The Meat-axe and f-cyclic matrices

S. P. GLASBY

Abstract. Let $\mathrm{M}(d,F)$ denote the algebra of $d\times d$ matrices over a field F, and denote by $m_X(t)$ and $c_X(t)$ the minimal and the characteristic polynomials of $X\in\mathrm{M}(d,F)$. We call X an f-*cyclic* matrix if f is an irreducible factor of $m_X(t)$ which does not divide $c_X(t)/m_X(t)$. We present a version of the Meat-axe algorithm that uses f-cyclic matrices. One advantage of f-cyclic matrices is that they unify and generalize previous work of Parker, Holt and Rees, Ivanyos and Lux, Neumann and Praeger. The greater abundance of f-cyclic matrices may lead to an improved probability/complexity analysis of the Meat-axe. The difficulties that occur when the Schur index exceeds one are explored.

Dedicated to Charles Leedham-Green on the occasion of his 65th birthday

2000 Mathematics subject classification: 15A52, 20C40

1. Introduction

Let A denote a finitely generated F-subalgebra of the algebra $\mathrm{M}(d,F)$ of all $d\times d$ matrices over a field F. Computational representation theory is concerned with the design, analysis and implementation of algorithms for elucidating the geometric/algebraic structure of A. Two basic geometric questions are: (1) Does A act irreducibly on the vector space F^d? If not, can a proper nonzero A-invariant subspace be found? If A is known to act irreducibly on $V:=F^d$, then another basic question is: (2) What is the space $\mathrm{Hom}_A(V,W)$ of all A-homomorphisms from V to an A-module W?

It is common to use the same name for different but related concepts. Thus one may say that problems (1) and (2) are solved using *the* Meat-axe algorithm. There are now a number of Meat-axe algorithms. Versions [**13**], [**3**], [**10**, **11**, **12**], [**4**] are concerned with the case when F is a finite field, and extensions such as [**2**, **14**] consider certain characteristic zero fields particularly $F=\mathbb{Q}$, see also [**9**]. We shall make a small step towards unifying and generalizing existing algorithms with the goal of providing a better understanding of the complexity and probability analysis of the Meat-axe.

Each version of the Meat-axe algorithm proves irreducibility by selecting random matrices $X\in A$ until one in a suitable subset S is found. The subset S varies with the version of the algorithm. Denote by S_P, S_{HR}, S_{NP}, and S_{fc} the subsets relevant to [**13**, **14**], [**3**], [**10**, **11**, **12**], and the present paper respectively. Each subset comprises certain $X \in A$ for which the characteristic polynomial $c_X(t)$, and the minimal polynomial $m_X(t)$ of X satisfy certain properties. Let S_P be the set of $X \in A$ for which $c_X(t)$ has an unrepeated linear factor in $F[t]$, as in [**13**, **14**]. Let S_{HR} be the set of $X \in A$ for which $c_X(t)$ has an unrepeated irreducible factor (of arbitrary degree) in $F[t]$, as in [**3**]. Let S_{NP} be the set of cyclic matrices in A, i.e. those for which $c_X(t) = m_X(t)$ as in [**10**, **11**, **12**]. In the present paper, S_{fc} comprises the *f-cyclic* matrices in A (see the abstract or §2 for a definition). The larger the subset S, the more likely that the MEAT-AXE will find a suitable matrix $X \in S$. For the purpose of this discussion it is important that S_{fc} properly contains S_{HR} and S_{NP}; clearly $S_P \subseteq S_{HR}$. Thus f-cyclic matrices unify existing work on the MEAT-AXE. We shall prove a general version of Simon Norton's irreducibility theorem for f-cyclic matrices over an arbitrary field F. As $S_{HR} \cup S_{NP} \subseteq S_{fc}$, we hope that a more precise probabilistic analysis of the MEAT-AXE algorithm can be given in the important case when F is a finite field. Neumann and Praeger [**10**, **11**, **12**] have begun an extensive program to better understand the complexity of, and probability analysis for, the finite field MEAT-AXE using cyclic matrices. In [**10**] they show that the proportion of $X \in \mathrm{M}(d, \mathbb{F}_q)$ that are not cyclic is $q^{-3} + O(q^{-4})$. It appears that f-cyclic matrices are appreciably more abundant. For example, when $d = 3$ the proportion of X that are not f-cyclic (for any f) is $q^{-4} + O(q^{-5})$, see §6. For larger d, this proportion is likely even smaller.

Clever arguments in [**3**, **4**] show that the MEAT-AXE will find, with high probability, an invariant subspace in the case that A acts reducibly and F is finite. As $S_{HR} \subseteq S_{fc}$, it follows that an f-cyclic matrix version of the MEAT-AXE will succeed in the reducible finite field case with at least this probability. It is not hard to construct f-cyclic matrices that do not lie in S_{HR} or S_{NP}. In the examples below X is $(t - \lambda)$-cyclic:

- $X \notin S_{HR}$ if $c_X(t) = m_X(t) = (t - \lambda)^2$. If $\lambda \neq \mu \in F$, then
- $X \notin S_{NP}$ if $c_X(t) = (t - \lambda)(t - \mu)^2$ and $m_X(t) = (t - \lambda)(t - \mu)$, and
- $X \notin S_{HR} \cup S_{NP}$ if $c_X(t) = (t-\lambda)^2(t-\mu)^2$ and $m_X(t) = (t-\lambda)^2(t-\mu)$.

It is desirable to develop a theory of module-splitting in the most general (natural) setting. As the module-splitting problem subsumes the polynomial factorization problem, it is natural to consider the MEAT-AXE algorithm only for fields F where practical algorithms exist for factoring polynomials into irreducibles. In practice, this presently means that F is either a finite field, or a relatively small degree extension of $\mathbb{Q}$, see [**1**, **6**]. Although in many of our examples F is galois (even abelian) over its prime field, we shall not assume that this is so.

The paper is organized as follows. Notations and conventions are described in §2. The vector spaces in this paper are always over a field, although in §5 vector spaces over division rings are implicit. We take particular care with left and right actions of both scalars and functions. A generalization of Norton's irreducibility theorem is given in §3. The density of f-cyclic matrices in the image of blow-up monomorphisms is considered in §§4-5. The density is close to 1 if the commuting algebra is commutative, and is 0 otherwise. Some preliminary remarks regarding the density of f-cyclic matrices in $\mathrm{M}(d, F)$ are given in §6.

2. Conventions and Notation

The material in this section is known, or is part of the folklore. See for example [**16**, **8**, **3**, **11**]. As different conventions can be employed with regards to left or right actions of scalars, functions and matrices we shall explicitly state our conventions, and define our notation *en route*. Not all of the remarks below hold when F is a division ring. As we only need F to be a field, we can not justify the extra space required to generalize to division algebras.

Let V denote the *left* F-vector space of $1 \times d$ matrices over the field F. View V as a *right* module for the ring $\mathrm{M}(d, F)$ of $d \times d$ matrices over F. We identify $V^* = \mathrm{Hom}_F(V, F)$ with the *right* F-vector space of $d \times 1$ matrices over F, and view V^* as acting on V on the *right*. Scalar multiplication in V^* is defined by

$$v(f\lambda) = (vf)\lambda \qquad (v \in V, f \in V^*, \lambda \in F). \tag{1}$$

Note that scalar multiplication in V^* satisfies $f(\lambda\mu) = (f\lambda)\mu$. In addition, V^* becomes a *left* $\mathrm{M}(d, F)$-module where the $\mathrm{M}(d, F)$-action is via matrix multiplication.

Given $u \in V$ and $v^* \in V^*$ we identify the 1×1 matrix $uv^* = [\lambda]$ with the scalar λ. The bilinear form $V \times V^* \to F$ defined by $(u, v^*) \mapsto uv^*$ is nondegenerate. Thus for each basis $v_1, \ldots, v_d$ of V, there exists a *dual* basis $v_1^*, \ldots, v_d^*$ of V^* satisfying $v_i v_j^* = \delta_{ij}$. As usual, δ_{ij} equals 1 if $i = j$, and 0 otherwise. Abbreviate "is a subspace of" by $\leq$. If $U \leq V$ and $W \leq V^*$, then $U^\perp \leq V^*$ and $W^\perp \leq V$ are defined by

$$U^\perp = \{v^* \in V^* \mid Uv^* = 0\} \quad \text{and} \quad W^\perp = \{v \in V \mid vW = 0\}. \tag{2}$$

Then $\dim_F(U^\perp) = d - \dim_F(U)$ and $\dim_F(W^\perp) = d - \dim_F(W)$.

Let U and V be left F-spaces. Then $f \in \mathrm{Hom}_F(U, V)$ satisfies

$$(u_1 + u_2)f = u_1 f + u_2 f \qquad (\lambda u)f = \lambda(uf) \qquad\qquad (\lambda \in F, u \in U).$$

Make $\mathrm{Hom}_F(U,V)$ a *right* F-space by defining f_1+f_2 and $f\lambda$ by

(3) $$u(f_1+f_2)=uf_1+uf_2 \quad\text{and}\quad u(f\lambda)=\lambda(uf) \qquad (u\in U).$$

Then $\mathrm{End}_F(V):=\mathrm{Hom}_F(V,V)$ is an F-algebra. Given $f\in\mathrm{Hom}_F(U,V)$ and $v^*\in V^*$ define $f^*(v^*)\in U^*$ by

(4) $$u(f^*(v^*))=(uf)v^* \qquad (u\in U, f\in\mathrm{Hom}_F(U,V), v^*\in V^*).$$

If follows from Eq. (3) and (4) that

(5) $$f^*(v_1^*+v_2^*)=f^*(v_1^*)+f^*(v_2^*) \quad\text{and}\quad f^*(v^*\lambda)=f^*(v^*)\lambda$$

and hence $f^*\in\mathrm{Hom}_F(V^*,U^*)$. Similarly Eq. (3) and (4) imply

$$f_1^*+f_2^*=(f_1+f_2)^* \quad\text{and}\quad \lambda f^*=(f\lambda)^*.$$

Hence $\mathrm{Hom}_F(V^*,U^*)$ is a left F-space and $\mathrm{End}_F(V^*)$ is an F-algebra.

Given bases (u_i) for U and (v_j) for V, we can define a matrix for $f\in\mathrm{Hom}_F(U,V)$. Assume that the matrix for $f^*\in\mathrm{Hom}_F(V^*,U^*)$ is relative to the bases (v_j^*) for V^* and (u_i^*) for U^*. The following calculation shows that the matrices of f and f^* are the *same* (not transposed). Suppose that

$$u_if=\sum_k f_{ik}v_k \quad\text{and}\quad f^*v_j^*=\sum_k u_k^*f_{kj}^*$$

where $(f_{ik}),(f_{kj}^*)\in\mathrm{M}(d,F)$ are uniquely determined. Using the equations $u_iu_k^*=\delta_{ik}$, $v_kv_j^*=\delta_{kj}$ and Eq. (4) gives

$$f_{ij}^*=u_i(\sum_k u_k^*f_{kj}^*)=u_i(f^*v_j^*)=(u_if)v_j^*=(\sum_k f_{ik}v_k)v_j^*=f_{ij}.$$

If $f\in\mathrm{Hom}_F(U,V)$ and $g\in\mathrm{Hom}_F(V,W)$, then $fg\in\mathrm{Hom}_F(U,W)$ is composed from left to right, while $f^*g^*\in\mathrm{Hom}_F(W^*,U^*)$ is composed from right to left. Repeated use of Eq. (4) gives:

$$\begin{aligned} u((fg)^*(w^*)) &= (u(fg))w^*=((uf)g)w^*=(uf)(g^*(w^*))\\ &= u(f^*(g^*(w^*))) \qquad\qquad (u\in U, w^*\in W^*)\end{aligned}$$

and hence $(fg)^*=f^*g^*$ holds. Thus $\mathrm{End}_F(V)\to\mathrm{End}_F(V^*)\colon f\mapsto f^*$ is an F-algebra isomorphism. Moreover, the maps $f\mapsto(f_{ij})$ and $f^*\mapsto(f_{ij}^*)$ define F-algebra isomorphisms from $\mathrm{End}_F(V)$ and $\mathrm{End}_F(V^*)$ to $\mathrm{M}(d,F)$. In our context neither map is an anti-isomorphisms, c.f. [**16**].

We call $e_1,\dots,e_d$ the *standard* basis for V. This basis has the additional property that the dual basis $e_1^*,\dots,e_d^*$ for V^* is also standard. Given a basis $x_1,\dots,x_d$ for V, the transposed basis $x_1^T,\dots,x_d^T$ coincides with the dual basis $x_1^*,\dots,x_d^*$ for V^* if and only if $XX^T=I$

where X denotes the matrix whose ith row is x_i. The standard basis has $X = I$.

Henceforth, A denotes an F-subalgebra of $\mathrm{M}(d, F)$. Assume without loss of generality that $1 \in A$. Suppose that V is an A-module, and U is an A-submodule of V. Then $U^\perp$ is an A^*-submodule of V^*. (Since $U(A^*U^\perp) = (UA)U^\perp = UU^\perp = 0$ by Eq. (4), it follows that $A^*U^\perp = U^\perp$.) Similarly, if W is an A^*-submodule of V^*, then $W^\perp$ is an A-submodule of V.

Fix a matrix $X \in \mathrm{M}(d, F)$. It is standard, c.f. [**7**], to view V as a right $F[t]$-module where scalar multiplication is defined by

$$vf(t) = vf(X) \qquad (v \in V, f(t) \in F[t]).$$

Alternatively, V may be viewed as a right $F[X]$-module where

$$F[X] := \{f(X) \mid f(t) \in F[t]\}$$

is isomorphic to the quotient ring $F[t]/m_X(t)F[t]$. (If F were a non-commutative division ring, then $F[X]$ need not be closed under multiplication, as F need not commute with X.)

The *minimal polynomial* of an X-invariant subspace U of V is defined to be the minimal polynomial of the restriction $X|U$, i.e. $m_{X|U}(t)$. The minimal polynomial of the cyclic subspace $uF[X]$ is also called the *order polynomial* of the vector $u \in V$. It is important in the sequel that V is an internal direct sum $V = V_1 \dotplus \cdots \dotplus V_r$ where each $V_i = v_iF[X]$ is cyclic, and d_{i+1} divides d_i for $1 \le i < r$ where $d_i(t) = m_{X|V_i}(t)$ is the order polynomial of v_i. The characteristic polynomial and the minimal polynomials of X are $c_X(t) = d_1(t)d_2(t)\cdots d_r(t)$ and $m_X(t) = d_1(t)$ respectively.

Definition. A matrix $X \in \mathrm{M}(d, F)$ is called f*-cyclic* if $f(t) \in F[t]$ is a monic irreducible divisor of $m_X(t)$ that does not divide $c_X(t)/m_X(t)$.

Put differently, X is f-cyclic if and only if f divides $c_X(t)$ and $m_X(t)$ with the same (positive) multiplicity. It is clear, therefore, that a cyclic matrix X is f-cyclic for *all* irreducible divisors f of $m_X(t)$. Also, if f is an unrepeated irreducible factor of $c_X(t)$, then X is f-cyclic. In summary, $S_{NP} \cup S_{HR} \subseteq S_{fc}$.

Let $m_X(t) = \prod f^{\mu(f)}$ be the factorization of $m_X(t)$ as a product of powers of distinct monic irreducible polynomials $f \in F[t]$. We may write $V = \dotplus V(f)$ where the sum is over monic irreducible divisors f of $m_X(t)$, and where the f-primary submodule, $V(f)$, of V has minimal polynomial $m_{X|V(f)}(t) = f(t)^{\mu(f)}$. It is clear that X is f-cyclic if and only if $V(f)$ is a cyclic $F[X]$-submodule, and X is cyclic if and only if $V(f)$ is a cyclic $F[X]$-submodule for *each* irreducible divisor f of $m_X(t)$.

Finally, we define the kernel and image of X and X^*:

$$\begin{aligned} \ker_V X &= \{v \in V \mid vX = 0\}, & \mathrm{im}_V X &= \{vX \mid v \in V\} = VX, \\ \ker_{V^*} X^* &= \{v^* \in V^* \mid X^* v^* = 0\}, & \mathrm{im}_{V^*} X^* &= \{X^* v^* \mid v^* \in V^*\}. \end{aligned}$$

3. Norton's Irreducibility Theorem

The following lemma is used to prove a version of Norton's irreducibility theorem for f-cyclic matrices.

Lemma 1. *Let $X \in \mathrm{M}(d, F)$ be an f-cyclic matrix. Then*

(a) $\ker_V f(X) = \mathrm{im}_V g(X)$ *where $g(t) = m_X(t)/f(t)$, and*

(b) *the restriction Y of X to* $\ker_V f(X)$ *has $c_Y(t) = m_Y(t) = f(t)$.*

Proof. (a) Using the notation at the end of §2, $V = V_1 \dot{+} \cdots \dot{+} V_r$ where $V_i = v_i F[X]$ is cyclic, and the $d_i := m_{X|V_i}(t)$ satisfy $d_r | \cdots | d_2 | d_1$. Since X is f-cyclic, f divides both $m_X(t) = d_1$ and $c_X(t) = d_1 d_2 \cdots d_r$ with the same multiplicity. As f is irreducible, it is coprime to $d_2, \ldots, d_r$. Since $0 = m_X(X) = g(X)f(X)$, it follows that $\mathrm{im}_V g(X) \subseteq \ker_V f(X)$. Conversely, let

$$v = v_1 h_1(X) + v_2 h_2(X) + \cdots + v_r h_r(X) \in \ker_V f(X)$$

where $h_1, h_2, \ldots, h_r \in F[t]$. Then $vf(X) = 0$ implies $v_i h_i(X) f(X) = 0$ for each i. If $i > 1$, then d_i divides $h_i f$ and hence d_i divides h_i. A similar argument shows that d_1 divides $h_1 f$, and hence $h_1 = k_1 g$ for some $k_1 \in F[t]$. In summary, $v = v_1 k_1(X) g(X) \in \mathrm{im}_V g(X)$. Thus $\ker_V f(X) \subseteq \mathrm{im}_V g(X)$, and equality obtains.

(b) Let Y be the restriction of X to $\mathrm{im}_V g(X)$. By part (a), $d_i | g$ for $i > 1$, and hence $V_i g(X) = 0$ for $i > 1$. Thus $\mathrm{im}_V g(X) = V_1 g(X)$ is cyclic generated by $v_1 g(X)$. Since $v_1 g(X) \neq 0$ and $v_1 g(X) f(X) = 0$, we see that $c_Y(t) = m_Y(t) = f$ as desired. □

The following theorem is influenced by [**3**] and [**11**].

Theorem 2. *Let A be an F-subalgebra of $\mathrm{M}(d, F)$. Suppose $X \in A$ is f-cyclic, and*

(a) *there exists $v \in \ker_V f(X)$ such that $vA = V$, and*

(b) *there exists $v^* \in \ker_{V^*} f(X^*)$ such that $A^* v^* = V^*$.*

Then V is an irreducible right A-module, and V^* is an irreducible left A^*-module.

Proof. Let U be a proper A-submodule of V. We shall prove that $U = \{0\}$ is the zero subspace. By Lemma 1(b), $\ker_V f(X)$ is an irreducible $F[X]$-submodule. As $U \cap \ker_V f(X)$ is an $F[X]$-submodule of $\ker_V f(X)$, it equals $\{0\}$ or $\ker_V f(X)$. The latter does not happen as by assumption (a), $v \in \ker_V f(X)$ satisfies $V = vA \subseteq U$, contradicting the fact that U is proper. Therefore $U \cap \ker_V f(X) = \{0\}$, and so $Uf(X) = U$. By Eq. (4), $U(f(X)^* v^*) = (Uf(X))v^*$, and by assumption (b), $f(X^*)v^* = 0$. Therefore

$$0 = U0 = U(f(X^*)v^*) = U(f(X)^* v^*) = (Uf(X))v^* = Uv^*.$$

Hence $v^* \in U^\perp$. The condition $A^* v^* = V^*$ implies that $U^\perp = V^*$, and hence that $U = \{0\}$. This proves that V is an irreducible A-module. The fact that V^* is an irreducible left A^*-module follows by considering perpendicular subspaces of A^*-submodules of V^*. $\square$

Theorem 2 suggests the following procedure:

f-cyclic irreducibility procedure.

Input. A finitely generated F-subalgebra A of $\mathrm{M}(d, F)$.
Output. A boolean value for IsIrreducible.

1. Choose a random $X \in A$ until an f-cyclic matrix is found.
2. Find $0 \neq v \in \ker_V f(X)$. If $vA \neq V$, then IsIrreducible := False, and stop.
3. Find $0 \neq v^* \in \ker_{V^*} f(X^*)$. If $A^* v^* \neq V^*$, then IsIrreducible is set False, and stop.
4. IsIrreducible := True, and stop.

It is clear that this procedure terminates correctly when it does terminate: it correctly returns False in Steps 2 or 3, and correctly reports True in Step 4 by Theorem 2. Unfortunately, it may fail to find an f-cyclic $X \in A$ in Step 1. This can happen when A acts reducibly, for example when V is a direct sum of isomorphic irreducible A-submodules. In this case no $X \in A$ is f-cyclic, and Step 1 fails to terminate. One solution to this conundrum is to recast our procedure along the lines of [**3**].

f-cyclic Meat-axe procedure.

Input. A finitely generated F-subalgebra A of $\mathrm{M}(d, F)$.
Output. A boolean value for IsIrreducible, and a witness.

1. Choose a random $X \in A$.
2. For each irreducible factor f of $c_X(t)$ do

2a. Select a random $0 \neq v \in \ker_V f(X)$. If $vA \neq V$, then

IsIrreducible := False; Witness := vA; and stop.

2b. Select a random $0 \neq v^* \in \ker_{V^*} f(X^*)$. If $A^*v^* \neq V^*$, then IsIrreducible := False; Witness := A^*v^*; and stop.

2c. If X is f-cyclic, then IsIrreducible := True; set Witness to be (X, f, v, v^*); and stop.

3. Return to Step 1.

According to Knuth [**5**, pp. 4–6] and *algorithm* is a definite sequence of instructions which terminates after finitely many steps with provably correct output. This definition is arguably too restrictive. It is somewhat awkward to describe [**6**, p. 748, p. 751] a Las Vegas algorithm or a Monte Carlo algorithm as a "computational method" or "procedure" because these do not satisfy the strict definition of an algorithm. In this paper a *Monte Carlo algorithm* is a definite sequence of instructions involving random selections which terminates *with high probability* yielding output that is *with high probability* provably correct. More precisely, given real numbers $\varepsilon_1, \varepsilon_2$ satisfying $0 < \varepsilon_1, \varepsilon_2 < 1$ there is a number N depending on $\varepsilon_1, \varepsilon_2$ and the size of the input, such that if N random selections are made then the probabilities of non-termination, and of incorrect termination, are at most ε_1 and ε_2 respectively. A *Las Vegas algorithm* is defined similarly except that when it terminates, it terminates correctly. Knuth's definition [**5**, pp. 4–6] of an algorithm essentially has $\varepsilon_1 = \varepsilon_2 = 0$, while a Las Vegas algorithm has $\varepsilon_2 = 0$.

In the case that F is finite, and Berlekamp's algorithm [**6**, p. 441] or the Cantor-Zassenhaus Las Vegas algorithm [**6**, p. 447] is used to factor $c_X(t)$, then our f-cyclic Meat-axe procedure is a Las Vegas algorithm. If A is irreducible, then it follows from [**3**, **10**] (as $S_{HR} \cup S_{NP} \subseteq S_{fc}$) that an f-cyclic $X \in A$ will be found with high probability after N random selections, and hence IsIrreducible will be correctly set True. If A is reducible, then it follows from [**3**, **4**] that with high probability a proper nonzero subspace will be found in Step 2a or 2b after N random selections.

Another solution to the conundrum of non-termination of the f-cyclic irreducibility procedure is to recast it as a Las Vegas algorithm for proving irreducibility. If F is finite and A is irreducible, then it follows from [**11**, **12**] (as $S_{NP} \subseteq S_{fc}$) that the procedure will correctly, and likely, set IsIrreducible to be True. That is, incorrect termination is impossible, and the probability of non-termination can be made arbitrarily small by choosing N sufficiently large. The case when F is infinite, however, presents a challenge to both of the above procedures. The example at the end of §5 shows that A can be irreducible and yet no

$X \in A$ is f-cyclic. In the case that F is finite it was correct to view non-termination of the f-cyclic irreducibility procedure as evidence that A is reducible, however, it is more complicated when F is infinite.

The selection of a random $X \in A$ implies the existence of a probability measure on A. If F is finite, then it is natural to use the uniform measure on A where the probability of selecting any matrix is $|A|^{-1}$. If F is infinite, then so is A, and the choice of a probability measure on A is less obvious.

It is not the purpose of this paper to study the complexity of, and probability analysis for, these procedures. This would require very careful statements of our assumptions, and the analysis depends heavily on whether or not F is infinite. We shall however, make progress in §§4–5 towards understanding the conditional probability that an f-cyclic $X \in A$ is not found after N selections, given that A is irreducible.

4. Blowing up fields

By the Wedderburn-Artin structure theorem [**8**] for rings, $A/\mathrm{rad}(A)$ a direct sum of semisimple rings, and if A is simple then $A \cong \mathrm{M}(r, D)$ where D is a division algebra over F.

Let D denote a division ring, and let F be a subfield of finite index of the center of D. Choose a basis for D over F, and let ϕ and Φ denote the corresponding blow-up monomorphisms

$$\phi\colon D \to \mathrm{M}(|D:F|, F) \quad \text{and} \quad \Phi\colon \mathrm{M}(r, D) \to \mathrm{M}(r|D:F|, F).$$

In §§4-5 we consider the density of f-cyclic matrices in $\mathrm{im}\Phi$. Let E denote a maximal subfield of D. In this section $D = E$, and the proportion of $X \in \mathrm{M}(r, D)$ such that $\Phi(X)$ is f-cyclic is close to 1, see Theorem 4. In §5, $D \neq E$ holds, and the above proportion is 0.

The finite extension $E : F$ may be assumed to be separable by [**16**, Theorem 7.15]. (In the cases of most interest to us, namely when F is finite or of characteristic zero, each maximal subfield E of D is obviously separable over F.) To simplify our exposition, we shall make the stronger assumption that $E : F$ is *galois*. Our results generalize readily to the finite separable case, as outlined later.

Given $f \in E[t]$ and $\sigma \in G := \mathrm{Gal}(E/F)$, denote by $\sigma(f)$ the polynomial obtained by applying σ to the coefficients of f. Define maps $L, N\colon E[t] \to F[t]$ by

$$L(f) = \mathrm{lcm}\{\sigma(f) \mid \sigma \in G\} \quad \text{and} \quad N(f) = \prod_{\sigma \in G} \sigma(f).$$

The map L may not have a standard name, however, N is called the *norm map* from $E[t]$ to $F[t]$. If the domains or codomains of L, N are ambiguous, we write $L_{E/F}$ and $N_{E/F}$.

Lemma 3. *Let $E : F$ be a finite galois extension with group G. Let $f, f_1, f_2 \in E[t]$ be monic polynomials. Then*

(a) *$L(\sigma(f)) = L(f)$ and $N(\sigma(f)) = N(f)$ for $\sigma \in G$.*
(b) *$L(f), N(f) \in F[t]$ and $L(f)$ divides $N(f)$.*
(c) *If $f \in E[t]$ is irreducible, then $L(f) \in F[t]$ is irreducible and $N(f) = L(f)^{|E:F(f)|}$ where $F(f)$ is the field generated by F and the coefficients of f.*
(d) *Let $f_1, f_2 \in E[t]$ be irreducible. The following are equivalent:*
(1) *$\gcd(L(f_1), L(f_2)) \neq 1$,*
(2) *$L(f_1) = L(f_2)$, and*
(3) *$\sigma(f_1) = f_2$ for some $\sigma \in G$.*
(e) *$N(f_1)N(f_2) = N(f_1 f_2)$ and $\operatorname{lcm}\{L(f_1), L(f_2)\}$ divides $L(f_1 f_2)$.*
(f) *If $f \in E[t]$ is irreducible, then $L(f^n) = L(f)^n$ for $n \in \mathbb{Z}$, $n \geq 0$.*

Proof. Parts (a) and (b) are clear. Let $g \in F[t]$ be an irreducible divisor of $L(f)$ where $f|g$. Since $\sigma(f)|\sigma(g)$ and $\sigma(g) = g$, it follows that $L(f)|g$, and hence $L(f) = g$. Let $\{f_1, f_2, \dots f_r\}$ be the orbit of f under G. For each i, there are $|E : F(f)|$ choices for $\sigma \in G$ such that $\sigma(f) = f_i$, and hence $N(f) = L(f)^{|E:F(f)|}$. This proves (c). Part (1) implies part (2) by (c). Also (2) implies (3) by comparing factorizations in $E[t]$. Finally (3) implies f_2 divides $L(f_1)$ and $L(f_2)$, and this implies (1). This proves (d). The multiplicative property of N follows from $\sigma(f_1 f_2) = \sigma(f_1)\sigma(f_2)$. Since $f_1|L(f_1 f_2)$ it follows that $L(f_1)|L(f_1 f_2)$. Similarly, $L(f_2)|L(f_1 f_2)$ and hence $\operatorname{lcm}\{L(f_1), L(f_2)\}$ divides $L(f_1 f_2)$. This proves (e). Finally, part (f) follows as the orbits of f and f^n under G are $\{f_1, f_2, \dots f_r\}$ and $\{f_1^n, f_2^n, \dots f_r^n\}$. ☐

We need a stronger result than [**10**, Corollary 5.2] in order to deal with f-cyclic matrices.

Theorem 4. *Let $E : F$ be a finite galois extension with group G, and let $\Phi \colon \mathrm{M}(r, E) \to \mathrm{M}(r|E : F|, F)$ be a blow-up monomorphism. Then*

(a) *$m_{\Phi(X)}(t) = L(m_X(t))$ and $c_{\Phi(X)}(t) = N(c_X(t))$ for $X \in \mathrm{M}(r, E)$.*
(b) *$\Phi(X)$ is g-cyclic for some irreducible divisor $g \in F[t]$ of $m_{\Phi(X)}(t)$ if and only if X is f-cyclic where $f := \gcd(g, m_X(t))$ is irreducible in $E[t]$.*

Proof. (a) As $m_{\Phi(X)}(X) = 0$, it follows that $m_X(t)$ and thus $L(m_X(t))$ divides $m_{\Phi(X)}(t)$. Conversely, $L(m_X(t)) \in F[t]$ and $L(m_X(\Phi(X))) = 0$.

Hence $m_{\Phi(X)}(t)$ divides $L(m_X(t))$, and so equality holds, c.f. [**10**, Lemma 5.1]. See [**16**, Theorem 9.10] for a proof that $c_{\Phi(X)}(t) = N(c_X(t))$.

(b) Let $c_X(t) = \prod f^{c(f)}$ and $m_X(t) = \prod f^{m(f)}$ be the factorizations of $c_X(t)$ and $m_X(t)$ as a product of powers of distinct monic irreducible polynomials in $E[t]$. Let $c_{\Phi(X)}(t) = \prod g^{C(g)}$ and $m_{\Phi(X)}(t) = \prod g^{M(g)}$ be corresponding factorizations in $F[t]$. By Lemma 3(c,d,e)

$$C(g) = \sum\{c(f) \mid L(f) = g\} \quad \text{and} \quad M(g) = \max\{m(f) \mid L(f) = g\}.$$

Assume that $\Phi(X)$ is g-cyclic, equivalently that $C(g) = M(g)$. Using the above displayed equation and $c(f) \geq m(f)$, there is only one irreducible divisor f of $m_X(t)$ such that $L(f) = g$, and for this divisor $c(f) = m(f)$. Therefore, X is f-cyclic and $f|g$. This proves the 'if' part of (b), the 'only if' part is proved by reversing the above arguments. □

Theorem 4(b) may be rephrased: X is f-cyclic and

$$\gcd(L(f), m_X(t)) = f \tag{6}$$

holds if and only if $\Phi(X)$ is $L(f)$-cyclic. Eq. (6) holds if and only if $\gcd(\sigma(f), m_X(t)) = 1$ for $1 \neq \sigma \in G$, c.f. [**10**, Corollary 5.2].

We return to our weaker assumption that $E : F$ is finite and separable. There is a finite extension K of E which is galois over F, see [**17**]. If $m = |E : F|$, then there are m monomorphisms, say $\sigma_1, \ldots, \sigma_m$, from E into K. For each i, there are precisely $|K : E|$ automorphisms $\sigma \in \mathrm{Gal}(K/F)$ such that the restriction $\sigma|E$ equals σ_i. Given $f \in E[t]$ define the maps $L, N \colon E[t] \to F[t]$ as follows:

$$L_{E/F}(f) = \mathrm{lcm}\{\sigma_1(f), \ldots, \sigma_m(f)\} \quad \text{and} \quad N_{E/F}(f) = \sigma_1(f) \cdots \sigma_m(f).$$

The connection between $L_{E/F}, N_{E/F}$ and $L_{K/F}, N_{K/F}$ is

$$L_{K/F}(f) = L_{E/F}(f) \quad \text{and} \quad N_{K/F}(f) = N_{E/F}(f)^{|K:E|}.$$

With the preceding remarks, Lemma 3 and Theorem 4 can be generalized by replacing "galois" by "separable". Minor modifications are required to the statements and proofs. For example, $\sigma \in G$ becomes $\sigma \in \mathrm{Gal}(K : F)$ where K is the galois closure of $E : F$. The details are left to the reader. Compare with [**16**, Exercise 9.4].

The density of f-cyclic matrices $X \in \mathrm{M}(r, E)$ such that $\Phi(X)$ is $L(f)$-cyclic in the case that $|F| = q$ is finite, is at least the density given in [**10**, **12**] because cyclic matrices are f-cyclic for each f. The density in the cyclic case is at least $1 - q^{-1} + O(q^{-2})$. A stronger bound exists when $|E : F| > 2$ see [**10**, Theorem 5.5]. While the density of f-cyclic matrices in $\mathrm{im}(\Phi)$ exceeds the density of cyclic matrices, it is unclear whether or not higher powers of q^{-1} are involved for most r, E.

5. Blowing up division rings

Now consider the case when D is a *noncommutative* division algebra with center F of finite index. Let E be a maximal subfield of D containing F. Then $|D : E| = |E : F| = m$ is the (Schur) index of D, and $m > 1$ as $D \neq E$. Let ϕ be a blow-up monomorphism $\phi\colon D \to \mathrm{M}(m^2, F)$, and define $\Phi\colon \mathrm{M}(r, D) \to \mathrm{M}(rm^2, F)$ by $\Phi((\lambda_{ij})) = (\phi(\lambda_{ij}))$. Our main result is:

Theorem 5. *No element of $\Phi(\mathrm{M}(r, D))$ is f-cyclic. Indeed, $m_Y(t)^m$ divides $c_Y(t)$ for each $Y = \Phi(X) \in \mathrm{im}(\Phi)$.*

Proof. Set $A := \Phi(\mathrm{M}(r, D))$, $B := \mathrm{M}(rm^2, F)$, and $C := C_B(A)$ where $C_B(A)$ is the centralizer in B of A. We shall prove first that $C \cong D^{\mathrm{op}}$.

Consider the product AC of subrings of the ring B. Then AC is a subring of B whose elements are finite sums $\sum a_i c_i$, where $a_i \in A$, $c_i \in C$. Recall that an F-algebra A is called *central* if its center equals $\{\lambda 1 \mid \lambda \in F\}$. Accordingly, A and B are central simple F-algebras. By [**8**, Theorem 4.7], $B \cong A \otimes_F C$. The map $A \otimes_F C \to B$ defined by $\sum a_i \otimes c_i \to \sum a_i c_i$ is a homomorphism whose image is AC. Since $A \otimes_F C \cong B$ is simple, the homomorphism is injective. Hence $A \otimes_F C \cong AC$ and $B = AC$.

By [**8**, Theorem 4.6], $D \otimes D^{\mathrm{op}} \cong \mathrm{M}(m^2, F)$, and hence the centralizer of $\phi(D)$ in $\mathrm{M}(m^2, F)$ is isomorphic to D^{op}. Therefore the centralizer C of A in B contains a subring isomorphic to D^{op}. However,

$$|C : F| = \frac{|B : F|}{|A : F|} = \frac{(rm^2)^2}{r^2 m^2} = m^2 = |D^{\mathrm{op}} : F|. \tag{7}$$

It follows from $D^{\mathrm{op}} \subseteq C$ and Eq. (7) that $C \cong D^{\mathrm{op}}$ as desired.

We view $\mathrm{M}(r, D)$ as a subring of $\mathrm{M}(r|D : E|, E)$ by blowing up over E (rather than over F). The center of $\mathrm{M}(r|D : E|, E)$ comprises scalar matrices over E, and thus we may view E as a subring of C. Let $\lambda_1, \dots, \lambda_m$ be a basis for C as a left E-space. Then

$$C = E\lambda_1 \dotplus \cdots \dotplus E\lambda_m \quad \text{and} \quad B = AC = AE\lambda_1 \dotplus \cdots \dotplus AE\lambda_m.$$

Thus $V = VB = V_1 \dotplus \cdots \dotplus V_m$ where $V_i = VAE\lambda_i$ is a right A-module. (Note that $E\lambda_i \subseteq C \cong D^{\mathrm{op}}$ and A commutes with $E\lambda_i$.) The map $V_i \to V_j$ defined by $v \mapsto v\lambda_i^{-1}\lambda_j$ is an A-module isomorphism. We view each V_i as an F-space. For $X \in A$ we write $X = X_1 \dotplus \cdots \dotplus X_m$ where X_i is the restriction of X to V_i. Accordingly, the minimal polynomial $m_{X_i}(t)$ lies in $F[t]$ (and not $E[t]$) and $m_X(t) = m_{X_1}(t)$ because $m_{X_1}(t) = m_{X_i}(t)$ for $i = 1, \dots, m$. Since $c_X(t) = c_{X_1}(t)^m$ and $m_{X_1}(t)$ divides $c_{X_1}(t)$, it follows that $c_X(t) = h(t)^m m_X(t)^m$ where

$h(t) = c_{X_1}(t)/m_{X_1}(t) \in F[t]$. As $m > 1$, this proves that no element of A is f-cyclic. □

Using the terminology of [**16**, §9a], $c_{X_1}(t)$ is the *reduced characteristic polynomial* of $X \in A$. The proof of Theorem 5 gives alternate proofs of Theorems 9.3 and 9.5 in [**16**].

Theorem 5 shows that the procedures in §3 both fail to find an f-cyclic matrix when V is irreducible and $D := \mathrm{End}_A(V)$ is noncommutative. This is surprising for the following reason. A heuristic argument in [**10**, §1] shows that almost all matrices in $\mathrm{M}(d, F)$ with $F \subseteq \mathbb{C}$ are separable, and hence f-cyclic. If $\mathrm{im}(\Phi)$ were randomly spread throughout $\mathrm{M}(rm^2, F)$, then we would expect that almost all matrices in $\mathrm{im}(\Phi)$ are f-cyclic, contrary to Theorem 5. Thus special care must be taken as the image and codomain of Φ have vastly different densities of f-cyclic matrices. Indeed, as $S_P \subseteq S_{fc}$, Parker's MEAT-AXE algorithm [**14**] also fails to terminate when V is irreducible and $D := \mathrm{End}_A(V)$ is noncommutative.

We construct an example with $m > 1$. Let F be a subfield of the real numbers $\mathbb{R}$, and let D be the quaternion algebra over F with elements

$$\lambda = \lambda_0 + \lambda_1 i + \lambda_2 j + \lambda_3 ij \qquad (\lambda_0, \lambda_1, \lambda_2, \lambda_3 \in F)$$

where $i^2 = j^2 = (ij)^2 = -1$. Consider the blow-up monomorphism

$$\phi\colon D \to \mathrm{M}(4, F) \quad \text{defined by} \quad \phi(\lambda) = \begin{pmatrix} \lambda_0 & \lambda_1 & \lambda_2 & \lambda_3 \\ -\lambda_1 & \lambda_0 & -\lambda_3 & \lambda_2 \\ -\lambda_2 & \lambda_3 & \lambda_0 & -\lambda_1 \\ -\lambda_3 & -\lambda_2 & \lambda_1 & \lambda_0 \end{pmatrix}$$

relative to the F-basis $1, i, j, ij$. Set $\lambda^* = \lambda_0 - \lambda_1 i - \lambda_2 j - \lambda_3 ij$, and define $f_\lambda(t)$ by

$$f_\lambda(t) = (t - \lambda)(t - \lambda^*) = t^2 - 2\lambda_0 t + (\lambda_0^2 + \lambda_1^2 + \lambda_2^2 + \lambda_3^2) \in F[t].$$

Then $c_{\phi(\lambda)}(t) = f_\lambda(t)^2$. As $f_\lambda(t)$ has discriminant $-4(\lambda_1^2 + \lambda_2^2 + \lambda_3^2) \leq 0$, we see $m_{\phi(\lambda)}(t) = f_\lambda(t)$ is irreducible if $\lambda \neq \lambda^*$, and $m_{\phi(\lambda)}(t) = t - \lambda_0$ if $\lambda = \lambda^*$. In either case, $m_{\phi(\lambda)}(t)^2$ divides $c_{\phi(\lambda)}(t)$, and $\phi(D)$ contains no f-cyclic matrices.

6. PROPORTIONS OF f-CYCLIC MATRICES

In this section we compare the proportion of $X \in \mathrm{M}(3, \mathbb{F}_q)$ that are f-cyclic with the proportion that are cyclic.

There are three *types* of matrix $X \in \mathrm{M}(3, \mathbb{F}_q)$ that are not cyclic. These are listed below according the the values of $c_X(t)$ and $m_X(t)$. The first two types are not f-cyclic, while the third is $(t - \mu)$-cyclic. Set $G := \mathrm{GL}(3, \mathbb{F}_q)$. Then

$c_X(t)$	$(t-\lambda)^3$	$(t-\lambda)^3$	$(t-\lambda)^2(t-\mu) \quad \lambda \neq \mu$
$m_X(t)$	$(t-\lambda)^2$	$t-\lambda$	$(t-\lambda)(t-\mu)$
$\lvert C_G(X)\rvert$	$q^2(q-1)^2$	$\lvert G\rvert$	$(q^2-1)(q^2-q)(q-1)$
$\lvert G : C_G(X)\rvert$	$(q^3-1)(q+1)$	1	$q^2(q^2+q+1)$
$\#c_X(t)$	q	q	$q(q-1)$

The dot product of the last two rows is the number of non-cyclic matrices in $\mathrm{M}(3, \mathbb{F}_q)$, namely $q^6+q^5+q^4-q^3-q^2$. By comparison the number of non-f-cyclic matrices is $q^5+q^4-q^2$. Hence the density of cyclic matrices and f-cyclic matrices is respectively

$$1-q^{-3}-q^{-4}-q^{-5}+q^{-6}+q^{-7} \quad \text{and} \quad 1-q^{-4}-q^{-5}+q^{-7}.$$

Although it appears that the density of f-cyclic matrices in $\mathrm{M}(d, \mathbb{F}_q)$ may increase [**15**] as a function of d, it is unclear at this stage whether or not an f-cyclic Meat-axe algorithm will be more efficient than a cyclic Meat-axe algorithm.

Acknowledgement

I am grateful to the referee: his/her comments improved this paper.

References

[**1**] Henri Cohen, *A Course in Computational Algebraic Number Theory*, Graduate Texts in Mathematics **138**, Springer, 1993.

[**2**] D.F. Holt, The Meataxe as a tool in computational group theory, in *The atlas of finite simple groups: ten years on (Birmingham, 1995)*, London Math. Soc. Lecture Note Series, vol. 249, Cambridge Univ. Press, Cambridge, pp. 74–81.

[**3**] D.F. Holt and S. Rees, *Testing modules for irreducibility*, J. Austral. Math. Soc. Ser. A **57** (1994), no. 1, 1–16.

[**4**] G. Ivanyos and K. Lux, *Treating the exceptional case of the MeatAxe*, Experiment. Math. **9** (2000), no. 3, 373–381.

[**5**] Donald E. Knuth, *The Art of Computer Programming, vol. 1: Fundamental Algorithms*, 3rd Ed., Addison-Wesley, Reading, MA, 1997.

[**6**] Donald E. Knuth, *The Art of Computer Programming, vol. 2: Seminumerical Algorithms*, 3rd Ed., Addison-Wesley, Reading, MA, 1997.

[**7**] N. Jacobson, *Basic Algebra I*, W.H. Freeman and Company, 1974.

[**8**] N. Jacobson, *Basic Algebra II*, W.H. Freeman and Company, 1980.

[**9**] G. Michler and O. Solberg, *Testing modules of even order for simplicity*, J. Algebra **2002** (1998), 229–242.

[**10**] P. M. Neumann and C. E. Praeger, *Cyclic matrices over finite fields*, J. London Math. Soc. **52** (1995), no. 2, 263–284.

[**11**] P.M. Neumann and C.E. Praeger, *Cyclic matrices and the* Meataxe, in Groups and Computation III, Ohio State Univ. Math. Res. Inst. Publ. **8**, (2001), 291–299.

[**12**] P.M. Neumann and C.E. Praeger, *Exploiting cyclic matrices in computer algebra: sharpening the Meataxe*, Preprint June 2005.

[**13**] R. A. Parker, *The computer calculation of modular characters (the meat axe)*, in computational group theory, M. D. Atkinson (ed.), Proc. London Math. Soc. Symposium on Computational Group Theory in Durham, Academic Press, London, 1984, 267–274.

[**14**] R. A. Parker, *An integral meat-axe*. The atlas of finite groups: ten years on (Birmingham, 1995), 215–228, London Math. Soc. Lecture Note Ser. **249**, Cambridge Univ. Press, Cambridge, 1998.

[**15**] Cheryl E. Praeger, Personal Communication, May 2005.

[**16**] I. Reiner, *Maximal orders*, London Math. Soc. Monographs **5**, Academic Press, 1975.

[**17**] Steven Roman, *Field Theory*, Graduate Texts in Mathematics **158**, Springer-Verlag, 1995

S.P. Glasby
Department of Mathematics
Central Washington University
WA 98926-7424, USA
http://www.cwu.edu/∼glasbys/