Treating the Exceptional Cases of the MeatAxe

Gábor Ivanyos and Klaus Lux

CONTENTS

We show that the Holt–Rees extension of the standard MeatAxe procedure finds submodules of modules over finite algebras with positive probability in more cases than originally claimed. For the case when the Holt–Rees method fails we propose a further, but still simple and efficient extension.

1. INTRODUCTION

Finding the irreducible composition factors of a finite module M for a finite dimensional associative algebra A over a finite field F is one of the fundamental tasks in computational modular representation theory. The most commonly used practical approach to this problem is the MeatAxe algorithm [Parker 1984], which solves the problem of proving constructively that M is irreducible. Originally, the method did not perform satisfactorily when the ground field F is large. Holt and Rees [1994] have proposed an extension to Parker's method based on factoring the characteristic polynomial of random elements from A. They provided an accurate analysis and showed that their approach proves efficiently that a given module is irreducible regardless of the size of the ground field. Furthermore, in most cases they also have a definite chance of finding a nontrivial submodule. In this note we prove that the extension works in more cases than claimed in [Holt and Rees 1994]. However, there is still one type of module where the algorithm definitely fails; we propose a method for this case. The implementation of M. Ringe as part of the C-MeatAxe shows that our algorithm is also practically feasible. The reader interested in the theoretical complexity of related problems is referred to the survey [Rónyai 1993].

We restrict our attention to a fundamental subtask which can be interpreted as an effective version of testing irreducibility. Let $F = \mathrm{GF}(q)$ be the finite field consisting of q elements. We assume, without

The paper was written while Ivanyos was a visitor at the department of Mathematics and Computing Science, Eindhoven University of Technology, supported by NWO-OTKA Grant N26673. The research was also supported by FKFP Grant 0612/1997, OTKA Grants 016503, 022925, and EC Grant ALTEC-KIT.

loss of generality, that A contains an identity element denoted by 1_A or by 1 for short; that the module is $M = F^d$, the space of column vectors of length d; and that the action of A on M is faithful and is given in terms of matrices for generators of A. (In contrast to the MeatAxe implementation in GAP [Schönert et al. 1996], which is based on row vectors and right action, our discussion is presented in terms of column vectors and left action.) The procedure either concludes that M is an irreducible A-module or returns a nontrivial submodule of M. Furthermore, we identify A with its image in $\mathcal{M}_d(F)$. Finally, we assume that we are provided with an auxiliary procedure which generates random elements of A (independently and uniformly).

The paper is structured as follows. In Section 2 we briefly comment on the original algorithm proposed in [Holt and Rees 1994] in order to extend the probability analysis to slightly more cases and to describe a class of algebras A which contains all the situations where the method fails. The algorithm for this class of algebras is outlined in Section 3. The probability of success will be estimated in Section 4. Finally, in Section 5 we provide some experimental results with the C-MeatAxe implementation of the algorithm.

For the standard notions and facts related to finite dimensional algebras and modules, see [Pierce 1982], for example. We adopt the following conventions. Modules are assumed to be left modules. If H is a subset of an algebra A and K is a subset of an A-module (which can be A itself) then by HK we denote the linear span of all the products θv where $\theta \in H$ and $v \in K$. Also, if θ and η are two elements of the algebra A then $[\theta, \eta]$ stands for the additive commutator $\theta\eta - \eta\theta$. Again, if H and K are two subsets of A then $[H, K]$ denotes the linear span of all the commutators $[\theta, \eta]$ (with $\theta \in H$ and $\eta \in K$). By $\operatorname{Rad} A$ and by $\operatorname{Rad} M$ we denote the (Jacobson) radical of the algebra A and the module M. $\operatorname{Rad}^2 M$ stands for the iterated radical $\operatorname{Rad}\operatorname{Rad} M$. We also use the standard notation $\mathrm{C}_H(K)$ and $\mathrm{Z}(H)$ for the centralizer of the subset K in a subalgebra H and the center of H, respectively.

We assume that A is a fixed algebra. It will be convenient to introduce some additional notation. By the Wedderburn–Malcev principal theorem [Pierce 1982, Section 11.6], A can be written as

$$A = S + \operatorname{Rad} A, \quad \text{where } S \cong A/\operatorname{Rad} A.$$

Since the complementary subalgebra S is unique up to a conjugation by an inner automorphism of A, we can speak about the structural properties of A in terms of S even if S is not specified explicitly.

2. THE EXCEPTIONAL ALGEBRAS

In [Holt and Rees 1994], the extension of MeatAxe is proved to succeed in constructing a nontrivial submodule with probability at least 0.144 in many cases. In particular, it recognizes irreducible modules, finds a nontrivial submodule if $M/\operatorname{Rad} M$ is decomposable or M contains non-isomorphic composition factors. The submodule is generated from the kernel of $p(\theta)$, where θ is a random element and $p(x)$ is an appropriate irreducible factor of the characteristic polynomial of θ on M (see Lemma 2.1 below). The probability analysis of success is based on the following observation, which will be useful in the analysis of the present paper as well.

Lemma 2.1. *Let W be an irreducible A-module and $E = \operatorname{End}_A(W)$, the algebra consisting of the A-endomorphisms of the module W. Then for at least 21.4% of the elements $\theta \in A$ the characteristic polynomial over F of θ on the module W has an unrepeated irreducible factor of degree $\dim_F E$.*

Proof. By Schur's lemma and Wedderburn's theorem on finite division algebras, E is a finite extension field of F. Note also that if W as an E-module is isomorphic to E^n and $I = \{\theta \in A \mid \theta W = (0)\}$ is the annihilator ideal of W, then $A/I \cong \mathcal{M}_n(E)$. Since uniform selection of elements in A corresponds to uniform selection in the factor A/I, we may assume throughout the proof that $I = (0)$ and identify A with $\mathcal{M}_n(E)$. The statement for the case $E = F$ is proved in [Holt and Rees 1994] (with a somewhat bigger constant), therefore we may restrict ourselves to the case $e = \dim_F E > 1$. The argument given in by Holt and Rees for this case appears to contain a minor mistake, therefore we give a corrected proof below.

The condition is equivalent to saying that θ, considered as a matrix over E, has an unrepeated eigenvalue λ such that λ is not contained in any proper subfield E' with $F \leq E' < E$ and for every automorphism $\sigma \in \operatorname{Gal}(E|F)$ such that $\lambda^\sigma \neq \lambda$, λ^σ

is not an eigenvalue of θ. (This follows from the fact that the characteristic polynomial of θ over F is $\prod_{\sigma\in\mathrm{Gal}(E|F)} c(x)^\sigma$, where $c(x) \in E[x]$ is the characteristic polynomial of θ, regarded as a matrix over E. See [Reiner 1975, Theorem 9.10 and Exercise 9.4], for example.)

Note that at most half of the elements in E can be contained in a proper subfield of E. This establishes the case $n = 1$. For the rest of the proof we assume $n > 1$.

Let $F = \mathrm{GF}(q)$ and $E = \mathrm{GF}(q^e)$. Following the arguments given in [Holt and Rees 1994], let H denote the number of matrices $\theta \in \mathcal{M}_n(E)$ such that a specific $\lambda \in E$ is an unrepeated eigenvalue of θ. Also, let H' stand for the number of matrices with two distinct specific unrepeated eigenvalues $\lambda, \mu \in E$. In [Holt and Rees 1994] it is shown that H and H' are independent of the particular choice of λ and μ, and

$$H = \frac{1}{q^e-1}\prod_{i=0}^{n-1}(q^{en}-q^{ei}) \quad \text{and} \quad H' \le \frac{H}{q^e-1}.$$

Let R denote the set of elements $\lambda \in E$ such that λ has exactly e conjugates over $\mathrm{Gal}(E|F)$ and let $r = |R|$. By inclusion-exclusion, at least

$$rH - \binom{r}{2}H' \ge \left(r - \frac{r(r-1)}{2(q^e-1)}\right)H$$

matrices have some unrepeated eigenvalue from R. For the number of matrices having at least two eigenvalues from some orbit of $\mathrm{Gal}(E|F)$ on R we have the crude upper bound

$$\frac{r}{e}\binom{e}{2}H' \le \frac{r}{e}\binom{e}{2}\frac{H}{q^e-1}.$$

Hence the number of matrices with the required property is at least

$$\begin{aligned}\left(r - \frac{r(r-1)+r(e-1)}{2(q^e-1)}\right)H &= \left(1 - \frac{r+e-2}{2(q^e-1)}\right)rH \\ &\ge \left(1 - \frac{q^e+e-4}{2(q^e-1)}\right)\frac{q^e}{2}H \\ &= \left(\frac{1}{2} - \frac{e-3}{2(q^e-1)}\right)\frac{q^e}{2}H \\ &\ge \tfrac{7}{30}q^eH.\end{aligned}$$

The first inequality follows from

$$q^e/2 \le r \le q^e - q \le q^e - 2,$$

while the second follows from the fact that the maximal value of $(e-3)/(2q^e-2)$ for the integers $q, e \ge 2$ is $\frac{1}{30}$ (taken at $q = 2$, $e = 4$). Hence the proportion of such matrices is at least

$$\begin{aligned}\frac{7}{30}q^eH/q^{en^2} &= \frac{7}{30}\prod_{i=2}^{n}(1-q^{-ei}) \\ &\ge \frac{7}{30}\prod_{i=2}^{\infty}(1-4^{-i}) \ge 0.214. \qquad \square\end{aligned}$$

Remark. The mere assumption that θ, regarded as a matrix over E contains an unrepeated eigenvalue λ which is not contained in any proper subfield [Holt and Rees 1994] appears to be insufficient even for the purposes of the MeatAxe. Indeed, if an algebraic conjugate λ' of λ, different from λ, is also an eigenvalue of θ, then the characteristic polynomial of θ over F contains the minimal polynomial $p(x)$ of λ at least twice and therefore the dimension of the kernel of $p(\theta)$ over E is at least 2.

The only possible situations when the Holt–Rees extension of MeatAxe may fail are modules M such that $\operatorname{Rad} M \ne (0)$, $M/\operatorname{Rad} M$ is irreducible and all the composition factors are isomorphic to $M/\operatorname{Rad} M$. Since M is faithful, this implies that every irreducible A-module is isomorphic to $M/\operatorname{Rad} M$. Let $E = \mathrm{End}_A(M/\operatorname{Rad} M)$, as in Lemma 2.1. Then E is a finite extension field of F and $M/\operatorname{Rad} M$ is isomorphic to E^n as an S-module for some integer n, where S is a subalgebra of A isomorphic to $A/\operatorname{Rad} A$. Note that the multiplicity of E^n in M is d/en, where $e = \dim_F E$ and $S \cong \mathcal{M}_n(E)$. The center of S is therefore isomorphic to E. We may and shall identify E with $\mathrm{Z}(S)$. In summary:

$$\begin{cases}\operatorname{Rad} A \ne (0), \quad S \cong \mathcal{M}_n(E), \\ E = \mathrm{Z}(S) \text{ is an extension field of } F.\end{cases} \tag{2–1}$$

The Holt–Rees extension of the MeatAxe is shown to succeed even in this case provided that $E = F$. We extend the proof given in [Holt and Rees 1994] to the more general case where $E \le \mathrm{Z}(A)$.

Proposition 2.2. *Assume that* (2–1) *holds*, $M/\operatorname{Rad} M$ *is irreducible and* $E \le \mathrm{Z}(A)$. *Then, for at least* 14.4% *of the elements* θ *in* A, *there exists a factor* $p(x) \in F[x]$ *of the characteristic polynomial of* θ *on* M *such that the kernel of* $p(\theta)$ *is a nonzero subspace of* $\operatorname{Rad} M$.

Proof. The case $E = F$ is proved in [Holt and Rees 1994]. Assume that $E > F$. Note that every element $\theta \in A$ can be uniquely written in the form $\theta = \theta_0 + \theta_1$ where $\theta_0 \in S$ and $\theta_1 \in \operatorname{Rad} A$. Assume that the characteristic polynomial (over F) of $\theta_0 \in S$ on the irreducible S-module $M/\operatorname{Rad} M \cong E^n$ has an unrepeated irreducible factor $p(x) \in F[x]$ of degree $e = \dim_F E$. By Lemma 2.1, this is the case for at least 21.4% of the possible choices for θ_0. Let $\lambda_1, \dots, \lambda_e$ be the roots of $p(x)$ in E. Then there exists an element $\lambda \in \{\lambda_1, \dots, \lambda_e\}$, say $\lambda = \lambda_1$, such that the kernel of $\theta_0 - \lambda$ is an E-submodule of $M/\operatorname{Rad} M$ of rank 1 (i.e., a one dimensional E-linear subspace). Furthermore, $\theta_0 - \lambda_i$ is a unit in S for $i = 2, \dots, e$.

Obviously, for every $\theta_1 \in \operatorname{Rad} A$, the kernel of $\theta_0 + \theta_1 - \lambda$ in M is nonzero, since the quotient map on $M/\operatorname{Rad} M$ is $\theta_0 - \lambda$. Let L stand for the set consisting of $\theta_1 \in \operatorname{Rad} A$ for which this kernel is not contained in $\operatorname{Rad} M$. We claim that L is contained in a proper E-submodule of $\operatorname{Rad} A$. To this end consider M as an S-module. Since S is a simple algebra there exists an S-submodule M_0 complementary to $\operatorname{Rad} M$. Then M_0, as an S-module, is isomorphic to $M/\operatorname{Rad} M$. In particular, there exists a nonzero element $v \in M_0$ such that $(\theta_0 - \lambda)v = 0$. Then for every element $\theta_1 \in \operatorname{Rad} A$, the kernel of $\theta_0 + \theta_1 - \lambda$ is contained in the E-submodule $Ev + \operatorname{Rad} M$. Assume now that $\theta_1 \in L$, i.e., this kernel contains an element $u \in M \setminus \operatorname{Rad} M$. Then $u = \varepsilon v + w$ for some unit $\varepsilon \in E$ and some element $w \in \operatorname{Rad} M$. Multiplying by ε^{-1}, we may assume that $u = v + w$ with $w \in \operatorname{Rad} M$. Now

$$0 = (\theta_0 + \theta_1 - \lambda)(v + w) = \theta_1 v + (\theta_0 - \lambda)w + \theta_1 w,$$

and hence $\theta_1 v = -(\theta_0 - \lambda)w - \theta_1 w$ is in

$$\begin{aligned}(\theta_0 - \lambda)&\operatorname{Rad} M + \operatorname{Rad} A \operatorname{Rad} M \\ &= (\theta_0 - \lambda)\operatorname{Rad} M + \operatorname{Rad}^2 M.\end{aligned}$$

Thus

$$\begin{aligned}L &\subseteq L' \\ &= \{\theta_1 \in \operatorname{Rad} A \mid \theta_1 v \in (\theta_0 - \lambda)\operatorname{Rad} M + \operatorname{Rad}^2 M\}.\end{aligned}$$

Obviously L' is an E-submodule of $\operatorname{Rad} A$. Assume that $L' = \operatorname{Rad} A$. Then

$$\begin{aligned}\operatorname{Rad} M &= \operatorname{Rad}(Av) = \operatorname{Rad} Av \\ &= L'v \subseteq (\theta_0 - \lambda)\operatorname{Rad} M + \operatorname{Rad}^2 M.\end{aligned}$$

(Here the first equality holds because of $M = Av + \operatorname{Rad} M$ and Nakayama's lemma.) From this we infer that $\theta - \lambda$ acts surjectively on the factor module $\operatorname{Rad} M/\operatorname{Rad}^2 M$, and hence on its composition factors as well. Since all these composition factors are isomorphic to $M/\operatorname{Rad} M$, this is a contradiction to the fact that $\theta - \lambda$ is singular on $M/\operatorname{Rad} M$. Thus L is included in the proper E-submodule L' of $\operatorname{Rad} A$, as claimed.

By the claim, for at least $1 - 1/|E|$ of the possible choices for θ_1, the kernel of $\theta - \lambda = \theta_0 + \theta_1 - \lambda$ is a subspace of $\operatorname{Rad} M$. Let $\rho = \prod_{i=2}^{e}(\theta - \lambda_i)$. Then ρ is a unit modulo $\operatorname{Rad} A$ and hence ρ itself is a unit in A. Therefore the kernel of $(\theta - \lambda)\rho = p(\theta)$ is equal to the kernel of $\theta - \lambda$. Thus, the kernel of $p(\theta)$ is a nonzero subspace of $\operatorname{Rad} M$ provided that the kernel of $\theta - \lambda$ is. As the components θ_0 and θ_1 of θ are chosen independently, this gives $0.214(1 - 1/|E|) \geq 0.214 \cdot 3/4 > 0.16$, so at least $16\% > 14.4\%$ of the elements $\theta \in A$ satisfy the desired property. $\square$

This means that the Holt–Rees extension of Meat-Axe succeeds with probability at least 0.144 in this case. Hence we can restrict our attention to the case where E is not central, i.e., algebras A satisfying (2–1) and the additional hypothesis

$$[A, E] > (0). \tag{2–2}$$

3. THE ALGORITHM

We propose the method described below for treating algebras with properties (2–1) and (2–2). As described in the last section, the algorithm has been successfully incorporated into the program chop by M. Ringe, which is part of the C-MeatAxe version 2.3.

We assume that a random element $\theta \in A$ is selected and that the irreducible factors of the characteristic polynomial $c(x)$ of θ over F are computed. Note that these computations are carried out as a part of the original algorithm described in [Holt and Rees 1994]. We select a factor $p(x)$ of minimum degree among the factors of $c(x)$ of minimum multiplicity and do the following.

(i) Determine the polynomial $i(x)$, a representative of the primitive idempotent of the algebra

$$F[x]/(c(x))$$

corresponding to the factor $p(x)$. More precisely, by the Chinese Remainder Theorem,

$$F[x]/(c(x)) \cong F[x]/(p^l(x)) \oplus F[x]/(q(x)),$$

where l is the multiplicity of $p(x)$ in $c(x)$ and $q(x) = c(x)/p^l(x)$ and we want the identity element of the component isomorphic to

$$F[x]/(p^l(x)).$$

To be explicit, 1 can be expressed using the extended Euclidean algorithm in the form

$$1 = a(x)p^l(x) + b(x)q(x),$$

with polynomials $a(x)$ and $b(x)$. Then $i(x) \equiv b(x)q(x) \pmod{c(x)}$.

(ii) Choose another random element $\eta \in A$ as well as a random vector $v \in M$ and calculate the submodule N generated by $[\theta, i(\theta)\eta i(\theta)]v$. If this is a proper nonzero submodule then return N, otherwise report failure.

We make comments only on the costs of steps which are additional to the Holt–Rees extension of the MeatAxe procedure. The polynomial $i(x)$ can be determined with $O(d^2)$ operations in F (Note that l is less than d). The cost of computing the vector $[\theta, i(\theta)\eta i(\theta)]v$ is $O(d^3)$ arithmetical operations assuming that we use a method based on performing $O(d)$ matrix-by-vector multiplications. Using a method based on fast calculation of Krylov sequences [Bini and Pan 1994] the cost can be reduced to $O(\mathrm{MM}(d)\log d)$ operations, where $\mathrm{MM}(d)$ stands for the number of arithmetic steps required to multiply two d by d matrices. We remark that Eberly and Giesbrecht [1996, Lemma 3.1] give an efficient algorithm to compute all the primitive idempotents of the subalgebra generated by θ simultaneously in explicit matrix form. The method is based on computing the rational canonical form of θ [Giesbrecht 1995], and the running time is essentially $O(\mathrm{MM}(d)\log d)$.

Thus the total number of arithmetical steps required by the algorithm is dominated by the cost of computing the submodule N in step (ii), which is $O(d^3)$, provided that the number of generators of A is fixed.

4. PROBABILITY OF SUCCESS

We now give an estimate for the probability of finding a proper submodule in the situation where the algebra A satisfies conditions (2–1) and (2–2). Actually we show that the commutator $[\theta, i(\theta)\eta i(\theta)]$ has a positive chance for being a nonzero element of $\mathrm{Rad}\, A$.

Lemma 4.1. *Assume that the finite dimensional F-algebra A with identity satisfies conditions* (2–1) *and* (2–2). *Let ι be an idempotent of S. Then*

(a) $[\iota E\iota, \iota A\iota] = \iota[E, A]\iota$,
(b) $S(\iota[E, A]\iota)S = [E, A]$, *and*
(c) $(0) \subset [\iota E\iota, \iota A\iota] \subseteq \mathrm{Rad}\, A$.

Proof. First we note that since ι commutes with E, $\iota\varepsilon\iota = \iota\varepsilon = \varepsilon\iota$ for every $\varepsilon \in E$ and hence $\iota E\iota = \iota E = E\iota$. Part (a) is immediate from the following equalities which hold for every $\varepsilon \in E$ and $\alpha \in A$.

$$\begin{aligned}\iota\varepsilon \cdot \iota\alpha\iota - \iota\alpha\iota \cdot \iota\varepsilon &= \iota\varepsilon\iota \cdot \alpha\iota - \iota\alpha \cdot \iota\iota\varepsilon \\ &= \iota\varepsilon \cdot \alpha\iota - \iota\alpha \cdot \varepsilon\iota = \iota(\varepsilon\alpha - \alpha\varepsilon)\iota.\end{aligned}$$

To prove part (b), let $\sigma, \tau \in S$, $\varepsilon \in E$, $\alpha \in A$. Then

$$\begin{aligned}\sigma\iota[\varepsilon, \alpha]\iota\tau &= \sigma\iota\varepsilon\alpha\iota\tau - \sigma\iota\alpha\varepsilon\iota\tau \\ &= \varepsilon\sigma\iota\alpha\iota\tau - \sigma\iota\alpha\iota\tau\varepsilon = [\varepsilon, \sigma\iota\alpha\iota\tau],\end{aligned}$$

where the second equality holds because ε commutes with the elements $\iota, \sigma, \tau \in S$. From this we infer that $S\iota[E, A]\iota S = [E, S\iota A\iota S]$. It remains to establish the equality $S\iota A\iota S = A$. To this end observe that $S\iota S$ is a nonzero ideal in the simple algebra S, therefore $S\iota S = S$. Hence $S\iota A\iota S = S\iota SAS\iota S = SAS = A$. (The first and the last equalities are obvious because S contains 1_A.)

Part (c) follows from (a) and (b) and the fact that E is central modulo $\mathrm{Rad}\, A$. $\square$

After these preparations we are ready to give a lower bound on the probability of success of the algorithm.

Proposition 4.2. *Assume that the matrix algebra $A \leq \mathcal{M}_d(F)$ satisfies conditions* (2–1) *and* (2–2). *Then the proportion of the triples $(\theta, \eta, v) \in A \times A \times F^d$ for which the algorithm described in the preceding section finds a proper submodule is at least* 0.08.

Proof. Assume that $p(x)$ is an unrepeated irreducible factor of the characteristic polynomial of $\theta + \mathrm{Rad}\, A$ on E^n. Then the degree of $p(x)$ is the dimension (over F) of the kernel of $p(\theta + \mathrm{Rad}\, A)$. This subspace is obviously a $\mathrm{Z}(A/\mathrm{Rad}\,A)$-submodule of E^n, and hence the degree of $p(x)$ is at least

$$e = \dim_F \mathrm{Z}(A/\mathrm{Rad}\,A) = \dim_F E.$$

Assume that the degree of $p(x)$ is exactly e. By Lemma 2.1, such a factor does exist for at least 21.4% of the elements $\theta \in A$. Furthermore, all the factors of this kind are characterized as the minimum degree factors amongst the factors of minimal multiplicity of the characteristic polynomial of θ on the whole module M.

Referring to the homomorphism $F[x]/(c(x)) \to A$ induced by $x \mapsto \theta$, it is immediate that $\iota = i(\theta)$ is an idempotent. Let $\bar{\theta} = \theta + \mathrm{Rad}\,A$ and $\bar{\iota} = \iota + \mathrm{Rad}\,A$. Furthermore, the characteristic polynomial of $\bar{\iota}\bar{\theta}$ on E^n is $p(x)x^{(n-1)e}$. It follows that $\bar{\iota}\bar{\theta}$ and $\bar{\iota}$ have rank e, therefore $\bar{\iota}$ is a primitive idempotent of $A/\mathrm{Rad}\,A$. Hence $\bar{\iota}(A/\mathrm{Rad}\,A)\bar{\iota} = \bar{\iota}\mathrm{Z}(A/\mathrm{Rad}\,A)\bar{\iota}$. In particular, $\bar{\iota}\bar{\theta} \in \bar{\iota}\mathrm{Z}(A/\mathrm{Rad}\,A)$. On the other hand, the minimum polynomial of $\bar{\iota}\bar{\theta}$ on $\bar{\iota}Z(A/\mathrm{Rad}\,A)$ is of degree e, therefore $\bar{\iota}\bar{\theta}$ generates the whole $\bar{\iota}\mathrm{Z}(A/\mathrm{Rad}\,A)$.

Now $[\theta, \iota\eta\iota]$ is a nonzero element of $\mathrm{Rad}\,A$ for at least $1 - 1/|E| \geq \frac{3}{4}$ of the elements $\eta \in A$, see Lemma 4.3 below, and let us assume in the following that this is the case. Then $[\theta, \iota\eta\iota]$ is a nontrivial F-linear transformation and hence the kernel has codimension at least 1. Therefore for at least $1 - 1/|F| \geq \frac{1}{2}$ of the elements $v \in M$ the vector $[\theta, \iota\eta\iota]v$ is a nonzero element of the proper submodule $\mathrm{Rad}\,AM = \mathrm{Rad}\,M$. Putting the bounds together, the algorithm finds a proper submodule with probability at least $0.214 \cdot \frac{3}{4} \cdot \frac{1}{2} > 0.08$. □

The proposed method, complemented with the Holt-Rees approach, gives an algorithm of Las Vegas type for every case.

We now give the promised proof of the statement used above.

Lemma 4.3. *Assume that the finite dimensional F-algebra A with identity satisfies conditions* (2–1) *and* (2–2). *Assume further that θ is an element of A and ι is an idempotent of the subalgebra of A generated by θ and 1_A such that the subalgebra of $A/\mathrm{Rad}\,A$ generated by $\iota\theta + \mathrm{Rad}\,A$ is $(\iota + \mathrm{Rad}\,A)\mathrm{Z}(A/\mathrm{Rad}\,A)$. Then $[\theta, \iota A\iota] \subseteq \mathrm{Rad}\,A$ and $[\theta, \iota\eta\iota] \neq 0$ for at least $1 - 1/|E|$ of the elements $\eta \in A$.*

Proof. Let A_θ denote the subalgebra of A generated by $\iota\theta$. We first note that ι is the identity element of A_θ. Indeed, $\iota\alpha = \alpha\iota = \alpha$ holds for every element $\alpha \in A_\theta$. On the other hand, it is straightforward to see that $A'_\theta = A_\theta + F\iota$ is a subalgebra and A_θ is an ideal of A'_θ. By the assumption

$$\begin{aligned}(A'_\theta + \mathrm{Rad}\,A)/\,\mathrm{Rad}\,A &= (A_\theta + \mathrm{Rad}\,A)/\,\mathrm{Rad}\,A \\ &\cong Z(A/\mathrm{Rad}\,A);\end{aligned}$$

thus A'_θ is a local algebra and A_θ is not a nilpotent ideal. But since in a local algebra every proper ideal is contained in the radical, $A_\theta = A'_\theta$, establishing the containment $\iota \in A_\theta$.

We will now replace S and E with appropriate conjugates in order to achieve the situation where $\iota \in S$ and ιE is a subalgebra of A_θ. By the Wedderburn–Malcev principal theorem, $A_\theta = S_\theta + \mathrm{Rad}\,A_\theta$, where S_θ is a semisimple subalgebra of A_θ. Since every maximal semisimple subalgebra of A is a conjugate of S by an inner automorphism [Malcev 1942], there exists a unit $\sigma \in A$ such that $S^\sigma = \sigma^{-1}S\sigma \geq S_\theta$. Because conditions (2–1) and (2–2) are invariant under automorphisms, we may replace S with S^σ and E with E^σ, or, equivalently, assume $S_\theta \leq S$. Note that ι is just the identity element of S_θ.

By the assumption,

$$(A_\theta + \mathrm{Rad}\,A)/\mathrm{Rad}\,A = (\iota + \mathrm{Rad}\,A)\mathrm{Z}(A/\mathrm{Rad}\,A) \cong E$$

is a simple algebra; therefore $\mathrm{Rad}(A_\theta + \mathrm{Rad}\,A) = \mathrm{Rad}\,A$. On the other hand, $\mathrm{Rad}\,A_\theta + \mathrm{Rad}\,A$ is obviously a nilpotent ideal of $A_\theta + \mathrm{Rad}\,A$. It follows that $\mathrm{Rad}\,A_\theta \leq \mathrm{Rad}\,A$, $A_\theta + \mathrm{Rad}\,A = S_\theta + \mathrm{Rad}\,A$ and $S_\theta = \iota E$.

Observe that, since the idempotent ι commutes with θ, for every $\eta \in A$ we have

$$\begin{aligned}[\theta, \iota\eta\iota] &= \theta\iota\eta\iota - \iota\eta\iota\theta = \theta\iota\iota\eta\iota - \iota\eta\iota\iota\theta \\ &= \iota\theta\iota\eta\iota - \iota\eta\iota\iota\theta = [\iota\theta, \iota\eta\iota].\end{aligned}$$

The equality $\iota E = S_\theta$ and the preceding lemma give $[S_\theta, \iota A\iota] \subseteq \mathrm{Rad}\,A$. Since $S_\theta \leq A_\theta \leq S_\theta + \mathrm{Rad}\,A$, we have

$$[A_\theta, \iota A\iota] \subseteq [S_\theta, \iota A\iota] + \mathrm{Rad}\,A \subseteq \mathrm{Rad}\,A.$$

The first inclusion of the formula above holds because $\mathrm{Rad}\,A$ is a two-sided ideal and hence

$$[\iota A\iota, \mathrm{Rad}\,A] \subseteq \mathrm{Rad}\,A.$$

In particular, $[\theta, \iota A\iota] = [\iota\theta, \iota A\iota] \subseteq \mathrm{Rad}\, A$. So we have proved the first part of the statement.

In order to see the second part, notice that, since $S_\theta \leq A_\theta$,

$$\mathrm{C}_{\iota A\iota}(\theta) = \mathrm{C}_{\iota A\iota}(A_\theta) \leq \mathrm{C}_{\iota A\iota}(S_\theta) < \iota A\iota.$$

The latter inclusion is strict because not the whole $\iota A\iota$ commutes with $S_\theta = \iota E$ by Lemma 4.1. Obviously, $\mathrm{C}_{\iota A\iota}(S_\theta)$ is an S_θ-submodule of $\iota A\iota$ (multiplication by elements from S_θ from the left hand side). The set of elements η such that $[\theta, \iota\eta\iota] = 0$ is the F-linear subspace $(1_A - \iota)A + A(1_A - \iota) + \mathrm{C}_{\iota A\iota}(\theta)$. By the preceding argument the codimension of this subspace is at least $\dim_F S_\theta = \dim_F E$, whence the second part of the assertion follows. □

5. EXPERIMENTAL RESULTS

We conclude with a running time comparison for two versions of the composition factor program chop in the C-MeatAxe written by M. Ringe. For the details of the algorithm used for determining the composition factors of a given representation, see for example [Lux 1997]. For our purposes the following rough outline of the algorithm is sufficient: the main subroutine of chop takes as input a sequence of matrices for the generators of the algebra A. Its aim is to find a proper A-invariant subspace or to prove that there is no such subspace. If it finds an A-invariant subspace, it determines matrices for the generators on the subspace and the quotient space and calls itself with the matrices obtained. The search for the invariant subspace is done by looking for the kernel of words in the matrices for the generators using the Holt–Rees approach; see [Holt and Rees 1994]. In the old implementation, if a word θ did not lead to splitting of the given representation and irreducibility could not be shown in reasonable time using Norton's lemma, then the algorithm would just take the next word. In the new implementation, before taking the next word, we check for the exceptional case. This is done as follows:

For all factors $p(x)$ of the characteristic polynomial of θ of degree at least two, do the following: determine $i(x)$ as defined above. Then choose a second word η and random vector v. Furthermore, determine the vector $[\theta, i(\theta)\eta i(\theta)]v$. If it is nonzero, check whether it lies in a proper invariant subspace. If it does, call the main subroutine with the two new representations obtained.

This method (actually, an implementation based on an earlier version of the paper) differs from the algorithm described in Section 3 in the sense that just one factor is selected in order to assure that the algorithm *never* performs more than roughly $O(d^3)$ operations. As the number of factors of the characteristic polynomial is usually small (probably around $O(\log n)$, where n is the number given in (2–1)), the exhaustive search given here does not cause too much loss of efficiency in practice. Furthermore, the algorithm succeeds usually with the first factor with a probability much higher than the modest estimate given in Section 4.

We now compare the running times of the old and the new version of chop. The new version is part of the MeatAxe 2.3.2 release and the old is part of the MeatAxe release 2.2 as delivered with GAP 3.4.

In order to test the two programs we proceed as follows. We first construct a reducible representation for a finite group that has two isomorphic composition factors with a large endomorphism ring. This is done using GAP [Schönert et al. 1996].

We first take the two generators A, B for $\mathrm{SL}(n, F_q)$, $q = p^e$ with p prime, as produced by the command SpecialLinearGroup in GAP. We then form the matrices

$$a_1 = \begin{pmatrix} A & I \\ 0 & A^\Phi \end{pmatrix}$$

and

$$b_1 = \begin{pmatrix} B & I \\ 0 & B^\Phi \end{pmatrix}.$$

Here I is the n by n identity matrix over F_q and Φ denotes the Frobenius automorphism of F_q mapping $x \in F_q$ to x^p. Note that Φ is applied to the entries of A, B. The GAP command BlowupSQ is used to perform the Galois descent, i.e. to replace all entries in the matrices by the corresponding matrices in the regular representation of F_q over F_p. In this way we construct two $2en$ by $2en$ matrices a_1 and b_1 over F_p.

We then proceed by conjugating a_1 and b_1 by a random invertible matrix produced by GAP, the resulting matrices are a and b. Furthermore, let

$$x(a, b) = ababbabababbabb,$$

$$y(a, b) = ababbabbabababb$$

be given words in a and b. We test the programs on the representation generated by the pairs of matrices $(x^i(a,b), y^i(a,b))$, where

$$x^0(a,b) = a, \qquad y^0(a,b) = b,$$
$$x^i(a,b) = x(x^{i-1}(a,b)), \qquad y^i(a,b) = y(y^{i-1}(a,b)),$$

for $i = 1, \ldots, 5$. As already mentioned in [Holt and Rees 1994], the old chop program can only successfully split a representation if it finds a null vector of a word in the generators contained in an invariant subspace. If the representation is given by the matrices a and b, it follows from the construction of a and b that the first basis vector of a singular word in the generators will lie in the invariant subspace generated by the first ne basis vectors. This justifies the conjugation with a random invertible matrix. The reason for taking random elements in the group is given by the observation that if the chosen generators are not random enough there will be short words in these generators whose nullspace is contained in the proper invariant subspace.

The tables below compare the running time. The first column gives the running time for the old version, the second column the running time for the new version. As one can see, if we input the original matrices a and b to chop, the old version has no problem in splitting the representation.

The first example is of dimension 8 over $F_{3^{10}}$, so the resulting representation over F_3 is of dimension 160.

generators	old	new
$a,\ b$	0.1 s	0.1 s
$x^1(a,b),\ y^1(a,b)$	66.4 s	0.6 s
$x^2(a,b),\ y^2(a,b)$	999.3 s	1.9 s
$x^3(a,b),\ y^3(a,b)$	1216.5 s	1.6 s
$x^4(a,b),\ y^4(a,b)$	558.8 s	0.6 s
$x^5(a,b),\ y^5(a,b)$	1081.9 s	0.6 s

The next example is of dimension 8 over $F_{2^{12}}$, so the resulting representation over F_2 is of dimension 192.

generators	old	new
$a,\ b$	0.4 s	0.5 s
$x^1(a,b),\ y^1(a,b)$	7.0 s	2.3 s
$x^2(a,b),\ y^2(a,b)$	272.8 s	2.2 s
$x^3(a,b),\ y^3(a,b)$	117.5 s	6.1 s
$x^4(a,b),\ y^4(a,b)$	551.4 s	2.4 s
$x^5(a,b),\ y^5(a,b)$	813.5 s	2.3 s

Similar runs for other primes indicated the same tendency. The running time of the old version is longer, the fluctuation is greater, and in principle it would be no problem, as predicted by the theory, to produce examples, where the quotient between the old running time and the new running time gets arbitrarily large.

ACKNOWLEDGEMENTS

We are indebted to Jan Draisma, Richard Parker, Lajos Rónyai, Magdolna Szőke, and to the anonymous referee for their useful remarks and suggestions.

REFERENCES

[Bini and Pan 1994] D. Bini and V. Y. Pan, *Polynomial and matrix computations, v. 1: Fundamental algorithms*, Birkhäuser, Boston, 1994.

[Eberly and Giesbrecht 1996] W. Eberly and M. Giesbrecht, "Efficient decomposition of associative algebras", pp. 170–178 in *ISSAC'96: Proceedings of the International Symposium on Symbolic and Algebraic Computation* (Zurich, 1996), edited by Y. N. Lakshman, ACM Press, New York, 1996.

[Giesbrecht 1995] M. Giesbrecht, "Nearly optimal algorithms for canonical matrix forms", *SIAM J. Comput.* **24**:5 (1995), 948–969.

[Holt and Rees 1994] D. F. Holt and S. Rees, "Testing modules for irreducibility", *J. Austral. Math. Soc. Ser. A* **57**:1 (1994), 1–16.

[Lux 1997] K. Lux, *Algorithmic methods in modular representation theory*, Habilitationschrift, Rheinisch-Westfälische Technische Hochschule, Aachen, Germany, 1997.

[Malcev 1942] A. Malcev, "On the representation of an algebra as a direct sum of the radical and a semisimple subalgebra", *Doklady Acad. Sci. URSS (N.S.)* **36** (1942), 42–45.

[Parker 1984] R. A. Parker, "The computer calculation of modular characters (the Meat-Axe)", pp. 267–274 in *Computational group theory* (Durham, 1982), edited by M. D. Atkinson, Academic Press, London, 1984.

[Pierce 1982] R. S. Pierce, *Associative algebras*, Graduate Texts in Math. **88**, Springer, New York, 1982.

[Reiner 1975] I. Reiner, *Maximal orders*, London Mathematical Society Monographs **5**, Academic Press, London and New York, 1975.

[Rónyai 1993] L. Rónyai, "Computations in associative algebras", pp. 221–243 in *Groups and computation* (New Brunswick, 1991), edited by L. Finkelstein and W. M. Kantor, DIMACS Ser. Discrete Math. Theoret. Comp. Sci. **11**, Amer. Math. Soc., Providence, 1993.

[Schönert et al. 1996] M. Schönert et al., *GAP: Groups, algorithms, and programming*, 5th ed., Lehrstuhl D für Mathematik, RWTH Aachen, 1996. See http://www.math.rwth-aachen.de/~GAP/WWW/.

Gábor Ivanyos, Computer and Automation Institute, Hungarian Academy of Sciences, Lágymányosi u. 11., H-1111 Budapest, Hungary (Gabor.Ivanyos@sztaki.hu)

Klaus Lux, Department of Mathematics, University of Arizona, 617 N. Santa Rita, Tucson 85721, Arizona, United States (klux@math.arizona.edu)

Received July 30, 1999; accepted in revised form November 8, 1999