_i.

Proof For $\alpha \in A_i$ we show that $\alpha \in S_i t'$ for some $t' \in T'_i$. Since T is a transversal for S in A there exists $t \in T$ such that $\alpha \in St$. Since $St \cap A_i$ is not empty there is some $t' \in T'$ such that $t' \in A_i$ and $St = St'$. Therefore $t' \in T'_j$ for some $j \geq i$, and hence $t' \in T'_i$ is such that $\alpha \in S_i t'$. ∎

Definition 6.16 Let T be a transversal for S in A and define T_i, S_i and A_i as above. We say that T is *compatible* with the series of subgroups (6.10) if T_i is a transversal for S_i in A_i for each $i \in \{0, \ldots, l-1\}$.

Corollary 6.17 If T is compatible with (6.10), then

$$\mathcal{A}'_i = \{t\alpha(\overline{t\alpha})^{-1} \mid t \in T_i, \alpha \in \mathcal{A}_i\}$$

is a generating set for S_i.

If we ensure that the orbit-stabiliser algorithm constructs a transversal that is compatible with (6.10), then the resulting generating set for G/L will exhibit restrictions to its LG-series. Recall that the orbit-stabiliser algorithm computes the transversal T by constructing a sequence of partial transversals $U_1, U_2, \ldots, U_k$. The partial transversal U_k is obtained from U_{k-1} by multiplying each element of U_{k-1} by the elements of the generating set $\mathcal{A}$ and then appending to U_k some of these new elements. When the U_k obtained in this way is identical to U_{k-1} at the end of a step the transversal is complete and $T = U_k$. Given a set $\mathcal{A}$ that exhibits restrictions to the LG-series (6.9) of G we first construct a transversal T_{l-1} for S_{l-1} in A_{l-1}, by performing the usual orbit-stabiliser calculation on the set $\mathcal{A}_{l-1}$. We then construct a transversal T_{l-2} for S_{l-2} in A_{l-2} by performing the usual calculation on the set $\mathcal{A}_{l-2}$ but starting with the partial transversal T_{l-1}. After $l-1$ of these steps we have a transversal T for S in A that is of the appropriate form. A psuedo-code description follows, which can be contrasted with the usual one described in Chapter 2.

```
Set T = {1} and i = l − 1
while i ≥ 1  do
    repeat
        Note the size of T
        for each t ∈ T and α ∈ A_i  do
            Add u = tα to T if Su ∩ T = {}
        end do
    until the size of T has not increased
    Decrement i by 1
end do
Return T
```

The generating set for the stabiliser can be built up as the transversal is built up. When $u = t\alpha$ is not added to the transversal, because $Su \cap T = \{t'\}$ for some $t' \in T$, we append $u(t')^{-1}$ to the generating set for S and note that it is an element of $\mathcal{A}'_i$. Therefore, when the orbit-stabiliser calculation terminates, we have a set $\mathcal{A}' = \{\alpha'_1, \ldots, \alpha'_t\}$ generating S and also the index sets I_i for which $\mathcal{A}'_i = \{\alpha'_j \mid j \in I_i\}$.

6.2.1 Lifting to a new nilpotent layer

A simple algorithm for computing the automorphism group of a finite soluble group H simply descends the LG-series for H and applies the algorithms of the previous sections to each quotient of H. A substantial improvement is obtained by using the fact that the generating sets of automorphisms can be chosen to exhibit restrictions to the LG-series. Let N be the last non-trivial term of the nilpotent series of H. Assume that we have the automorphism group of $G = H/N$. We know that N is a direct product of its Sylow subgroups, each of which is characteristic in H. We prove that we can compute the automorphism group of H by considering the extensions of G by each of the Sylow subgroups of N separately. The subgroup of Aut G that lifts to H is the intersection of the subgroups that lift to each of the extensions by Sylow subgroups of N.

Let $P_1, \ldots, P_t$ be the Sylow subgroups of N corresponding to primes $p_1, \ldots, p_t$ respectively. Consider the group extension

$$1 \longrightarrow N \xrightarrow{\iota} H \xrightarrow{\sigma} G \longrightarrow 1$$

where ι is the inclusion map for $N < H$. Let τ be a transversal function from G into H, and $\varphi\colon G \times G \to N$ the factor set associated with τ. Note that an element h of H can be written uniquely in the form

$$h = g^{\tau} v_1 v_2 \ldots v_t$$

with $g \in G$ and $v_i \in P_i$ for $i \in \{1, \ldots, t\}$. We write v^g for $v^{g^{\tau}}$.

Let π be a set of prime numbers. Define a subgroup P_π of N by

$$P_\pi = \prod_{p_k \in \pi} P_k.$$

Therefore P_π is the Hall π-subgroup of N. Let $p = p_i$ for some $p_i \notin \pi$ and define $P_p = P_i$. For $v = v_1 \ldots v_t$ an element of N, let v_π be the product of v_k for $p_k \in \pi$, and $v_p = v_i$. Since N is the direct product of P_π and $P_{\pi'}$ we can identify N/P_π with $P_{\pi'}$. We shall make extensive use of this identification since it eliminates an excessive amount of notation that would otherwise be required.

Define the quotient groups $H_\pi = H/P_{\pi'}$ and $H_p = H/P_{p'}$ where $P_{p'} = P_{\{p\}'}$. The group H_π is an extension of G by P_π and H_p is an extension of G by P_p. The corresponding exact sequences are

$$1 \longrightarrow P_\pi \longrightarrow H_\pi \longrightarrow G \longrightarrow 1$$
$$1 \longrightarrow P_p \longrightarrow H_p \longrightarrow G \longrightarrow 1.$$

Let $\tau_\pi : G \to H_\pi$ and $\tau_p : G \to H_p$ be the transversal functions from G into H_π and H_p respectively defined by $\tau_\pi = \tau/P_{\pi'}$ and $\tau_p = \tau/P_{p'}$. Let φ_π and φ_p be the factor sets associated with τ_π and τ_p respectively. If $(g, g')\varphi = v$ for $g, g' \in G$, then, by the definition of the factor sets φ_π and φ_p, we have

$$(g, g')\varphi_\pi = v_\pi \quad \text{and} \quad (g, g')\varphi_p = v_p.$$

Now assume that π contains all but one prime p from $\{p_1, \ldots, p_t\}$, so that $N = P_\pi P_p$ and H is an extension of G by $P_\pi P_p$. The transversal function for H is related to those of H_π and H_p by

$$v_\pi^{g^\tau} = v_\pi^{g^{\tau_\pi}} \quad \text{and} \quad v_p^{g^\tau} = v_p^{g^{\tau_p}}$$

for all $g \in G$ and $v_\pi, v_p \in N$. The factor set φ for H satisfies

$$(g, g')\varphi = (g, g')\varphi_\pi \; (g, g')\varphi_p \tag{6.11}$$

for all $g, g' \in G$. We shall show that an automorphism of G lifts to H precisely when it lifts to both H_π and H_p.

Let α be an automorphism of G that lifts to both H_π and H_p and denote these liftings by γ_π and γ_p respectively. Let β_π be the restriction of γ_π to P_π. If $g^{\tau_\pi} v_\pi$ is an element of H_π, then γ_π acts by

$$(g^{\tau_\pi} v_\pi)^{\gamma_\pi} = g^{\alpha\tau_\pi} \; (g)\psi_\pi \; v_\pi^{\beta_\pi}$$

for some function $\psi_\pi : \; G \to P_\pi$ (see Lemma 5.2). Computing the image of the equation

$$(g^{\tau_\pi} v_\pi) \; (k^{\tau_\pi} u_\pi) = (gk)^{\tau_\pi} \; (g, k)\varphi_\pi \; v_\pi^k u_\pi$$

under γ_π we get the equation

$$(g^\alpha, k^\alpha)\varphi_\pi \; (g)\psi_\pi^{k^\alpha} \; (v_\pi^{\beta_\pi})^{k^\alpha} \; (k)\psi_\pi = (gk)\psi_\pi \; (g, k)\varphi_\pi^{\beta_\pi} (v_\pi^k)^{\beta_\pi}. \tag{6.12}$$

satisfied by ψ_π. Similarly, if β_p is the restriction to P_p of γ_p, then γ_p acts on an element of H_p by

$$(g^{\tau_p} v_p)^{\gamma_p} = g^{\alpha\tau_p} \; (g)\psi_p \; v_p^{\beta_p}$$

for some function $\psi_p \colon G \to P_p$, and we have an analogous equation to (6.12) arising from γ_p.

We construct an automorphism of H from the automorphisms of H_π and H_p. Define a map ψ from G into N by

$$(g)\psi = (g)\psi_\pi \; (g)\psi_p$$

for all $g \in G$.

Theorem 6.18 The map $\gamma\colon H \to H$ defined by

$$\gamma\colon g^{\tau} v_{\pi} v_{p} \longmapsto g^{\alpha\tau}\ (g)\psi\ v_{\pi}^{\beta_{\pi}} v_{p}^{\beta_{p}}$$

is an automorphism of H.

Proof First note that γ is a well defined map from H into H. To verify that γ defines an endomorphism of H we compute

$$\begin{aligned}
A &= (g^{\tau} v_{\pi} v_{p})^{\gamma} (k^{\tau} u_{\pi} u_{p})^{\gamma} \\
&= (g^{\alpha\tau}\ (g)\psi\ v_{\pi}^{\beta_{\pi}} v_{p}^{\beta_{p}})(k^{\alpha\tau}\ (k)\psi\ u_{\pi}^{\beta_{\pi}} u_{p}^{\beta_{p}}) \\
&= (gk)^{\alpha\tau}\ (g^{\alpha}, k^{\alpha})\varphi\ (g)\psi^{k^{\alpha}}\ (v_{\pi}^{\beta_{\pi}})^{k^{\alpha}} (v_{p}^{\beta_{p}})^{k^{\alpha}}\ (k)\psi\ u_{\pi}^{\beta_{\pi}} u_{p}^{\beta_{p}}
\end{aligned}$$

and

$$\begin{aligned}
B &= ((gk)^{\tau}\ (g,k)\varphi\ v_{\pi}^{k} u_{\pi} v_{p}^{k} u_{p})^{\gamma} \\
&= (gk)^{\alpha\tau}\ (gk)\psi\ (g,k)\varphi_{\pi}^{\beta_{\pi}}\ (g,k)\varphi_{p}^{\beta_{p}}\ (v_{\pi}^{k})^{\beta_{\pi}} (v_{p}^{k})^{\beta_{p}} u_{\pi}^{\beta_{\pi}} u_{p}^{\beta_{p}}.
\end{aligned}$$

If $A = B$ for all $g, k \in G$ and $v_{\pi} v_{p}, u_{\pi} u_{p} \in N$, then γ is an endomorphism of H. Since elements of P_{π} and P_{p} commute with each other, we have $A = (gk)^{\alpha\tau} A_{\pi} A_{p}$ and $B = (gk)^{\alpha\tau} B_{\pi} B_{p}$ where

$$\begin{aligned}
A_{\pi} &= (g^{\alpha}, k^{\alpha})\varphi_{\pi}\ (g)\psi_{\pi}^{k^{\alpha}}\ (v_{\pi}^{\beta_{\pi}})^{k^{\alpha}}\ (k)\psi_{\pi}\ u_{\pi}^{\beta_{\pi}} \\
B_{\pi} &= (gk)\psi_{\pi}\ (g,k)\varphi_{\pi}^{\beta_{\pi}}\ (v_{\pi}^{k})^{\beta_{\pi}} u_{\pi}^{\beta_{\pi}}
\end{aligned}$$

and A_{p}, B_{p} have analogous definitions. Clearly $A_{\pi} = B_{\pi}$ is a consequence of (6.12), and similarly for $A_{p} = B_{p}$. Therefore γ is an endomorphism of H. Clearly γ restricts to an automorphism of N and satisfies $\gamma/N = \alpha$. Hence γ is an automorphism of H. ∎

Lemma 6.19 If α is an automorphism of G that does not lift to H_{p}, then α does not lift to H.

Proof If γ is a lifting of α to H, then $\gamma/P_{p'}$ is a lifting of α to H_{p}. ∎

For brevity, denote by H_{i} the group $H_{p_{i}}$.

Corollary 6.20 For each $i \in \{1, \ldots, t\}$ let $S_i \leq \operatorname{Aut} G$ be the subgroup of automorphisms of G that lift to H_i. Then

$$S = \bigcap_{i=1}^{t} S_i$$

is the subgroup of automorphisms of G that lift to H.

We use this last corollary as the basis for an algorithm that computes the automorphism group of the extension to the next nilpotent layer, given a description of the automorphism group of the quotient above this layer. It relies on the fact that we can compute a generating set that exhibits restrictions to the nilpotent series of a group. This is a consequence of the results from earlier this section.

Some additional notation will be convenient. For $\mathcal{B}$ a set of automorphisms of the group H, and K a characteristic subgroup of H, denote by $\mathcal{B}/K$ the set $\{\beta/K \mid \beta \in \mathcal{B}\}$. We have an algorithm for computing a generating set for $\operatorname{Aut} H$. It takes as input a generating set for $\operatorname{Aut} G$. The generating set of $\operatorname{Aut} H$ produced by the algorithm exhibits restrictions to N. We call this algorithm the *lifting algorithm*. Of course, if the algorithm is supplied with a set generating a proper subgroup S of $\operatorname{Aut} G$, the end result is a (possibly proper) subgroup S' of $\operatorname{Aut} H$. This subgroup S' is such that S'/N is the subgroup of S consisting of all those automorphisms in S that lift to H.

Let $\mathcal{A}$ be the generating set for $\operatorname{Aut} G$. Consider H_1, the extension of G by the Sylow subgroup P_1 of N. Apply our lifting algorithm with set $\mathcal{A}$ as input. The result is a set $\mathcal{A}_1$ that generates the automorphism group of H_1, such that $\mathcal{A}_1$ exhibits restrictions to P_1. The set $\mathcal{A}_1/P_1$ generates the subgroup of $\operatorname{Aut} G$ consisting of all those automorphisms of G that lift to H_1. We also have a subset $\mathcal{A}'_1 \subset \mathcal{A}_1$ that generates the full subgroup of liftings of the identity to H_1.

Next consider H_2, the extension of G by P_2. We apply the lifting algorithm the with set $\mathcal{A}_1/P_1$ as input, and compute a set $\mathcal{A}_2$ of automorphisms of H_2. This set contains a subset $\mathcal{A}'_2$ which generates the full subgroup of liftings of the identity to H_2. By the corollary, the set $\mathcal{A}_2/P_2$ generates the subgroup of $\operatorname{Aut} G$ consisting of all those automorphisms of G that lift to both H_1 and H_2.

Repeating the process for $H_3, \ldots, H_t$, we end up with a set $\mathcal{A}_t$ of automorphisms of H_t. This set is such that A_t/P_t is the subgroup of $\operatorname{Aut} G$ consistings of all those automorphisms that lift to each H_i, $i \in \{1, \ldots, t\}$. That is, A_t/P_t

generates the subgroup S of Corollary 6.20. Moreover, we have sets $\mathcal{A}'_1, \dots, \mathcal{A}'_t$ which contain generating sets for the liftings of the identity to each H_i.

There are two points to note. Given a lifting of the identity automorphism of G to H_i for some $i \in \{1, \dots, t\}$, we can construct an automorphism of H from it. This automorphism will be a lifting of the identity automorphism of G to H. The second point is that while computing $\mathcal{A}_2$ from $\mathcal{A}_1$, we can simultaneously compute a set $\mathcal{A}''_1$ of words in the elements of $\mathcal{A}_1$. This set is such that $\mathcal{A}''_1/P_1$ generates a subgroup of G that lifts to both H_1 and H_2. At the next step, computing $\mathcal{A}_3$, the words in $\mathcal{A}''_1$ are adjusted again. We repeat this process for each step and for all sets $\mathcal{A}''_i$. The resulting sets $\mathcal{A}''_i$ supply automorphisms of H that are not liftings of the identity. From these and the liftings of the identity, the automorphism group of H is constructed.

6.3 Inner automorphisms

Let G and H be successive quotients of the LG-series of a finite soluble group, so that H is an extension of G by an elementary abelian p-group M for some prime p. Define the transversal function τ from G into H as before. If α is an inner automorphism of G induced by $g \in G$, then α lifts to H and a lifting is given by the inner automorphism induced by $g^\tau \in H$. Since the test for whether an automorphism lifts is relatively expensive in both the split and non-split cases, it is important for computational reasons to keep track of inner automorphisms.

Let $\mathcal{A}$ be a generating set for $\operatorname{Aut} G$ and let S be the subgroup of $\operatorname{Aut} G$ that lifts to H. Let I be a subset of $\mathcal{A}$ containing inner automorphisms, and assume that for each $\alpha \in I$ we know an element g_α of G inducing α. It is a trivial task to compute liftings of each element of I. Moreover, while constructing the transversal for S in $\operatorname{Aut} G$ during the orbit-stabiliser calculation, the action of inner automorphisms on elements of the partial transversals can be ignored. Let T be a transversal for S in $\operatorname{Aut} G$, let t be an element of T and let α be an inner automorphism. Since the inner automorphisms form a normal subgroup of $\operatorname{Aut} G$, we have $t\alpha = \alpha' t$ for some inner automorphism α'. Clearly $\alpha' \in S$ and therefore the cosets $St\alpha$ and St are equal for all inner automorphisms α.

6.3.1 Detecting inner automorphisms

During the orbit-stabiliser calculation it is possible that an inner automorphism will arise as the product of two non-inner automorphisms, $t\alpha$ say. Since this automorphism will be one of the Schreier generators returned by the algorithm, it is important to identify it as an inner automorphism where possible, since lifting it to the next quotient is then much easier.

When this is not too expensive, we can solve the problem by computing a stabiliser chain in G. Let the generating set for G be $\mathcal{A}_G = \{g_1, \ldots, g_n\}$. Let $G_1 = G$ and define G_i for $i > 1$ to be the stabiliser of g_{i-1} in G_{i-1}. Let α be an automorphism of G. The following pseudo-code defines a test for whether α is an inner automorphism.

```
Set w to the identity of G
Set i = 1 and Outer = false
while i < n and not Outer  do
    Compute the orbit of g_i under the action of G_i
    if (g_i^α)^(w^-1) = g_i^(a_i) for some a_i ∈ G_i  then
        Set w to a_i w and increment i
    else
        Set Outer = true
    end if
end do
```

If α is inner, then at the end of this loop w contains an element of G that induces α. This test is effective when the conjugacy classes of the elements of the generating set of G are small. When these classes are very large the construction of the stabiliser chain becomes too expensive to make this test worthwhile.

6.3.2 Inner liftings of the identity

Let γ be a lifting of the identity to H, and suppose that it is an inner automorphism induced by $h \in H$. Let $g \in G$ be the element of G satisfying $g^\tau = h$. Clearly the inner automorphism of G induced by g is trivial. When we compute the automorphism group of the next quotient, if γ has not been identified as an inner automorphism of H, then γ will cause many unnecessary and expensive lifting calls.

Inner automorphisms that are liftings of the identity of G can be added to the generating set of Aut H in one of two ways. The first is as automorphisms arising from derivations, in which case they correspond to conjugation by elements of M. This case arises for both split and non-split extensions. We alter the computation of the generating set for these automorphisms as follows. Let B be a basis for the derivations from G to M. A basis C for the inner derivations from G to M is easily constructed by conjugating the generators of H by the generators of M. By appending linearly independent elements of B to C we obtain a basis for the derivations that is partitioned so that a basis for the inner derivations is exhibited. From this basis a generating set for the subgroup of Aut H can be constructed and all the automorphisms corresponding to inner derivations identified.

The second source of inner automorphisms arises only for split extensions. These are the liftings of the identity that correspond to module automorphisms. Since they act non-trivially on the abelian subgroup M, they cannot correspond to inner automorphisms induced by elements of M. Instead they correspond to conjugation by $h = g^\tau$ in H where g lies in the centre of G. There are effective algorithms for computing a generating set for the centre of a group given by a special power conjugate presentation (see Eick 1993, Chapter III). If $g \in G$ is an element of the generating set for the centre of G, we append the inner automorphism induced by g^τ to the generating set of liftings of the identity.

6.4 Results and comparisons

A prototype of the algorithm described in this thesis has been implemented in GAP (Schönert *et al.* 1993). While this prototype does not incorporate all of the features of the algorithm described in the thesis, its performance is already substantially better than the current alternative for finite soluble groups. This section presents evidence supporting this claim, as well as measurements of the performance of the prototype on various finite soluble groups. A full implementation of the algorithm is planned for the near future. This newer version will have improved performance over the prototype and will extend the range of application of the algorithm substantially.

There are two widely available implementations of algorithms for computing the automorphism group of a finitely presented group.

The first, described by O'Brien (1994), is a program for computing the automorphism group of a finite p-group. It forms part of the Standard Presentation algorithm in the ANU p-Quotient Program (O'Brien 1993). The implementation of this algorithm is in C and it performs extremely well. For small primes the automorphism group of p-groups with composition lengths up to 50 can be computed within a few minutes of CPU time on a Sparc Station 10/31 (O'Brien 1994).

The other is an implementation of an algorithm described by Robertz (1976). This implementation is in C and it is available as an intrinsic function in CAYLEY (Cannon 1984). It is a general purpose algorithm that may be applied to any group whose conjugacy classes can be computed; this includes groups defined by power conjugate presentations, permutation groups and matrix groups. The automorphism group is found by computing a base and strong generating set for a permutation representation of it. The automorphisms permute certain unions of conjugacy classes of the group. The algorithm attempts to find enough of these unions of conjugacy classes so that the automorphism group acts faithfully on them. The generating set for the automorphism group is returned as a set of permutations which can then be converted back to automorphisms of the original group. We shall refer to this algorithm as the General algorithm.

One of the major advantages of implementing in GAP, and similar computer algebra languages, is that many of the fundamental structures required for implementing algebraic algorithms exist as a fundamental part of the language. These include data structures for storing and computing with matrices, finite presentations, power conjugate presentations, and homomorphisms between groups. A disadvantage is that the language is interpreted, which means that an implementation in GAP is likely to run slower than an implementation in a compiled language like C. However, Celler *et al.* (1993) have shown that, at least for the case of an implementation of the p-quotient algorithm, it is possible to come reasonably close to compiled language speeds in GAP on some calculations.

The prototype implementation of the Soluble algorithm used to obtain the results below has two significant deficiencies that should be mentioned. Despite these two problems it is still a significant improvement over the use of the General algorithm for computing the automorphism group of a finite soluble group.

The first deficiency was mentioned earlier. It is that the prototype does not incorporate all the features of the algorithm. One key omission is the ability to compute the lifting of the automorphism group to a new nilpotent layer by considering each of the extensions by Sylow subgroups separately. However, the performance of an improved version that incorporates this feature can be measured using the prototype. Perform the computations for each of these separate extensions by the Sylow subgroups using the prototype and record the times. An upper limit on the time that would be taken by a version that computes the nilpotent case properly is given by the sum of these times. An example of this calculation is given later.

The other deficiency is in the performance of some crucial sections of the algorithm. This is primarily due to some poor choices for data structures that were made early in the development of the implementation. With help from Alice Niemeyer, some of these problems were identified and removed, resulting in considerable performance improvements (more than an order of magnitude). Some more problems exist at too low a level to be sensibly replaced, and these will be addressed by moving to the new implementation.

The following naming system is adopted for the groups used in testing the Soluble and General implementations:

1) p^n is an elementary abelian group of order p^n.
2) $E_{p^n:p^e}$ is an extraspecial group of order p^n and exponent p^e.
3) S_n is the symmetric group of degree n.
4) Q_n is the quaternion group of order n.
5) The binary octohedral group, denoted by BO, is the group of order 48 defined by the following consistent power conjugate presentation.

$$\begin{array}{llllll}\{\, a,b,c,d,e \mid & a^2 = e, & & & & \\ & b^a = b^2, & b^3, & & & \\ & c^a = de, & c^b = d, & c^2 = e, & & \\ & d^a = ce, & d^b = cd, & d^c = de, & d^2 = e, & \\ & e^a = e, & e^b = e, & e^c = e, & e^d = e, & e^2 \}.\end{array}$$

6) $H \ltimes K$ is a semi-direct product of the groups H and K, with the convention that $H \ltimes K \ltimes N = (H \ltimes K) \ltimes N$.

Some of the groups used in the following tables were constructed by taking extensions of the groups listed above. The remaining groups were obtained as soluble quotients of some finitely presented groups. Niemeyer (1993) has produced an effective program, called the ANU Soluble Quotient program (SQ), that computes a power conjugate presentation for a soluble quotient of a given finitely presented group. The remaining test groups were all constructed using the SQ program.

Listed below are the initial finite presentations and the order and derived length of each of soluble groups obtained from these.

1) $K_1 = \{\, a, b \mid (ab)^2 b^{-6}, a^4 b^{-1} a b^{-9} a^{-1} b\}$ of order 1296 and derived length 6.

2) $K_2 = \{\, a, b \mid a^2 b a^{-1} b a^{-1} b^{-1} a b^{-2}, a^2 b^{-1} a b^{-1} a b a^{-1} b^2\}$ of order 2400 and derived length 5.

3) $K_3 = \{\, a, b \mid ab^2 (ab^{-1})^2, (a^2 b)^2 a^{-1} b a^2 (bab)^{-1}\}$ of order 3000 and derived length 5.

4) $K_4 = \{\, a, b \mid ab^2 a^{-1} b^{-1} ab^3, ba^2 b^{-1} a^{-1} ba^3\}$ of order 1320 and derived length 5.

5) $K_5 = \{\, a, b \mid ab^3 a^{-1} b^{-1} ab^3, ba^3 b^{-1} a^{-1} ba^3\}$ of order 5832 and derived length 5.

6) K_6 of order 1296 and derived length 6 with presentation

$$\begin{aligned} \{\, a, b \mid\ & (ab)^2 (ab^3)^2 (ab^{-1})^3 b^{-1} ab^{-1} (ab^{-2})^2 b^{-1}, \\ & (ab)^2 (ab^2)^2 bab (ab^{-2})^2 b^{-1} ab^2 (ab^{-2})^2, a^2 b^6\} \end{aligned}$$

7) $K_7 = \{\, a, b \mid (ab)^3 b^6, aba^{-1} bab^3 aba^{-1} bab^{-1}, a^2 b^{-2} aba^{-1} b^3 ab^{-1} a^{-1} b^2\}$ of order 41472 and derived length 6.

8) K_8 of order 82944 and derived length 7 obtained as a quotient of the group with presentation

$$\begin{aligned} \{\, a, b, c \mid\ & a^{-1} c^{-1} ac, b^{-1} c^{-1} bc, c^2, a^2 b^6, \\ & (ab)^2 (ab^3)^2 (ab^{-1})^3 b^{-1} ab^{-1} (ab^{-2})^2 b^{-1} c, \\ & (ab)^2 (ab^2)^2 bab (ab^{-2})^2 b^{-1} ab^2 (ab^{-2})^2 c\} \end{aligned}$$

9) K_9 of order $3^1 2^{24}$ and derived length 11 obtained as a soluble quotient of the group with presentation

$$\begin{aligned} \{\, a, b, c, d, e \mid\ & AcDeCDeADeCbeCDAeBc, \\ & CDAeBAcEBcDeCDAeBAD^2 eaEadcEdCae, \\ & acdcEdEdaEAeAEd^2 abAEd^2 abEadabEadcE^2 d, \\ & dcEBADeDeCDAeAEd^2 aceCDAeBAcEBcDeCaEa, \\ & CeCDAeBADAeAEdacDeCabcEdaEAeAEd^2 abEad, \\ & cEdEdaEAeAEd^2 abEadabcEdaEAeAEd^2 abEad^2, \\ & aeCDAeBADAeAEda^2 CbADEdaEAeAEd^2 abEadcE^2 d, \\ & CDAeAEd^2 abEaEadabEadcECaeCDAeBAcEBcDBcDe, \\ & CDAeAEdeCDAeBADAecEdaEAeAEd^2 abEadabEadcEBcDe\}, \end{aligned}$$

where an upper-case letter denotes the inverse of a generator.

10) K_{9a}, K_{9b} and K_{9c} are quotients of K_9 by terms of its LG-series. They have derived lengths 7, 8 and 10 respectively.

The consistent power conjugate presentations resulting from the SQ program were converted to special power conjugate presentations using the algorithms described and implemented by Eick (1993).

All computations in GAP and CAYLEY were performed using 10 Megabytes of workspace.

In Table 1 we list some finite soluble groups together with their order and the order of their automorphism group. The time taken to compute the automorphism group is measured in CPU seconds on a Sparc Station 10/51.

An interesting feature of the execution times listed in Table 1 is the difference between the running times for the groups K_{9a} and K_{9b}. These groups are quotients of K_9 by successive terms of its LG-series. Therefore part of the calculation of the automorphism group of K_{9b} is the construction of a generating set for the subgroup of the automorphism group of K_{9a} that lifts to K_{9b}. The index of this subgroup is much larger than the indexes involved in any of the preceding or subsequent lifting steps. As a consequence, much more computation is involved in computing the transversal for this step, resulting in the larger running time. Note that the execution times for K_{9c} and K_9 increase more slowly after this step.

Consider the two groups in Table 1 that are split extensions of $E_{3^3:3}$ by elementary abelian groups 2^6 and 7^3. The automorphism groups were computed in 7 seconds and 25 seconds respectively. Consider the related split extension of $E_{3^3:3}$ by the direct product of these elementary abelian groups. It is a group of order 592704. An improved prototype that handles nilpotent layers correctly should compute the automorphism group of this larger group within 32 seconds on a Sparc Station 10/51.

Table 2 compares the performance of the Soluble algorithm to the General algorithm. The latter is designed for a much wider class of groups than just finite soluble groups, and consequently it does not take advantage of the large amount of information available from a power conjugate presentation for a finite soluble group. We therefore expect it to be of limited use when applied to finite soluble groups. Conversely, an algorithm designed for computing the automorphism group of finite soluble groups should at least perform better than a general purpose one. The results in Table 2 show that the Soluble implementation does indeed perform better than the General algorithm. As the results show, the Soluble algorithm performs better by several orders of magnitude on soluble groups of even moderate size. All times are in CPU seconds for a Sparc Station ELC. Also listed is the degree of the permutation representation computed by the General algorithm.

Table 1 Timings for the Soluble algorithm

G	$\lvert G\rvert$	$\lvert A\rvert$	time
$BO \ltimes E_{7^3:7}$	16 464	7 056	7
$2 \ltimes E_{7^3:7^2}$	686	2 058	7
$S_4 \ltimes 3^4$	1 944	11 664	2
$S_4 \ltimes E_{3^3:3} \ltimes 11^2$	78 408	1 568 160	4
$E_{3^3:3} \ltimes 2^6$	1 728	1 254 113 280	7
$E_{3^3:3} \ltimes 7^3$	9 261	2 084 963 328	26
K_1	1 296	3 888	4
K_2	2 400	2 400	6
K_3	3 000	600	10
K_4	1 320	2 640	5
K_5	5 832	17 496	27
K_6	1 296	3 888	3
K_7	41 472	82 944	66
K_8	82 944	82 944	59
K_{9a}	49 152	393 216	64
K_{9b}	786 432	3 145 728	2 741
K_{9c}	12 582 912	201 326 592	3 000
K_9	50 331 648	805 306 368	3 316

Table 2 Comparison of Soluble and General algorithms

	Soluble	General	
G	time	time*	degree
$BO \ltimes 7^2$	7	4 250	1 371
$BO \ltimes E_{7^3:7}$	19	640	1 665
$S_4 \ltimes 3^4$	6	57	737
$S_4 \ltimes E_{3^3:3} \ltimes 11^2$	12	5 700	9 033
K_2	18	33 700	1 349
K_3	30	—	
K_4	15	97	714
K_5	80	—	
K_6	9	25	269

* "—" indicates that the calculation did not complete within 2 hours of CPU time

Chapter 7

Infinitely many examples

This chapter contains a symbolic automorphism group calculation. The end result is a description of the automorphism group of an infinite number of groups. We define an infinite family of groups by a parameterised presentation. The index set for the parameter is the set of all odd primes. For each odd prime p the group defined by the presentation is an extension of the quaternion group of order 8 by an elementary abelian p-group of order p^2. The automorphism group of an arbitrary member of this family is computed. The result is a parameterised generating set for the automorphism group of each member of the family, as well as a function of p indicating the order of the automorphism group.

Let p be an odd prime. Let Q_8 denote the group defined by the presentation

$$\{\, u, v \mid u^4 = 1, v^2 = u^2, v^u = v^{-1} \},$$

which is a quaternion group of order 8. The automorphism group of Q_8 is S_4, generated by automorphisms

$$\begin{aligned} \alpha_1 :\ & u \mapsto uv \\ & v \mapsto v, \\ \alpha_2 :\ & u \mapsto v \\ & v \mapsto u. \end{aligned}$$

Let ζ be a primitive $(p-1)$-th root of unity in $\mathbb{F}_p$. Choose α and β in $\mathbb{F}_p$ satisfying the equation $\alpha^2 + \beta^2 = -1$ (a simple counting argument shows that this is always possible, see for example Herstein 1975, Lemma 7.1.7). Since $\mathbb{F}_p$ is finite, there exist algorithms that will find such elements ζ, α, and β of $\mathbb{F}_p$ for all odd primes p (Lidl & Niederreiter 1983; Knuth 1969). We may choose β to be non-zero for all p. We may choose α to be zero if p is congruent to 1 mod 4, and non-zero otherwise.

Let

$$\bar{u} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad \bar{v} = \begin{pmatrix} \alpha & \beta \\ \beta & -\alpha \end{pmatrix}$$

be elements of $\mathrm{GL}(2, p)$ and define the subgroup $W = \langle \bar{u}, \bar{v} \rangle$. Then $\xi \colon \mathrm{Q}_8 \to W$ defined by $u \mapsto \bar{u}$ and $v \mapsto \bar{v}$ is an isomorphism. Dickson (1901) shows that there is a single conjugacy class of subgroups of $\mathrm{GL}(2, p)$ which are isomorphic to Q_8 (for a more accessible source of Dickson's results, see Huppert 1967).

Let M be an elementary abelian group of order p^2 with generating set $\{a, b\}$. The group W has a natural embedding into the automorphism group of M which we may use to define a semi-direct product, G, of Q_8 by M via the isomorphism $\xi \colon \mathrm{Q}_8 \to W$. From the presentation for Q_8, a presentation for M on generating set $\{a, b\}$, and the embedding of W into $\operatorname{Aut} M$ we can write down the parameterised presentation

$$\begin{aligned} \{\, u, v, a, b \mid\ & u^4 = 1,\ v^2 = u^2,\ v^u = v^{-1}, \\ & a^p = 1,\ b^p = 1,\ b^a = b, \\ & a^u = b^{-1},\ b^u = a,\ a^v = a^\alpha b^\beta,\ b^v = a^\beta b^{-\alpha} \} \end{aligned}$$

for G. We have already observed that there is a single conjugacy class of subgroups of $\mathrm{GL}(2, p)$ that are isomorphic to Q_8. Therefore all semi-direct products of M by Q_8 with faithful action are isomorphic to this group. The group G has a characteristic elementary abelian subgroup $\langle a, b \rangle$ which we identify with M. The quotient of G by M is isomorphic to Q_8, and we have a generating set $\{\alpha_1, \alpha_2\}$ for the automorphism group of the quotient. We apply the automorphism group algorithm described in Chapter 4 to calculate a generating set for the automorphism group of G.

7.1 Lifting non-identity automorphisms

The presentation for G exhibits a complement for M in G and we identify Q_8 with this complement. Recall that ξ is the representation of Q_8 into $\mathrm{GL}(2, p)$, and its image is the matrix group W.

Now α_1 lifts to an automorphism of G if there exists an invertible matrix X in $\mathbb{M}(2, p)$ satisfying $u^\xi X = X u^{\alpha_1 \xi}$ and $v^\xi X = X v^{\alpha_1 \xi}$. Let X be the matrix

$$\begin{pmatrix} x_1 & x_2 \\ x_3 & x_4 \end{pmatrix}$$

in $\mathbb{M}(2,p)$. The two equations that determine whether α_1 lifts to G translate to the system of linear equations

$$\left.\begin{aligned} x_1 &= -\beta x_3 + \alpha x_4 \\ x_2 &= \alpha x_3 + \beta x_4 \\ -x_3 &= -\beta x_1 + \alpha x_2 \\ -x_4 &= \alpha x_1 + \beta x_2 \end{aligned}\right\} \quad u^\xi X = X(uv)^\xi$$

$$\left.\begin{aligned} \alpha x_1 + \beta x_3 &= \alpha x_1 + \beta x_2 \\ \alpha x_2 + \beta x_4 &= \beta x_1 - \alpha x_2 \\ \beta x_1 - \alpha x_3 &= \alpha x_3 + \beta x_4 \\ \beta x_2 - \alpha x_4 &= \beta x_3 - \alpha x_4 \end{aligned}\right\} \quad v^\xi X = X v^\xi$$

in the x_i. This homogeneous system of linear equations has a 1-dimensional solution space which corresponds to the subspace spanned by

$$\nu_1 = \begin{pmatrix} (1+\alpha) & \beta \\ \beta & (1-\alpha) \end{pmatrix}$$

in $\mathbb{M}(2,p)$. Note that ν_1 has non-zero determinant. Therefore the pair (α_1, ν_1) is a compatible pair and it can be used to construct a lifting, γ_1, of α_1 to G.

A similar calculation for the automorphism α_2 leads to a 1-dimensional space of solutions spanned by

$$\nu_2 = \begin{pmatrix} \alpha\epsilon & \beta - \epsilon \\ \beta\epsilon + 1 & -\alpha\epsilon \end{pmatrix}$$

in $\mathbb{M}(2,p)$, where ϵ is 1 if p is congruent to 3 mod 4, and 0 otherwise. The determinant of ν_2 is $2\epsilon - \beta(1-\epsilon)$, which is non-zero for all p. Therefore (α_2, ν_2) is a compatible pair leading to a lifting, γ_2, of α_2 to G.

7.2 Liftings of the identity

So far we have shown that both of the generators of the automorphism group of Q_8 lift to the extension group G. Therefore the whole automorphism group of Q_8 lifts. The particular liftings γ_1 and γ_2 that we obtained will generate a subgroup of $\operatorname{Aut} G$ that has $\operatorname{Aut} \mathrm{Q}_8$ as a quotient. All that is required now is a generating set for the subgroup I of $\operatorname{Aut} G$ that consists of those automorphisms which restrict to the trivial automorphism on the quotient G/M. This subgroup I has a semi-direct decomposition into two other subgroups. The normal subgroup of this decomposition is the subgroup D of automorphisms which correspond to derivations from Q_8 to M.

Since the complements of M in G have order co-prime to M, they are all conjugate in G and hence all the derivations are inner (for example, Robinson 1982, 11.1.3). The inner derivations correspond to the inner automorphisms of G induced by elements of M, and the action of these is apparent from the presentation of G. Despite this fact we go through the general calculation to exhibit the details of this part of the algorithm. A derivation δ has the form

$$\begin{aligned} \delta\colon\ u &\mapsto a^{x_1}b^{x_2} \\ v &\mapsto a^{y_1}b^{y_2} \end{aligned}$$

and corresponds to the automorphism

$$\begin{aligned} \gamma\colon\ u &\mapsto ua^{x_1}b^{x_2} \\ v &\mapsto va^{y_1}b^{y_2} \\ a &\mapsto a \\ b &\mapsto b \end{aligned}$$

where x_1, x_2, y_1, and y_2 are elements of $\mathbb{F}_p$. We compute the Fox derivatives of the relations of G and then construct the system of linear equations that x_1, x_2, y_1 and y_2 must satisfy for γ to define a derivation. This system of linear equations is

$$\begin{aligned} x_1 + x_2 &= (\alpha - 1)y_1 + \beta y_2 \\ -x_1 + x_2 &= \beta y_1 + (1-\alpha)y_2 \\ (1-\alpha)x_1 - \beta x_2 &= (1-\beta)y_1 + \alpha y_2 \\ -\beta x_1 + (1+\alpha)x_2 &= \alpha y_1 + (1+\beta)y_2. \end{aligned}$$

These equations reduce to

$$\begin{aligned} 2x_1 &= (1+\alpha-\beta)y_1 + (-1+\alpha+\beta)y_2 \\ 2x_2 &= (1+\alpha+\beta)y_1 + (1-\alpha+\beta)y_2. \end{aligned}$$

Therefore there is a 2-dimensional solution space spanned by $(y_1 = 2,\ y_2 = 0)$ and $(y_1 = 0,\ y_2 = 2)$. These two solutions lead to automorphisms γ_3 and γ_4 of G respectively.

The final step is to find a generating set for a complement in I of D. A complement for D can be chosen so that it fixes Q_8 pointwise. These automorphisms correspond to $\mathbb{F}_p\mathrm{Q}_8$-module automorphisms of M, or equivalently, to matrices centralising W in $\mathrm{GL}(2,p)$. The centraliser in $\mathrm{GL}(2,p)$ of $\bar{u}$ is the subgroup

$$\left\langle \begin{pmatrix} x_1 & x_2 \\ -x_2 & x_1 \end{pmatrix} \;\middle|\; x_1, x_2 \in \mathbb{F}_p, x_1^2 + x_2^2 \neq 0 \right\rangle,$$

and a matrix in this subgroup commutes with $\bar{v}$ if and only if $x_2 = 0$. So the required complement is cyclic of order $p-1$ and is generated by the automorphism γ_5 corresponding to the scalar matrix ζI_2, where I_2 is the identity element in $\mathrm{GL}(2,p)$.

Therefore the automorphism group of G is generated by automorphisms $\gamma_1, \ldots, \gamma_5$ which are defined as follows:

$$\begin{aligned} \gamma_1\colon\ & u \mapsto uv \\ & v \mapsto v \\ & a \mapsto a^{(1+\alpha)} b^{\beta} \\ & b \mapsto a^{\beta} b^{(1-\alpha)}, \end{aligned}$$

$$\begin{aligned} \gamma_2\colon\ & u \mapsto v \\ & v \mapsto u \\ & a \mapsto a^{\alpha\epsilon} b^{\beta-\epsilon} \\ & b \mapsto a^{\beta\epsilon+1} b^{-\alpha\epsilon}, \end{aligned}$$

$$\begin{aligned} \gamma_3\colon\ & u \mapsto ua^{\alpha-\beta+1} b^{\alpha+\beta+1} \\ & v \mapsto va^2 \\ & a \mapsto a \\ & b \mapsto b, \end{aligned}$$

$$\begin{aligned} \gamma_4\colon\ & u \mapsto ua^{\alpha+\beta-1} b^{-\alpha+\beta+1} \\ & v \mapsto vb^2 \\ & a \mapsto a \\ & b \mapsto b, \end{aligned}$$

$$\begin{aligned} \gamma_5\colon\ & u \mapsto u \\ & v \mapsto v \\ & a \mapsto a^{\zeta} \\ & b \mapsto b^{\zeta}. \end{aligned}$$

Since all automorphisms of Q_8 lift to G and the order of the subgroup of liftings of the identity is $(p-1)p^2$, the order of the automorphism group of G is $24(p-1)p^2$.

The group G can be realised as a subgroup Y of $\mathrm{GL}(3,p)$ whose elements have the form

$$\begin{pmatrix} A & 0 \\ m & 1 \end{pmatrix},$$

where A lies in W and m is a 2-dimensional row vector over $\mathbb{F}_p$. A subgroup of the automorphism group of Y is obtained as the quotient of the normaliser of Y by the centraliser of Y in $\mathrm{GL}(3,p)$. While the structures of the normaliser and centraliser of Y vary depending on whether p is congruent to 1, 3, 5 or 7 modulo 8 (Short 1992), the quotient in all cases has order $24(p-1)p^2$. This is precisely the order of $\mathrm{Aut}\, G$ that was calculated above. Therefore the whole automorphism group of Y is obtained in this way. Furthermore, since very detailed descriptions of the normaliser and centraliser of Y are easily obtained from Short (1992) much more of the structure of $\mathrm{Aut}\, G$ is available via this construction.

7.3 Observations

We have a parameterised infinite family of groups, $\mathcal{F} = \{G_p \mid p \in \mathcal{P}\}$, where $\mathcal{P}$ is the set of odd prime numbers. The groups G_p can be described by a single parameterised presentation indexed by $p \in \mathcal{P}$ but also requiring some parameters that are algorithmically dependent on p; namely α and β in $\mathbb{F}_p$ satisfying $\alpha^2+\beta^2 = -1$. The calculation described in this chapter proved that we can write down a description of the automorphism groups of the groups in $\mathcal{F}$. This description is in terms of a parameterised generating set and the order as a function of p. The description of the automorphism groups requires additional parameters which again depend algorithmically on p: the primitive root $\zeta \in \mathbb{F}_p$ and ϵ taking values of 0 or 1 as p is congruent to 1 or 3 mod 4 respectively.

All the steps in the above calculation could in principle be performed by a sufficiently powerful computer implementation of an automorphism lifting algorithm. Systems such as MAPLE (Char *et al.* 1988) and MATHEMATICA (Wolfram 1988) routinely deal with complex expressions in indeterminates, using prespecified equations satisfied the indeterminates to reduce intermediate expressions to simpler forms. The row reductions involved in solving the systems of linear equations could be solved in the general case by a computer algorithm using these techniques, despite the fact that the matrices contain unspecified "parameters" such as p, α, β and ζ. Incorporating such symbolic calculators into an automorphism lifting algorithm, or indeed into other group theoretic algorithms, would allow a single calculation to be performed resulting in information about all members of a family of groups.

It should be noted, however, that the above calculation involves a very special situation. In this family of groups the automorphism group of the fixed quotient Q_8 lifts entirely to the whole group. In general only a subgroup of the automorphism group will lift and determining this subgroup involves an orbit-stabiliser calculation. Such a calculation may well be required to find a transversal whose length is a function of the parameters of the family. With a sufficiently powerful implementation this may still be possible. In any event, for those situations where either the whole automorphism group lifts, as here, or where the liftable subgroup has a fixed index in the automorphism group irrespective of the parameters defining the family, automorphism lifting for an infinite parameterised family of groups could be performed by an implementation of this lifting algorithm.

Bibliography

Wieb Bosma and John Cannon (1993), *Handbook of* MAGMA *functions.* Sydney University: Department of Pure Mathematics.

John J. Cannon (1984), "An Introduction to the Group Theory Language, Cayley", *Computational Group Theory*, (Durham, 1982), pp. 145–183. Academic Press, London, New York.

Frank Celler and M.F. Newman and Werner Nickel and Alice C. Niemeyer (1993), "An algorithm for computing quotients of prime-power order for finitely-presented groups and its implementation in GAP", Research Report **SMS-127-93**, School of Mathematical Sciences, Australian National University.

B.W. Char and K.O. Geddes and G.H. Gonnet and M.B. Monagan and S.M. Watt (1988), *Maple reference manual.* New York.

L.E. Dickson (1901), *Linear Groups with an Exposition of the Galois Field Theory.* Reprinted (1958) by Dover, New York.

Bettina Eick (1993), "Spezielle PAG-Systeme in Computeralgebrasystem GAP". Diplomarbeit, RWTH Aachen.

Walter Feit (1982), *The Representation Theory of Finite Groups*, North-Holland Mathematical Library, **25**. North-Holland Publishing Co., Amsterdam, New York, Oxford.

V. Felsch and J. Neubüser (1968), "Über ein Programm zur Berechnung der Automorphismengruppe einer endlichen Gruppe", *Numer. Math.*, **11**, 277–292.

R.H. Fox (1953), "Free differential calculus I", *Math. Ann.*, **57**, 547–560.

L. Gerhards and E. Altmann (1970), "A computational method for determining the automorphism group of a finite solvable group", *Computational Problems in Abstract Algebra*, (Oxford, 1967), pp. 61–74. Pergamon Press, Oxford, London, Edinburgh.

I.N. Herstein (1975), *Topics in Algebra* (2nd edition). John Wiley & Sons, New York, Chichester, Brisbane, Toronto, Singapore.

B. Huppert (1967), *Endliche Gruppen I*, Grundlehren Math. Wiss., **134**. Springer-Verlag, Berlin, Heidelberg, New York.

Rodney James and M.F. Newman and E.A. O'Brien (1990), "The Groups of Order 128", *J. Algebra*, **129**(1), 136–158.

D.L. Johnson (1980), *Topics in the Theory of Group Presentations*, London Math. Soc. Lecture Note Ser., **42**. Cambridge University Press, Cambridge.

D.L. Johnson (1990), *Presentations of Groups*, London Math. Soc. Stud. Texts, **15**. Cambridge University Press, Cambridge.

Gregory Karpilovsky (1992), *Group Representations, Volume 1*, North-Holland Mathematics Studies, **175**. Elsevier Science Publishers B.V.

Donald E. Knuth (1969), *Seminumerical Algorithms*, The Art of Computer Programming, **2**. Addison-Wesley, Reading, Mass.

Rudolf Lidl and Harald Niederreiter (1983), *Finite Fields*, Encyclopedia of Mathematics and Its Applications, **20**. Addison-Wesley, Reading, Mass.

Bernard R. McDonald (1974), *Finite Rings with Identity*. Marcel Dekker, New York.

Brendan D. McKay (1978), "Computing automorphisms and canonical labellings of graphs", D.A. Holton and Jennifer Seberry (Eds.), *Combinatorial Mathematics*, Lecture Notes in Math., **686**. Springer-Verlag.

M.F. Newman (1977), "Determination of groups of prime-power order", *Group Theory*, Lecture Notes in Math., **573**, (Canberra, 1975), pp. 73–84. Springer-Verlag, Berlin, Heidelberg, New York.

M.F. Newman and E.A. O'Brien (in preparation), "Application of computers to questions like those of Burnside, II".

Alice C. Niemeyer (1993), *Computing Presentations for Soluble Groups*, PhD thesis. Australian National University.

E.A. O'Brien (1991), "The groups of order 256", *J. Algebra*, **143**(1), 219–235.

E.A. O'Brien (1993), "Isomorphism testing for p-groups", *J. Symbolic Comput.*, **17**.

E.A. O'Brien (1994), "Computing automorphism groups of p-groups", Wieb Bosma and A.J. van Poorten (Eds.), *Proceedings of CANT '92*, (Sydney, 1992). Kluwer.

R.A. Parker (1984), "The computer calculation of modular characters (the Meat-Axe)", M.D. Atkinson (Ed.), *Computational Group Theory*, (Durham, 1982), pp. 267–274. Academic Press, London, New York.

Richard S. Pierce (1982), *Associative Algebras*, Graduate Texts in Math., **88**. Springer-Verlag, New York, Heidelberg, Berlin.

W. Plesken (1987), "Towards a Soluble Quotient Algorithm", *J. Symbolic Comput.*, **4**, 111–122.

Heinrich Robertz (1976), "Eine Methode zur Berechnung der Automorphismengruppe einer endlichen Gruppe". Diplomarbeit, RWTH Aachen.

Derek J. Robinson (1981), "Applications of cohomology to the theory of groups", C.M. Campbell and E.F. Robertson (Eds.), *Groups – St Andrews 1981*, London Math. Soc. Lecture Note Ser., **71**, pp. 46–80. Cambridge University Press.

Derek J. Robinson (1982), *A Course in the Theory of Groups*, Graduate Texts in Math., **80**. Springer-Verlag, New York, Heidelberg, Berlin.

G.J. Schneider (1990), "Computing with Endomorphism Rings of Modular Representations", *J. Symbolic Comput.*, **9**.

Martin Schönert *et al.* (1993), GAP – *Groups, Algorithms and Programming.* RWTH, Aachen: Lehrstuhl D für Mathematik.

M.W. Short (1992), *The Primitive Soluble Permutation Groups of Degree less than* 256, Lecture Notes in Math., **1519**. Springer-Verlag, Berlin, Heidelberg, New York.

Charles C. Sims (1994), *Computation with finitely presented groups.* Cambridge University Press, Cambridge.

D.E. Taylor (1987), "Pairs of generators for Matrix Groups", *The Cayley Bulletin*, **3**.

Stephen Wolfram (1988), *Mathematica: a system for doing mathematics by computer.* Reading, Mass.