Testing modules for irreducibility

Derek F. Holt and Sarah Rees

Abstract

A practical method is described for deciding whether or not a finite-dimensional module for a group over a finite field is reducible or not. In the reducible case, an explicit submodule is found. The method is a generalisation of the Parker-Norton 'Meataxe' algorithm, but it does not depend for its efficiency on the field being small. The principal tools involved are the calculation of the nullspace and the characteristic polynomial of a matrix over a finite field, and the factorisation of the latter. Related algorithms to determine absolute irreducibility and module isomorphism for irreducibles are also described. Details of an implementation in the GAP system, together with some performance analyses are included.
1991 Mathematics subject classification (Amer. Math. Soc.): 20C40, 20-04.

1 Introduction

The purpose of this paper is to describe a practical method for deciding whether or not a finite dimensional FG-module M is irreducible, where $F = GF(q)$ is a finite field and G is a finite group. The module is assumed to be defined by matrices over F for a set of generators of the group. The method is a generalisation of the 'Meataxe' algorithm of Richard Parker, which is described in [8]. Unlike Parker's algorithm, however, it does not depend on the field F being small. In fact, on the assumption that the field operations within F are performed using look-up tables, its performance is virtually independent of the field. We stress that our priority is to achieve practical speed ahead of theoretical efficiency; in [9] it is proved that the problem of reducing a d-dimensional module over a finite field $GF(q)$ can be solved in time which is polynomial in $d\log(q)$.

When M is reducible, our procedure nearly always produces a basis for a proper submodule, together with the matrices for the group generators on the submodule and quotient module. Unfortunately, there is one particular configuration, which will be described later, in which our procedure will not find a basis for a proper submodule within a reasonable time. This configuration was discovered independently by Peter Neumann and Cheryl Praeger.

When M is irreducible, the procedure produces certain output, which enables us to test M for absolute irreducibility and, in the case when M is not absolutely irreducible, to find generators for the centralising ring of M. In any case, if M is irreducible, then we can test it for isomorphism with other FG-modules.

Parker's standard meataxe uses the following test, known as Norton's irreducibility test. Let A denote the F-algebra generated by the matrices for the elements of G, and let M^T denote the module defined by the transposes of these matrices. An element ξ of A is chosen, and the nullspaces N of ξ, and N' of its transpose ξ^T are computed. Then provided that

(a) N is nonzero,

(b) every nonzero vector v in N generates the whole of M as an FG-module, and

(c) at least one nonzero vector w of N' generates the whole of M^T as an FG−module,

M is proved irreducible. Otherwise M is reducible, and a proper submodule has been identified in (b) or (c).

Part (b) of the test is clearly the part which could be time-consuming if ξ were not well chosen. In Parker's standard meataxe, elements are randomly chosen until ξ is found with a nontrival nullspace of low dimension, preferably 1-dimensional. The problem with this is that, if the field has order q, then the probability of the nullspace being nontrivial is about $1/q$ and so, for large q, we might have to make a large number of choices. Another problem arises when the dimension, e say, of the centralising field of M is large, because we cannot find a ξ with nontrivial nullspace of dimension smaller than e.

Our strategy is different. Rather than generate ξ itself randomly, we select a random element θ of the F-algebra A, calculate the characteristic polynomial $c(x)$ of θ, and then factorise it. Then we set $\xi = p(\theta)$, where $p(x)$ is an irreducible factor of $c(x)$. In this way, ξ will always have nontrivial nullspace N. In the situation where N is irreducible as an $F\langle\theta\rangle$-module (which is the case, for example, whenever $p(x)$ is a non-repeating factor of $c(x)$), it is sufficient to carry out part (b) of the test for a single nonzero vector v in the nullspace of ξ. In other cases, examination of a single vector will not give a conclusive test for irreducibility, but might prove reducibility. If not, another factor $p(x)$ of $c(x)$, or another element θ is selected.

Note that the only additional calculations compared with Parker's algorithm are the calculation and factorisation of the characteristic polynomial. Both of these can be accomplished in time $O(d^3)$, where d is the dimension of M, which is the same complexity as that of the other parts of the process, such as calculating nullspaces. We are grateful to Peter Neumann for a suggestion which gave rise to these ideas.

As one might expect from the above description, this is a *Las Vegas* algorithm, which means that there is no certainty that it will ever stop. However, if it does stop, then it will always return the correct answer. To prove that it is a useful algorithm, we need to estimate the probability that it will return an answer for a particular random group algebra element. This we are able to do, except in one particular situation, namely when M is reducible, all composition factors of M are isomorphic, and $M/Rad(M)$ is irreducible but not absolutely irreducible. (Here $Rad(M)$ denotes the radical of M; that is, the intersection of all maximal submodules of M.) In all other cases we show that the correct answer is returned for at least 0.144 of the elements in FG, although the actual proportion is really much higher than this most of the time. In the bad situation described above, it will rapidly become clear that we are almost certainly in this situation, since the characteristic polynomial will be a proper power for every element θ that we find. In examples of this type, a high proportion of the elements θ of A result in elements ξ having a nullspace that is too large to be useful for finding a proper submodule of M explicitly.

These algorithms have been implemented by the authors in the GAP system (see [10]), and they have been tested for dimensions up to about 200 with fields up to order 2^{16} (that is, for those fields for which GAP has look-up tables stored). Several of the basic components of the procedure, such as factorising polynomials and calculating nullspaces, were already present in the GAP library, and so we were able to use these as they stood. We have tested it on a large variety of examples and observed that in many cases only a single element θ needs to be considered, and we have never come across an example, other than in the bad situation described above, where more than about five or six elements θ were necessary. (The number of different θ that need be considered is partly dependent on the procedure for choosing θ, and other implementation details.)

More recently, the algorithms have been implemented in the new language MAGMA which is being developed at the University of Sydney as a successor to CAYLEY (see [2]), and they performed satisfactorily for dimensions as high as 2000.

Although reducing modules is a basic computation in its own right, we consider the algorithms described here to be part of the more general project of recognising finite matrix groups computationally, in some appropriate sense. According to a result of Aschbacher [1], all such groups fall into one of nine categories. The first aim of this recognition project is to be able to place a given matrix group in at least one of these categories (they are not quite mutually exclusive). One of the categories consists of reducible groups, and so we can now recognise them satisfactorily. The tests for absolute irreducibility and for module isomorphism are crucial parts of the recognition of some of the other categories, including groups defined over an extension field, imprimitive groups, and tensor products. Parts of these procedures have already been implemented in GAP by the authors, C.R. Leedham-Green and E.A. O'Brien. Details will be described in forthcoming papers.

The paper is organised as follows. In Section 2, we describe the procedure for testing for irreducibility more precisely, we prove that any answer that it returns is correct, and we estimate the probability that it will return an answer for a particular element θ, other than in the exceptional situation described above. In Sections 3 and 4, we describe the additional algorithms for testing for absolute irreducibility and module isomorphism. In Section 5, we discuss some implementation issues and provide some sample timings.

2 The test for irreducibility

2.1 The algorithm

As input, we are given $d \times d$ matrices $x_1, x_2, \ldots, x_r$ over the finite field F of order q, which represent the elements of a generating set of G and so define the FG-module M. We shall denote the algebra generated by the matrices x_i by A. Then A is isomorphic to a quotient of the group algebra FG. The algorithm proceeds as follows.

Step 1. Choose a 'random' element θ in A.

Step 2. Calculate the characteristic polynomial $c(x)$ of θ.

Step 3. Extract the irreducible factors of $c(x)$ in order of increasing degree. For each such factor $p(x)$, do the following.

- **(i)** Calculate $\xi = p(\theta)$.
- **(ii)** Calculate the nullspace N of ξ. If $\dim(N) = \deg(p)$, then we call $p(x)$ a *good* factor of $c(x)$.
- **(iii)** Choose a nonzero vector in N and calculate a basis of the submodule of M generated by this vector under the action of the matrices $x_1, x_2, \ldots, x_r$. If this is a proper submodule, then return the answer **reducible**.
- **(iv)** Calculate the transposed matrices $x_1^T, x_2^T, \ldots, x_r^T$, if this has not been done already. Calculate the nullspace N' of ξ^T.
- **(v)** Choose a nonzero vector in N' and calculate a basis of the submodule of M^T generated by this vector under the action of the matrices $x_1^T, x_2^T, \ldots, x_r^T$. If this is a proper submodule, then return the answer **reducible**.
- **(vi)** If $p(x)$ is a good factor, then return the answer **irreducible**.

Step 4. Go back to Step 1.

The individual computations involved in this procedure are mostly fairly routine, and some of them were already present in GAP. The method for choosing random elements from A is important, since theoretically we require genuinely random elements, but for reasons of efficiency we cannot afford too many matrix multiplications for each such choice. Since taking linear sums of matrices is comparatively inexpensive, the problem is essentially the same as that of choosing random elements from G. There is a lengthy discussion of this question in an earlier work by the authors on matrix groups [5]. The method used in our implementation is due to Charles Leedham-Green, and is justified in [7]. The basic idea is as follows. Initially we have r generators $x_1, x_2, \ldots, x_r$ of G, but this generating set is enlarged by the addition of a new generator before the choice of each random element θ from A. The new generator is chosen as a random product $x_i x_j$ of two of the existing generators, where $i \neq j$. So, in general, we have generators $x_1, x_2, \ldots, x_s$ of G, where $s \geq r$, and $s = r$ initially. The element θ is then chosen as a linear sum $\alpha_1 x_1 + \alpha_2 x_2 + \ldots + \alpha_s x_s$, where the α_i are random elements of F.

Of course, to get genuinely random elements we would need to do an enormous amount of preprocessing, which we cannot afford. Our method seems to perform reasonably satisfactorily in practice although, particularly when there are few initial generators and the field is small, the first few elements θ chosen are sometimes too atypical to be useful.

The algorithm used to calculate the characteristic polynomial of a matrix is described in [5]. It is in fact quite straightforward, and just involves calculating the orbits of vectors under the matrix. The factorisation and nullspace algorithms were already implemented in GAP. The factorisation algorithm uses the Cantor-Zassenhaus method (see [3]). It works by successively extracting the factors of $p(x)$ of degree e, for $e = 1, 2, 3 \ldots$. The nullspace calculation essentially involves column-reducing the matrix to reduced echelon form, and the submodule computation is routine. If we find a proper submodule L' of M^T in the transposed case, then we have to find a corresponding submodule L of M. In fact, in terms of row vectors, L is simply the orthogonal complement of L', and so this is also straightforward.

2.2 The correctness of an answer

The algorithm only returns the answer **reducible** when it has specifically found a submodule, so the correctness of that answer is clear. So assume that the answer **irreducible** is returned. The proof below is an adaptation of the correctness proof of Parker's original Meataxe, but this does not seem to have ever appeared in print.

Now the answer **irreducible** can only be returned when $p(x)$ is a good factor of $c(x)$. So, in order to prove correctness, we need only show that, if M is reducible, with a proper submodule L, then an application of the test with an element θ, whose characteristic polynomial has a good factor $p(x)$, will produce a proper submodule, and hence the correct answer **reducible**.

So suppose we have a submodule L, an element θ and a corresponding good factor $p(x)$. Since $\theta|_N$ acts with minimal polynomial $p(x)$ on N and $dim(N) = deg(p)$, $\theta|_N$ must act irreducibly on N. It follows that the subspace $L \cap N$, which is fixed by θ, must either be trivial or equal to the whole of N. In the latter case, the chosen nonzero vector in N will lie in L, and so a proper submodule of M (contained in L) will be found, and the answer **reducible** returned at Step 3(iii).

Suppose, on the other hand, that $L \cap N = \{0\}$. Let M have F-basis $e_1, e_2, \ldots, e_d$, where $e_1, \ldots, e_c$ is a basis of L, and let $f_1, f_2, \ldots, f_d$ be the basis orthonormal to $e_1, e_2, \ldots, e_d$ (in terms of the usual scalar product). Then, with respect to the first of these bases, all matrices ζ in A have the form

$$\begin{pmatrix} \zeta^{(1)} & 0 \\ \zeta^{(2)} & \zeta^{(3)} \end{pmatrix},$$

where $\zeta^{(1)}$ and $\zeta^{(3)}$ are $c \times c$ and $(d-c) \times (d-c)$ matrices, respectively. With respect to the second basis, the form of ζ^T is then

$$\begin{pmatrix} \zeta^{(1)T} & \zeta^{(2)T} \\ 0 & \zeta^{(3)T} \end{pmatrix}.$$

Since $L \cap N = \{0\}$ and N is the nullspace of $\xi = p(\theta)$, $\xi^{(1)}$ must have nullity 0, and so rank c, and so $\xi^{(3)}$ has rank $d - c - \dim(N)$ and nullity $\dim(N)$. It follows that the nullspace N' of ξ^T is wholly contained in the submodule of dimension $d - c$ spanned by $f_{c+1}, f_{c+2}, \ldots, f_d$. Hence any vector in N' lies in that submodule, and so the correct answer **reducible** will be returned in Step 3(v). This completes the proof that an answer **irreducible** is correct.

2.3 The likelihood of a decision being reached

We turn now to the question of how quickly the procedure can be expected to produce an answer. In other words, we want to see what proportion of the elements θ will result in a decision. We shall derive lower bounds for this, although they will not be very accurate ones, and by working harder, we could obtain considerably better estimates. These would be of limited value, however, given the inevitable deficiencies in our procedure for choosing random elements θ. Furthermore, in the one bad situation mentioned in the introduction, we are not able to compute a general sensible lower bound.

First we shall summarise the representation theory that we need to carry out these arguments. This material can all be found in Chapter IV of [4], for example. A minor difference from that treatment is that we are assuming that M is a right rather than a left FG-module, so we talk of right ideals rather than left.

Recall that A is defined to be the F-algebra of matrices generated by the matrices of G, which is also known as the *enveloping algebra* of the module M. Clearly M is a faithful right A-module. Let $\bar{A}$ denote $A/Rad(A)$, where the radical $Rad(A)$ of A is defined as the sum of all nilpotent right ideals (see (24.5) of [4]). Then $\bar{A}$ is semisimple and is isomorphic to a direct sum of simple rings $\oplus_{i=1}^{r}\bar{A}_i$ ([4] (25.15)). Furthermore, any irreducible right $\bar{A}$-module is isomorphic as a module to a minimal right ideal of $\bar{A}$, and any such ideal is contained in one of the simple components $\bar{A}_i$. All minimal right ideals contained in a particular $\bar{A}_i$ are isomorphic as $\bar{A}$-modules. Thus any irreducible right $\bar{A}$-module is annihilated under right multiplication by all components $\bar{A}_i$ except one. In fact by (25.24) of [4], if L is any irreducible right A-module, then $LRad(A) = 0$ and so L can be regarded as an irreducible $\bar{A}$-module, and the above applies. In particular, this applies to any irreducible constituent of M.

If M itself is irreducible then, since it is a faithful A-module, we must have $Rad(A) = 0$ and A is a simple ring. Then by Wedderburn's Theorem ([4] (26.4)), A is isomorphic to the full ring $\mathcal{M}_n(E)$ of $n \times n$ matrices over the ring $E = Hom_F(M, M)$, and E is a division ring with F in its centre. In our situation, F and E are finite, so E must be isomorphic to an extension field $GF(q^e)$ of F for some e. The natural right $\mathcal{M}_n(E)$-module of row vectors has dimension n over E, and hence dimension ne over F. As an F module it must be isomorphic to M, so it follows that $d = ne$. In the special case $e = 1$, A is simply equal to the set $\mathcal{M}_d(F)$ of all $d \times d$ matrices over F. When $e > 1$, the matrices in A (with respect to an appropriate basis) can be regarded as $d/e \times d/e$ matrices over E, where the elements of E correspond to $e \times e$ submatrices.

We split our examination of the successful termination of the procedure into four cases.

In the first case the module M is irreducible. Thus the only way in which our algorithm can reach a decision is by finding an element θ of A for which the characteristic polynomial has a good factor $p(x)$, with the dimension of the nullspace of $p(\theta)$ equal to $\deg(p)$. To estimate the likelihood of a decision we need to estimate how likely we are to find such a θ and $p(x)$.

In the remaining cases M is reducible. We might prove reducibility without a good factor. It would be sufficient to find θ and a corresponding factor $p(x)$ with the nullspace of $p(\theta)$ contained in a submodule.

Case (i): M is irreducible.
Let $E = Hom_F(M, M)$, as above. Then E is isomorphic to $GF(q^e)$ for some e. Suppose first that $e = 1$, and so $A = \mathcal{M}_d(F)$. If our chosen element θ from

A has an unrepeated eigenvalue λ in F, then $c(x)$ will have an unrepeated linear factor $p(x) = x - \lambda$, and the nullspace of $p(\theta)$ will have dimension 1. Thus $p(x)$ will be a good factor of $c(x)$ for the element θ, and the correct answer will be returned by the procedure. In fact it can be shown that for large n and q, the proportion of elements in $\mathcal{M}_d(F)$ with this property is about $1 - exp(-1)$, which is greater than $1/2$. We will be content with a much cruder estimate, however, as follows.

The number of elements of $\mathcal{M}_d(F)$ with a specific unrepeated eigenvalue $\lambda \in F$ is the same as the number with 0 as an unrepeated eigenvalue. It is therefore equal to the product of the number of one-dimensional spaces $\langle v \rangle$ by the number of elements which map $\langle v \rangle$ to zero and act non-singularly on $M/\langle v \rangle$. Thus for this number, which we'll call $n_{d,q}$, we have that

$$n_{d,q} = \frac{q^d - 1}{q - 1} q^{d-1} |GL(d-1, q)|.$$

Then we observe that

$$n_{d,q} = \frac{1}{q-1} \prod_{i=1}^{d} (q^d - q^i) = \frac{|GL(d,q)|}{q-1}.$$

Now the number of elements of $GL(d-1, q)$ with a specific unrepeated eigenvalue μ is at most equal to the number of such elements in $\mathcal{M}_{d-1}(F)$; that is, $n_{d-1,q}$. So the number of elements of $\mathcal{M}_d(F)$ with specific distinct unrepeated eigenvalues λ and μ is at most

$$\frac{q^d - 1}{q - 1} q^{d-1} n_{d-1,q} = \frac{q^d - 1}{q - 1} q^{d-1} \frac{|GL(d-1, q)|}{q - 1} = \frac{n_{d,q}}{q - 1}.$$

Hence, by the inclusion/exclusion principle, at least $qn_{d,q} - \binom{q}{2} n_{d,q}/(q-1) = qn_{d,q}/2$ elements have an unrepeated eigenvalue. Now

$$qn_{d,q}/2|\mathcal{M}_d(F)| = (1 - q^{-2})(1 - q^{-3}) \ldots (1 - q^{-d})/2.$$

The smallest values of this expression occur for $q = 2$ with d large, when it tends towards the limit 0.2888. Thus, in all cases, at least 0.288 of the possible elements θ of A will have a linear good factor, and so we expect to find one reasonably quickly. Of course, there are many other types of good factor.

In the case when $e > 1$, the same argument holds, except that an (unrepeated) linear factor over E corresponds to a factor of $c(x)$ of degree e over the field F. For this factor to be irreducible (and hence a good factor), it is necessary that the eigenvalue λ does not lie in any proper subfield of E containing F. But this is the case for at least half of the elements λ in E (and usually a much higher proportion), and so we still expect to find a suitable element θ reasonably quickly. In fact, in the worst case $q = 4$, our counting argument shows that at least 0.234 of the elements in $\mathcal{M}_n(E)$ have characteristic polynomial with a good quadratic irreducible factor. However, this case demonstrates that we cannot always rely on being able to use a linear factor of $p(x)$.

Case (ii): $M/Rad(M)$ has nonisomorphic irreducible direct summands.
We recall that $Rad(M)$ is equal to the intersection of all maximal submodules of M, and $\bar{M} = M/Rad(M)$ is isomorphic to a direct sum of irreducible submodules. Suppose that $\bar{M}$ has nonisomorphic irreducible direct summands $\bar{M}_1$ and $\bar{M}_2$. We aim to estimate the number of θ in A for which the image in $\bar{M}_2$ has trivial nullspace, but the image in $\bar{M}_1$ has nontrivial nullspace. We shall see that, for such elements θ, the algorithm will always return a proper submodule. By the above, we have $\bar{A} = \oplus_{i=1}^{r}\bar{A}_i$, where $\bar{M}_i\bar{A}_j = 0$ for $i = 1$ and 2 and for all $j \neq i$. Let $\bar{\theta} = \theta + Rad(A)$. Then $\bar{\theta} = \bar{\theta}_1 + \bar{\theta}_2 + \bar{\theta}_3$ with $\bar{\theta}_1 \in \bar{A}_1$, $\bar{\theta}_2 \in \bar{A}_2$ and $\bar{\theta}_3 \in \bar{B} = \oplus_{i=3}^{r}\bar{A}_i$. Then, since θ is a random element of A, $\bar{\theta}_1$ and $\bar{\theta}_2$ will be random elements of $\bar{A}_1$ and $\bar{A}_2$, which are isomorphic to full matrix rings over extension fields of F. Furthermore the characteristic polynomial $c(x)$ of θ will be divisible by the product of the characteristic polynomials $c_1(x)$ and $c_2(x)$ of the induced actions of θ on the constituents $\bar{M}_1$ and $\bar{M}_2$. Now, since $\bar{A}_1$ and $\bar{A}_2$ are simple algebras, they each have a single isomorphism type of nonzero irreducible module. It follows that $c_1(x)$ and $c_2(x)$ are precisely the characteristic polynomials of $\bar{\theta}_1$ and $\bar{\theta}_2$ respectively, where, using Wedderburn's Theorem, $\bar{A}_1$ and $\bar{A}_2$ are regarded as matrix rings over extension fields of F. Now, if $c_1(x)$ has an irreducible F-factor $p(x)$ which is not a factor of $c_2(x)$, then when we deal with the factor $p(x)$ of $c(x)$ in the procedure, the nullspace N of $p(\theta)$ will be nontrivial, but will map onto the zero subspace of M_2. Any nonzero vector in N will therefore lie in a proper submodule of M, and the procedure will return the correct answer **reducible**. We claim that, for any choice of $c_1(x)$ and any particular irreducible factor $p(x)$ of $c_1(x)$, the characteristic polynomial $c_2(x)$ of $\bar{\theta}_2$ will have $p(x)$ as a factor for at most 0.712 of the possible choices of $\bar{\theta}_2$. This proves that the procedure will return an answer for at least 0.288 of the elements θ in this case.

To establish the claim, note that $\bar{\theta}_2$ is a random element of a full matrix ring $\mathcal{M}_t(E)$, for some extension field $E = GF(q^e)$ of F. For simplicity assume that $e = 1$ (the lower bound is in fact larger when $e > 1$). Suppose that $p(x)$ has degree s. Then, if $p(x)$ divides $c_2(x)$, the underlying t-dimensional vector space over F has at least one basis $b_1, b_2, \ldots, b_t$ for which $\bar{\theta}_2$ fixes the subspace W generated by $b_1, b_2, \ldots, b_s$, and acts on it as the companion matrix of $p(x)$. For a particular $\bar{\theta}_2$ with this property, any of the nonzero vectors in W can be chosen as b_1, but $b_2, \ldots b_s$ are determined by b_1. Thus the total number of such $\bar{\theta}_2$ is equal to at most

$$(q^t - 1)(q^t - q)\dots(q^t - q^{(s-1)})q^{t(t-s)}/(q^s - 1),$$

and this number divided by $|\mathcal{M}_t(F)|$ is equal to

$$(1 - q^{-t})(1 - q^{1-t})\dots(1 - q^{s-1-t})/(q^s - 1).$$

This is clearly less than $1/2$, except when $q = 2$ and $s = 1$. In that case, the precise proportion is equal to the proportion of singular matrices in $\mathcal{M}_t(F)$, which is

$$1 - (1 - 2^{-1})(1 - 2^{-2})\dots(1 - 2^{-t}),$$

and this tends to the limit 0.7112 as t approaches infinity.

Case (iii): $M/Rad(M)$ is reducible and all of its irreducible direct summands are isomorphic.
Suppose that $\bar{M} = M/Rad(M) = \bar{M}_{1,1} \oplus \bar{M}_{1,2} \oplus \dots \oplus \bar{M}_{1,t}$ with $t > 1$ and each $\bar{M}_{1,i}$ isomorphic to the irreducible module $\bar{M}_1$. Then $A/Rad(A) = \bar{A}_1 \oplus \bar{B}$, where $\bar{B}$ annihilates $\bar{M}_1$ and $\bar{A}_1$ is a full matrix ring over an extension field E of F. Let $|E : F| = e$, and let $\bar{\theta} = \bar{\theta}_1 + \bar{\theta}_2$ with $\bar{\theta}_1 \in \bar{A}_1$ and $\bar{\theta}_2 \in \bar{B}$. Now, $\bar{\theta}_1$ is a random element in $\bar{A}_1$ and, as we saw in Case (i), the characteristic polynomial $c_1(x)$ of $\bar{\theta}_1$ will have an unrepeated linear good factor $p(x)$ over E having degree e over F for at least 0.234 of the elements $\bar{\theta}_1$. In this case, the nullspace N of $p(\theta)$ maps onto subspaces $\bar{N}_i$ of $\bar{M}_{1,i}$ having dimension zero or one over E. Clearly, for any i and j such that $dim_E(\bar{N}_i) = dim_E(\bar{N}_j) = 1$, there is an EG-isomorphism (and hence an FG-isomorphism) from $\bar{M}_{1,i}$ to $\bar{M}_{1,j}$ that maps $\bar{N}_i$ to $\bar{N}_j$ and hence, for any nonzero elements $\bar{n}_i$ of $\bar{N}_i$, there is an FG-isomorphism from $\bar{M}_{1,i}$ to $\bar{M}_{1,j}$ that maps $\bar{n}_i$ to $\bar{n}_j$. Suppose that the chosen nonzero element n of N maps onto $\bar{n}_i$ in $\bar{N}_i$. If all of the $\bar{n}_i$ are zero, then n will lie in $Rad(M)$. If some of the $\bar{n}_i$ are nonzero, then $n + Rad(M)$ will lie in an irreducible submodule of $\bar{M}$. In either case n will lie in a proper submodule of M, and the procedure will return the correct answer **reducible**.

Case (iv): $M/Rad(M)$ is irreducible.
Let $M/Rad(M) = \bar{M}_1$. Now if $Rad(M)$ has an irreducible component not isomorphic to $\bar{M}_1$, then by similar arguments to those used in Case (ii), we can show that, for at least 0.288 of the elements θ of A, $c(x)$ will have an irreducible factor $p(x)$ such that the nullspace of $p(\theta)$ lies in $Rad(M)$, in which case the procedure will return the correct answer **reducible**.

Suppose therefore that all composition factors of M are isomorphic to $\bar{M}_1$ and let L be a submodule of $Rad(M)$ with $Rad(M)/L$ irreducible, and hence isomorphic to $\bar{M}_1$. As usual, let $E = Hom_F(\bar{M}_1, \bar{M}_1)$, regarded as an extension field $GF(q^e)$ of F.

Suppose first that $e = 1$, and so $E = F$. Let $\bar{A} = A/Ann_A(L)$, where $Ann_A(L)$ is the annihilator in A of L. Then $\bar{A}$ is isomorphic to the algebra generated by the matrices of G in their induced action on M/L. Relative to an appropriate basis, the elements α of $\bar{A}$ have the form

$$\begin{pmatrix} \alpha^{(1)} & 0 \\ \alpha^{(2)} & \alpha^{(1)} \end{pmatrix},$$

and the subspace of matrices α for which $\alpha^{(1)} = 0$ is an ideal $\bar{A}_1$ of $\bar{A}$. Furthermore, $\bar{A}/\bar{A}_1$ is isomorphic to the ring $\mathcal{M}_n(F)$ of $n \times n$ matrices over F, where $n = dim_F(\bar{M}_1)$. Since M/L is not completely reducible, $\bar{A}$ cannot be simple, and so $\bar{A}_1$ must be nonzero. Let $B = \{ \alpha^{(2)} \mid \alpha \in \bar{A}_1 \}$. Then, since $\bar{A}_1$ is an ideal of $\bar{A}$, B is closed under left and right multiplication by elements of $\mathcal{M}_n(F)$, and so B must itself be equal to $\mathcal{M}_n(F)$. Let $\theta \in A$ map onto $\bar{\theta} \in \bar{A}$. As we saw in Case (i), $\bar{\theta}^{(1)}$ will have an unrepeated eigenvalue $\lambda \in F$ for at least 0.288 of the elements θ of A. Now, amongst those matrices with a given value of $\bar{\theta}^{(1)}$ having the unrepeated eigenvalue λ, the nullspace of $p(\bar{\theta})$ (where $p(x) = x - \lambda$ for some λ) will lie in $Rad(M)$ for $(q-1)/q$ of the possible matrices $\bar{\theta}^{(2)}$, and for these values of $\bar{\theta}$, the procedure will return the correct answer **reducible**. Since $\bar{\theta}^{(2)}$ is a random element of $\mathcal{M}_d(F)$, the correct answer will be returned for at least 0.144 of the elements θ of A in all cases.

The case in which $F \subset E$ is the bad situation, which we currently cannot handle satisfactorily; in this case the proportion of elements θ for which the characteristic polynomial has a factor $p(x)$ which is either good or has nullspace completely contained in a submodule can be very low. If the procedure fails to find a useful choice of θ and $p(x)$, an exhaustive search through the whole nullspace of some $p(\theta)$ is necessary. In this case we do not know of any method of finding a basis for a proper submodule which has expected running time less than $O(|E|)$, which can of course be impractically large, even for quite moderate values of q and d.

To describe a typical situation in which this occurs, choose $e, f > 1$, and put $d = ef$. Let $F = GF(q), E = GF(q^e)$ and let $\iota : \mathcal{M}_f(E) \to \mathcal{M}_d(F)$ be an embedding. Let $D = \iota(\mathcal{M}_f(E))$ and $H = \iota(SL(f, q^e))$. Finally, let γ be a $d \times d$ matrix of order e normalising H and acting on H as a field automorphism of $GF(q^e)$. Then our module M is defined by the group of $2d \times 2d$ matrices of the form

$$\begin{pmatrix} \alpha & 0 \\ \beta\gamma & \alpha \end{pmatrix},$$

where $\alpha \in H$, $\beta \in D$. We tested the procedure on an example of this type with $d = 26$, $e = 13$, $f = 2$ and $q = 2$ and verified that this is indeed a problem.

As we have already mentioned, this example also gives problems with the standard meataxe, since every element has nullspace with dimension divisible by, and therefore greater than or equal to, the degree of the extension field.

We should like to thank Klaus Lux for pointing out an error in a version of this example which appeared in a preprint of this paper.

3 Testing for absolute irreducibility

Let us assume that, using the method described in the preceding section, we have proved that our FG-module M is irreducible. We assume also that we have stored the element θ of A (together with an expression for θ as a linear function in the generators), its characteristic polynomial $c(x)$, the good irreducible factor $p(x)$ of $c(x)$, and the matrix $p(\theta)$ having nullspace N with dimension equal to the degree f of p.

We shall now describe a method for determining the centralising field $E = GF(q^e)$ of M, together with a $d \times d$ matrix ζ which generates E as a field over F. (In other words, ζ centralises each of the matrices x_i and has minimal polynomial over F of degree $e = |E : F|$.) In particular, M is absolutely irreducible if and only if $E = F$ (see, for example, (29.13) of [4]).

The multiplicative group of E is cyclic of order $q^e - 1$. Let ρ be an element of E of order $q^e - 1$. In the method described here, we do not find such a ρ explicitly, since we only require an element ζ that generates E over F as a field. We know, however, that the element ζ that we are seeking will be a power of ρ.

Since ρ centralises M, it must centralise θ, and therefore must preserve the nullspace N of $p(\theta)$, and so $\rho|_N$ centralises the action $\theta|_N$ of θ on N. But, as an $F\langle\theta\rangle$-module, N is irreducible of dimension $f = deg(p)$. So, by Schur's lemma, the centraliser C of $\theta|_N$ in $GL(f, F)$ is isomorphic to the multiplicative group of $GF(q^f)$, which is cyclic of order $q^f - 1$. Thus $\rho|_N$ generates the unique cyclic subgroup of C of order $q^e - 1$. Let σ be an element of C of order dividing $q^e - 1$. Then σ is equal to a power $(\rho|_N)^k$ of $\rho|_N$. Let $\zeta = \rho^k$. Then $\zeta|_N = \sigma$, and, since ζ commutes with each generator x_i of G, and the translates of N under the action of G span the whole module, the full action of ζ on M can be calculated once σ is known. If σ has minimal polynomial of degree e, then so does ζ. Thus, we can find ζ explicitly by first finding elements of the appropriate order in C.

Since each centralising matrix can be written with respect to an appropriate basis as a matrix with identical $e \times e$ blocks down the diagonal, it is clear that e must be a common divisor of d and f. In our algorithm to determine e, we try all such common divisors e', in decreasing order, and test whether there is an element ζ with the above properties. If we find such a ζ, then we know that $e' = e$, and we stop; since we are considering the e' in decreasing order, we know at this stage that e cannot be any larger than e'.

For each such e', we select random matrices τ from the centraliser C of $\theta|_N$, starting with $\theta|_N$ itself. Random elements of C are easy to construct, since a centralising element is determined by its action on a single vector v_0, whose images under $\langle\theta\rangle$ span N. Thus we merely have to choose $v_0\tau$ to be a random vector v_1 in N, and then we can calculate τ from v_1. Next we calculate $\sigma = \tau^{(q^f-1)/(q^{e'}-1)}$. (Raising matrices to high powers is reasonably fast, and can be done in time proportional to the log of the power.) Then σ certainly has order dividing $q^{e'}-1$. If it has minimal polynomial of degree e' then we try and extend its action to compute ζ, as explained above. For at least half of the elements τ the degree of the minimal polynomial of σ will be equal to e' (this will be the case, for example, whenever τ has order q^f-1), so it should not be necessary to select many elements τ to get an appropriate σ. If E does have degree e' over F, then every such σ will extend to an element of the centraliser of M. On the other hand, if E has degree less than e' over F, then we shall not succeed in extending the action of σ to the full space to give a matrix which centralises every generator of G. Thus the test is conclusive in either case.

4 Testing for isomorphism

We turn now to the question of deciding isomorphism between two FG-modules M and M', at least one of which (M, say) has been proved irreducible. The method described here is essentially the same as that given by Parker in Section 6 of [8]. We assume that M' is defined by matrices $x'_1,\ldots,x'_r$ which give the actions on M' of the same group elements as the x_i do on M. Clearly we may assume that $\dim(M) = \dim(M')$. Before we start, we have to compute the centralising field E for M as described above. The test depends on the degree of the irreducible polynomial $p(x)$ being equal to $|E:F|$. In other words, the nullspace N of $p(\theta)$ must have dimension one over E. This is often true already but, if not, then we choose further random elements θ from A until we find one for which the characteristic polynomial $c(x)$ has an irreducible factor $p(x)$ with this property. (Since all of our probabilistic estimates in Section 2 were based on finding θ with this property, this search should not be too time consuming. Furthermore, we only need to consider irreducible factors of $c(x)$ of degree up to e.)

Assuming now that we have found $p(x)$ as above, we calculate the element θ' in the group algebra of M' that corresponds to θ, by computing the same linear sum in the generators x'_i as we did with the x_i to compute θ. Then we compute the characteristic polynomial of θ' and, if this is not equal to $c(x)$, we return the answer **false**. Otherwise, we compute the nullspace N' of $p(\theta')$ and, if this does not have dimension e over F, we can again return the answer **false**. Now, any isomorphism from M to M' must map N to N' and, since $\dim_E(N) = 1$, if there is such an isomorphism, then, for any nonzero elements n of N and n' of N', there must be an isomorphism that maps n to

n'. But we can test for this easily, by calculating the submodule L spanned by (n, n') in the direct sum $M \oplus M'$. There is an isomorphism if and only if L has dimension d over F, in which case the isomorphism ψ is defined explicitly, since $L = \{ (v, \psi(v)) \mid v \in M \}$.

5 Implementation issues and performance

For the GAP implemention, we introduced a record-type **GModule**. Any **GModule** must have components defining a field F, a dimension d, and a list of matrices x_i $(1 \leq i \leq r)$ that represent the images of the generators of a group G. Strictly speaking, we should also record a homomorphism φ from a specific group G with r generators g_i, where $\varphi(g_i) = x_i$ for $1 \leq i \leq r$, but, for reasons of efficiency and simplicity of coding, we do not insist that this map be specified explicitly. In fact in many examples $g_i = x_i$ and φ is the identity map.

After testing for irreducibility, there are two possible outcomes. If M turns out to be reducible, then a component is defined which contains a basis for a proper nontrivial submodule of M. Procedures can then be called, if required, to calculate the corresponding submodule and quotient module. If M is irreducible, then several record components are defined for M. Let θ be the element of the matrix algebra used in the irreducibility proof. One component contains θ itself, and another contains the formula by means of which θ is calculated from the original generators x_i. (This means that, since we do not keep the new generators of G, we have to record how the new generators are derived from the old, as described in Section 2.) We also have components for the characteristic polynomial $c(x)$ of θ, and the irreducible factor $p(x)$ of $c(x)$ that is used in the irreducibility proof. Finally, we record the dimension of the nullspace of $p(\theta)$ (which is equal to the degree of $p(x)$), and one vector from this nullspace.

After the absolute-irreducibility test, we introduce two new components, one containing the degree $e = |E : F|$ of the centralising field E over F, and (when $e > 1$) another containing a matrix which centralises each of the x_i and generates E as a field over F. If we wish to test modules for isomorphism, then it is essential that the nullspace of $p(\theta)$ should have dimension 1 over E; that is, dimension e over F. If this is not already the case, then we need to find a new element θ.

For the most part, our Meataxe implementation follows the algorithm described in Section 2 fairly closely. After experimentation, we made two minor adjustments, to improve the efficiency. Firstly, we do some preprocessing by adjoining a small number of new generators before we start. Particularly when r is small, this helps to prevent the first few elements θ chosen being too atypical to be useful. Secondly, for the i-th random element θ chosen, we give up at **Step 3** after considering factors $p(x)$ up to degree 2^i, and go on to a new element θ. This is because low degree factors are preferable for the later tests, and the evaluation of $p(\theta)$ for a high degree factor p can be quite time consuming. Of course, we do not know what the value of e is in advance, and no factor of degree less than e will work, so, if we wish the algorithms to be effective, we have to start considering higher degree factors eventually. The degree 2^i was chosen heuristically.

Below, we present a table of results, and timings. The times t_G (in seconds) are for our GAP implementation, and the times t_M are for the MAGMA implementation discussed briefly below. They should be viewed as a guideline only, because they can vary considerably with different runs on the same example. In fact, they are all averages over three runs but, in the case $GL(14,2)$, for example, the three times t_G were 8, 10 and 56 seconds. The GAP timings were on a Sun Sparcstation 10, and the MAGMA timings on a Solbourne machine (which is about half as fast as the Sun).

The bulk of the time is taken up by the calculations of the characteristic polynomial of θ, the nullspace of $p(\theta)$, and the spinning process. These calculations can all be made to run faster by coding them in 'C' (as in the MAGMA implementation) rather than in the GAP language. On the other hand, some of the times are artificially low, due to the fact that the given generators of the classical groups consist of very sparse matrices. The relatively large time t_G for $GL(200,10007)$ results from the fact that the speed of matrix operations in GAP seems to deteriorate considerably when the characteristic of the field grows large.

The column headed r/i means r for 'reducible' and i for 'irreducible'. The column headed $\#(\theta)$ gives the number of random elements θ that were considered (in the GAP run). This was equal to 1 for most trials on most examples. The examples $C_{19} \times Q_8$ have centralising field of degree 9 over F, and so, because of our policy of considering only factors $p(x)$ of degree up to 2^i for the i-th element θ, we could not hope to succeed with fewer than 4 such elements. The first of these has 8 mutually nonisomorphic irreducible composition factors of degree 9. The second is a direct sum of 10 isomorphic copies of an irreducible module of degree 18. (This example arose as a result of considering quotient groups of the Fibonacci group F(2,9).)

The example named C is cyclic of order $17^{40} - 1$, and the module is a direct sum of two irreducibles of dimension 20. The centralising field has dimension 20 over $F = GF(17^2)$. The final example, which did not complete, is an instance of the bad situation mentioned in Section 2, and the module has two isomorphic composition factors.

The original Meataxe has no hope at all of proving irreducibility for the examples $C_{19} \times Q_8$ and C, and for $GL(200,10007)$ it would take a long time.

All of these algorithms have also been implemented in MAGMA by Allan

G	q	d	r/i	$\#(\theta)$	t_G	t_M
M_{11}	7	44	i	1	1	
M_{11}	7	44	i	2	2	
M_{12}	7	54	i	1	2	
M_{12}	7	54	i	1	2	
$C_{19} \times Q_8$	5	72	r	4	26	
$L_2(81)$	41	82	i	1	8	
He	2	102	i	1	10	
He	2	102	i	2	30	
Ly	5	111	i	1	22	
M_{22}	7	154	i	1	50	6.3
$C_{19} \times Q_8$	5	180	r	4-5	38	5.7
$GL(14, 2)$	2	196	r	1	24	0.7
$GU(200, 9)$	9	200	i	1	54	3.6
$GL(200, 10007)$	10007	200	i	1	110	2.8
C	289	40	r	5	34	
$2^{52}.SL(2, 2^{13})$	2	52	r	?	?	?

Steel at the University of Sydney, together with some related procedures such as finding all minimal submodules of a module for a finite group over a finite field. Except for the case $q = 2$, they have been found to run faster than the original Meataxe algorithm. In particular, if many of the irreducible constituents of M are not absolutely irreducible, then they perform much more reliably in general. This situation occurs frequently in the analysis of finite soluble groups, where the algorithm can be applied to elementary abelian sections.

The last example in the table above did eventually complete in MAGMA and produce a proper submodule; the number of elements θ considered was about 5000.

References

[1] M. Aschbacher, 'On the maximal subgroups of the finite classical groups', Invent. Math. 76 (1984), 469-514.

[2] W. Bosma and J. Cannon, MAGMA *Handbook*, Sydney, 1993.

[3] D.G.Cantor and H. Zassenhaus, 'A new algorithm for factoring polynomials over finite fields', Math. of Computation 36, (1981), 587-592.

[4] C.W. Curtis and I. Reiner, *Representation Theory of Finite Groups and Associative Algebras*, John Wiley and Sons, 1962 (reprinted 1988).

[5] D.F. Holt and Sarah Rees, 'An implementation of the Neumann-Praeger algorithm for the recognition of special linear groups', Experimental Mathematics 1 (1993), 237-242.

[6] Donald E. Knuth, *The Art of Computer Programming (Vol. 2): Seminumerical Algorithms*, 2nd edn., Addison Wesley, Reading Mass., 1981.

[7] C.R. Leedham-Green, 'Generating random group elements', preprint.

[8] R. Parker, 'The computer calculation of modular characters. (The Meat-Axe)', in M. Atkinson (ed.), *Computational Group Theory*, Academic Press, London, 267-74, 1984.

[9] Lajos Rónyai, Computing the Structure of Finite Algebras, J. Symbolic Computation 9 (1990), 355-373.

[10] M. Schönert et. al., 'GAP – *Groups, Algorithms, and Programming*', Lehrstuhl D für Mathematik, RWTH Aachen, Germany, first edition, 1992.

Addresses:

Mathematics Institute
University of Warwick
Coventry CV4 7AL
Great Britain
e-mail: dfh@maths.warwick.ac.uk

Department of Mathematics and Statistics
University of Newcastle
Newcastle-upon-Tyne NE1 7RU
Great Britain
e-mail: sarah.rees@newcastle.ac.uk