Efficient Decomposition of Separable Algebras*

W. Eberly†
Department of Computer Science
University of Calgary
Calgary, Alberta, Canada T2N 1N4
Email: eberly@cpsc.ucalgary.ca
http://www.cpsc.ucalgary.ca/~eberly

M. Giesbrecht†
Department of Computer Science
University of Western Ontario
London, Ontario, Canada N6A 5B7
Email: mwg@csd.uwo.ca
http://www.csd.uwo.ca/~mwg

September 7, 1999

Abstract

We present new, efficient algorithms for computations on separable matrix algebras over infinite fields. We provide a probabilistic of the Monte Carlo type to find a generator for the centre of a given algebra $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ over an infinite field F using a number of operations that is within a logarithmic factor of the cost of solving $m \times m$ systems of linear equations. A Las Vegas algorithm is also provided under the assumption that a basis and set of generators for the given algebra are available. These new techniques yield a partial factorization of the minimal polynomial of the generator that is computed, which may reduce the cost of computing simple components of the algebra in some cases.

1 Introduction

A *finite-dimensional associative algebra* is a finite-dimensional vector space over a field F equipped with a multiplication operation under which the space forms an associative (though not necessarily commutative) ring with identity, and a *matrix algebra* is a subalgebra of the matrix ring $\mathsf{F}^{m\times m}$. Algebras over finite fields have been studied in an earlier paper (Eberly & Giesbrecht 1999). In this paper we propose efficient new algorithms for separable algebras over infinite fields.

Recall that the *(Jacobson) radical* $\mathrm{Rad}(\mathfrak{A})$ of an algebra $\mathfrak{A}$ over a field F is the intersection of all maximal left ideals in $\mathfrak{A}$, and that $\mathfrak{A}$ is *semi-simple* if $\mathrm{Rad}(\mathfrak{A}) = (0)$. Such an algebra is *separable* if the algebra $\mathfrak{A}^{\mathsf{E}} = \mathfrak{A} \otimes_{\mathsf{F}} \mathsf{E}$ obtained from $\mathfrak{A}$ by "extension of scalars" is semi-simple over E, for every field extension E of F. Curtis & Reiner (1962) and Pierce (1982) each discuss the properties of extensions of scalars and separable algebras that will be used in this paper.

*Extended abstract of some of this work appears in "Proceedings, International Symposium on Symbolic and Algebraic Computation," Zurich, 1996. ISSAC'96, pp. 170–178.

†Research was supported in part by the Natural Sciences and Engineering Research Council of Canada

Submitted to the Journal of Symbolic Computation.

As they note, any semi-simple algebra over a field of characteristic 0 and, more generally, over any perfect field is separable, so that our algorithms apply to all such algebras.

The first provably efficient algorithms for computing the structure of a matrix algebra are due to Friedl & Rónyai (1985), who gave polynomial time algorithms to find the Jacobson radical and to decompose a semi-simple algebra over a finite field or number field as a direct sum of simple algebras. Subsequent work by Rónyai (1987, 1990, 1992) and Ivanyos & Rónyai (1993) examined additional questions over number fields, and in particular showed that deciding whether an algebra over a number field possesses nontrivial idempotents has the same complexity as factoring integers. That is, it is (currently) intractable. The problem of finding such idempotents may be considerably more difficult: The algorithms of Rónyai (1992) and Ivanyos & Rónyai (1993) answer the decision problem without generating such idempotents and, to our knowledge, no bounds on the size of these idempotents are presently known.

Other work concerning these computations over large fields includes the algorithms of Cohen *et al.* (1997) and Ivanyos (1999) for computation of the radical of an associative algebra, and the randomized algorithm of Eberly (1991) for computation of the simple components of semi-simple algebras over large perfect fields.

More practical work has concerned computations over finite fields, including the heuristic of Parker (1984) to test irreducibility of an $\mathfrak{A}$-module over a small finite field and to split reducible modules, and the more recent extension of the technique (now effective over arbitrary finite fields) of Holt & Rees (1994), as well as the work of Schneider (1990) and Eberly & Giesbrecht (1999) to compute primitive idempotents in associative algebras. All these algorithms take advantage of the fact that primitive idempotents are easy to find in associative algebras over finite fields. As noted above, Rónyai (1987) has established that this is not the case at all for associative algebras over number fields so that other techniques must be used in this case.

We propose modifications of the method originally given by Friedl & Rónyai (1985) and adapted by Eberly (1991) to find the simple components of a semi-simple algebra by decomposing its centre. As we note, the technique is applicable to separable algebras over arbitrary fields. We provide more efficient Monte Carlo and Las Vegas algorithms for the first step in this process, namely, computation of a generator γ for the centre of a given separable matrix algebra $\mathfrak{A}$ over an arbitrary large field. The method also yields a partial factorization of the minimal polynomial of γ. A complete factorization of this minimal polynomial is required to compute the simple components of $\mathfrak{A}$, and the cost of this factorization tends to dominate the cost of the entire process. Thus, our modifications will not reduce the asymptotic worst-case complexity. However, the modifications may replace the need for a factorization of a single polynomial of large degree with factorizations of several polynomials of lower degree, and may reduce the cost of the computation in practice.

We will generally tie the complexity of our results to that of matrix multiplication. We define $\mathcal{MM}(m)$ such that $O(\mathcal{MM}(m))$ operations in a field F are sufficient to multiply two matrices in $\mathsf{F}^{m\times m}$ and to solve nonsingular systems of m linear equations in m unknowns over F. Using the standard algorithm requires $\mathcal{MM}(m) = m^3$ while the currently best known algorithm of Coppersmith & Winograd (1982) allows $\mathcal{MM}(m) = O(m^{2.376})$. As well, we define $\mathcal{M}(m)$ such that $O(\mathcal{M}(m))$ operations in F suffice to multiply two polynomials in $\mathsf{F}[x]$ of degree m. Using the standard algorithm allows $\mathcal{M}(m) = O(m^2)$, while the algorithm of Schönhage & Strassen (1971) and Schönhage (1977) allows $\mathcal{M}(m) = m \log m \log\log m$. For notational convenience we make the reasonable assumption that $O(m\mathcal{M}(m)) \subseteq O(\mathcal{M}\mathcal{M}(m))$.

To prove correctness of our probabilistic algorithms, we require some technical conditions on the presumed ability to select a random element α from the algebra $\mathfrak{A}$. One rigorous way of doing this will be to select a sufficiently large finite subset S of the field F, as well as a finite set of elements of $\mathfrak{A}$ whose F-linear span includes elements with the properties we need, and then to select elements uniformly from the S-linear span of these elements of $\mathfrak{A}$. We prove in the sequel that if F is infinite then the algebra always includes the elements we require, so that it will be sufficient to choose elements from the S-linear span of a basis for $\mathfrak{A}$. This requires $O(nm^2)$ operations in F if $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$, $\mathfrak{A}$ has dimension n over F, and a basis for $\mathfrak{A}$ is available.

Additional notation and results concerning the structure of separable algebras and their modules, and necessary computations of matrix normal forms, are included in Section 2. Section 3 introduces "self-centralizing elements" of algebras and the properties that we will need to decompose these algebras. Useful pairs of these elements, which we call "centering pairs," are introduced in Section 4, and are used in new algorithms to compute the centre of $\mathfrak{A}$. Finally, the Wedderburn decomposition of separable algebras is considered in Section 5.

2 Preliminaries

2.1 The Structure of a Semi-Simple Matrix Algebra

Suppose henceforth that $\mathfrak{A}$ is a separable algebra of dimension n over a field F, and that there exists a faithful $\mathfrak{A}$-module of dimension m. By the Wedderburn Structure Theorem (Wedderburn, 1907)

$$\mathfrak{A} = \mathfrak{A}_1 \oplus \mathfrak{A}_2 \oplus \cdots \oplus \mathfrak{A}_k$$

for simple algebras $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k \subseteq \mathfrak{A}$, and each simple component $\mathfrak{A}_i$ is isomorphic to a full matrix ring over a division ring D_i over F, so that

$$\mathfrak{A}_i \cong \mathsf{D}_i^{t_i \times t_i}$$

for some positive integer t_i, for $1 \le i \le k$. Furthermore as shown, for example, by Pierce (1982), the dimension of each simple algebra (such as $\mathfrak{A}_i$, or D_i) over its centre is a perfect square. Let E_i be the centre of $\mathfrak{A}_i$ (isomorphic to the centre of D_i as well and, consequently, a field extension of F); let $e_i = [\mathsf{E}_i : \mathsf{F}]$, let d_i^2 be the dimension of D_i over E_i, and let $n_i = d_i t_i$, so that $\mathfrak{A}_i$ has dimension $e_i d_i^2 t_i^2 = e_i n_i^2$ over F for all i and

$$n = e_1 n_1^2 + e_2 n_2^2 + \cdots + e_k n_k^2. \tag{2.1}$$

Suppose in addition that $\mathfrak{A}$ is a matrix algebra, so that $\mathfrak{A}$ is a subalgebra of $\mathsf{F}^{m\times m}$ for some positive integer m. Now the vector space $\mathsf{F}^{m\times 1}$ is an $\mathfrak{A}$-module in a natural way: For any element α of $\mathfrak{A}$ and vector $v \in \mathsf{F}^{m\times 1}$, the result αv of applying α to v is simply the matrix-vector product obtained by multiplying the matrix $\alpha \in \mathsf{F}^{m\times m}$ to the vector v.

Since $\mathfrak{A}$ is separable, and therefore semi-simple, $\mathsf{F}^{m\times 1}$ is a semi-simple $\mathfrak{A}$-module. That is, $\mathsf{F}^{m\times 1}$ is the direct sum of a set of simple $\mathfrak{A}$-modules, each of which is a faithful $\mathfrak{A}_i$-module for exactly one simple component $\mathfrak{A}_i$ of $\mathfrak{A}$ and which annihilates all the other simple components $\mathfrak{A}_j$. Suppose a decomposition of $\mathsf{F}^{m\times 1}$ as a direct sum of simple modules includes exactly s_i simple modules $M_1^{(i)}, M_2^{(i)}, \ldots, M_{s_i}^{(i)}$ such that $\mathfrak{A}_i M_j^{(i)} = \mathfrak{A}_i$ for $1 \leq j \leq s_i$ and $1 \leq i \leq k$, so that $s_i \geq 1$ for all i (since $\mathsf{F}^{m\times 1}$ is clearly a faithful $\mathfrak{A}$-module), and

$$\mathsf{F}^{m\times 1} = M_1^{(1)} \oplus M_2^{(1)} \oplus \cdots \oplus M_{s_1}^{(1)} \oplus \cdots \oplus M_1^{(k)} \oplus M_2^{(k)} \oplus \cdots \oplus M_{s_k}^{(k)}. \tag{2.2}$$

This decomposition is not unique. However, the values $s_1, s_2, \ldots, s_k$ certainly are. Furthermore it is well-known (see, for example, Curtis & Reiner 1962) that all simple modules that are faithful $\mathfrak{A}_i$-modules are isomorphic as $\mathfrak{A}$-modules, and as vector spaces over F,

$$M_j^{(i)} \cong \mathsf{D}_i^{t_i \times 1}$$

for $1 \leq j \leq s_i$. Consequently $M_j^{(i)}$ has dimension $e_i d_i^2 t_i = e_i n_i d_i$ over F for $1 \leq j \leq s_i$ and $1 \leq i \leq k$. Let $m_i = d_i s_i$ for $1 \leq i \leq k$; then an inspection of equation (2.2) and comparison of dimensions of modules confirms that

$$m = e_1 n_1 d_1 s_1 + e_2 n_2 d_2 s_2 + \cdots + e_k n_k d_k s_k = e_1 n_1 m_1 + e_2 n_2 m_2 + \cdots + e_k n_k m_k. \tag{2.3}$$

2.2 Distinguishing Elements by Matrix-Vector Products

Since $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ it is clear that one can check whether a given element α of $\mathfrak{A}$ is zero by inspecting the m^2 entries of the matrix α. It will be useful in the sequel to check this condition by computing and inspecting matrix-vector products instead. Therefore, let

$$N = N_{\mathfrak{A}} = \max_{1\leq i\leq k} \left\lceil \frac{n_i}{m_i} \right\rceil = \max_{1\leq i\leq k} \left\lceil \frac{t_i}{s_i} \right\rceil. \tag{2.4}$$

Definition 2.1. A set of vectors $v_1, v_2, \ldots, v_N \in \mathsf{F}^{m\times 1}$ is a *distinguishing set* for $\mathfrak{A}$ if there exists at least one vector v_i in this set such that $\alpha v_i \neq 0$, for every nonzero element α of $\mathfrak{A}$.

Clearly, if a distinguishing set $v_1, v_2, \ldots, v_N$ of vectors is available, then we can check whether $\alpha = 0$ for a given element α of $\mathfrak{A}$ by computing and inspecting the N matrix-vector products $\alpha v_1, \alpha v_2, \ldots, \alpha v_N$. We can also check whether two elements α and β are equal in $\mathfrak{A}$ by using these vectors to decide whether the difference $\alpha - \beta$ is nonzero.

Theorem 2.2. *If $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ is a semi-simple algebra and N is as defined in equation* (2.4), *above, then a distinguishing set of vectors $v_1, v_2, \ldots, v_N \in \mathsf{F}^{m\times 1}$ for $\mathfrak{A}$ exists.*

Proof. Suppose first that $\mathfrak{A}$ is simple and that $s_1 = 1$, so that $\mathsf{F}^{m\times 1}$ is a simple $\mathfrak{A}$-module. In this case, $k = 1$, $n = e_1 n_1^2$ for $n_1 = d_1 t_1$, $m = e_1 n_1 m_1$ for $m_1 = d_1$, and $N = n_1/m_1 = t_1$. Furthermore the centralizer of $\mathfrak{A}$ in $\mathsf{F}^{m\times m}$ (that is, the set of matrices commuting with all the elements of $\mathfrak{A}$) is a division algebra D with dimension $e_1 d_1^2$ over F, and $\mathsf{F}^{m\times 1}$ may be regarded as a module with dimension t_1 over this division algebra.

Now it suffices to choose $v_1, v_2, \ldots, v_N$ to be any basis for $\mathsf{F}^{m\times 1}$ over the centralizer D to ensure that $v_1, v_2, \ldots, v_N$ is an isolating set for $\mathfrak{A}$. For if $\gamma_1, \gamma_2, \ldots, \gamma_{e_1 d_1^2}$ is a basis for D over F then the set of vectors $\gamma_i v_j$ such that $1 \leq i \leq e_1 d_1^2$ and $1 \leq j \leq N = t_1$ forms a basis for $\mathsf{F}^{m\times 1}$ over F, and if $\alpha \in \mathfrak{A}$ such that $\alpha v_j = 0$ for $1 \leq j \leq N$ then, since γ_i commutes with α,

$$\alpha(\gamma_i v_j) = \alpha\gamma_i v_j = \gamma_i \alpha v_j = \gamma_i(\alpha v_j) = \gamma_i 0 = 0$$

for all i and j, implying that $\alpha = 0$ as well.

Suppose next that $\mathfrak{A}$ is simple and $s_1 > 1$, so that $\mathsf{F}^{m\times 1} = \mathcal{M}_1 \oplus \mathcal{M}_2 \oplus \cdots \oplus \mathcal{M}_{s_1}$ is a direct sum of simple $\mathfrak{A}$-modules $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_{s_1}$. The above argument can be applied to $\mathcal{M}_1$ instead of $\mathsf{F}^{m\times 1}$ to prove the existence of elements $u_1, u_2, \ldots, u_{t_1}$ of $\mathcal{M}_1$ such that αu_j is nonzero for at least one element u_j of this set whenever α is a nonzero element of $\mathfrak{A}$. Now, since $\mathfrak{A}$ is simple, the modules $\mathcal{M}_1, \mathcal{M}_2, \ldots \mathcal{M}_{s_1}$ are isomorphic as modules over $\mathfrak{A}$, so that there exist $\mathfrak{A}$-module isomorphisms $\phi_j : \mathcal{M}_1 \to \mathcal{M}_j$ for $2 \leq j \leq s_1$. Set $u_i = 0$ for $t_1 + 1 \leq i \leq Ns_1 = \lceil t_1/s_1 \rceil s_1$ and let

$$v_i = u_{(i-1)s_1+1} + \sum_{j=2}^{s_1} \phi_j(u_{(i-1)s_1+j}) \in \mathsf{F}^{m\times 1}$$

for $1 \leq i \leq N$. Since $\mathsf{F}^{m\times 1}$ is a direct sum of the $\mathfrak{A}$-modules $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_{s_1}$, and since the above maps $\phi_2, \phi_3, \ldots, \phi_{s_1}$ are $\mathfrak{A}$-module isomorphisms, if $\alpha \in \mathfrak{A}$ such that $\alpha v_i = 0$ then

$$\alpha u_{(i-1)s_1+1} = \alpha u_{(i-1)s_1+2} = \cdots = \alpha u_{is_1} = 0$$

as well. Thus if $\alpha \in \mathfrak{A}$ such that $\alpha v_i = 0$ for $1 \leq i \leq N$, then $\alpha u_j = 0$ for $1 \leq j \leq Ns_1$, implying that $\alpha = 0$ by the choice of $u_1, u_2, \ldots, u_{t_1}$. Thus $v_1, v_2, \ldots, v_N$ is also a distinguishing set in this case.

Finally, suppose that $\mathfrak{A}$ is semi-simple over F with simple components $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$. Let $\omega_1, \omega_2, \ldots, \omega_k \in \mathfrak{A}$ be the identity elements of $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$, respectively, so that these are orthogonal central idempotents in $\mathfrak{A}$, and so that

$$\mathsf{F}^{m\times 1} = \omega_1\mathsf{F}^{m\times 1} \oplus \omega_2\mathsf{F}^{m\times 1} \oplus \cdots \oplus \omega_k\mathsf{F}^{m\times 1}.$$

Now $\omega_i\mathsf{F}^{m\times 1}$ has a structure as an $\mathfrak{A}_i$-module and the above argument can be used to prove the existence of elements $v_{i,1}, v_{i,2}, \ldots, v_{i,\lceil t_i/s_i \rceil}$ of $\omega_i\mathsf{F}^{m\times 1}$ such that at least one of $\alpha_i v_{i,1}, \alpha_i v_{i,2}, \ldots, \alpha_i v_{i,\lceil t_i/s_i \rceil}$ is nonzero whenever α_i is a nonzero element of $\mathfrak{A}_i$.

For $1 \leq j \leq N$, set

$$v_j = \sum_{\lceil s_i/t_i \rceil \geq j} v_{i,j} \in \mathsf{F}^{m\times 1},$$

and recall that each element α of $\mathfrak{A}$ has a unique representation as a sum $\alpha = \alpha_1 + \alpha_2 + \cdots + \alpha_k$ where $\alpha_i \in \mathfrak{A}_i$ for $1 \leq i \leq k$. Furthermore, $\alpha = 0$ in $\mathfrak{A}$ if and only if $\alpha_i = 0$ in $\mathfrak{A}_i$ for all i, and this can be used to establish that the above elements $v_1, v_2, \ldots, v_N$ form a distinguishing set for $\mathfrak{A}$. $\square$

A consideration of the case that $\mathfrak{A}$ is simple and isomorphic to a full matrix ring over F suggests that this is the best we can hope for: In this case, if one chooses any set of fewer than N vectors, then there will exist a nonzero element of $\mathfrak{A}$ that annihilates all of them.

On the other hand, the news is not all bad. Suppose that $\mathfrak{A}$ is given by a set of structure constants that can be used to define a regular matrix representation of the algebra. In this case we have $m = n$ and, indeed, $s_i = t_i$ for $1 \leq i \leq k$, so that $N = 1$. One can then check whether $\alpha = \beta$ in $\mathfrak{A}$ by checking whether $\alpha v = \beta v$ for a single (well-chosen) vector. The Las Vegas algorithms given later in the paper will therefore perform quite well in this case (see in particular Theorems 3.14 and 4.7 below).

2.3 Minimal Polynomials of Elements

Let

$$d = d_{\mathfrak{A}} = e_1 n_1 + e_2 n_2 + \cdots + e_k n_k \tag{2.5}$$

for the values k, $e_1, e_2, \ldots, e_k$ and $n_1, n_2 \ldots, n_k$ as defined in Section 2.1. A comparison of equations (2.1), (2.3) and (2.5) confirms that $d \leq \min(m, n)$.

The following result is well-known (or, easily deducible), but it is important enough to this work be be mentioned here.

Lemma 2.3. *Let $\mathfrak{A} \subseteq \mathsf{F}^{m \times m}$ be semi-simple algebra over F, and let d be defined as above. Then the minimal polynomial of any element of $\mathfrak{A}$ has degree at most d over F.*

Proof. It will be useful to consider four successively more general cases, namely, that $\mathfrak{A}$ is isomorphic to a full matrix ring over F, central simple over F, simple and, finally, an arbitrary semi-simple algebra over F.

In the first case $k = e_1 = 1$, $n = n_1^2$ and, since elements of $\mathfrak{A}$ may be identified with $n_1 \times n_1$ matrices over F, the result follows from the fact that the minimal polynomial of a matrix is always a divisor of a polynomial in $\mathsf{F}[x]$ with degree n_1, namely, its characteristic polynomial.

In the second case $k = e_1 = 1$ and $n = n_1^2$ as above. Let E be an algebraic closure of F and consider the algebra $\mathfrak{A}^{\mathsf{E}} = \mathfrak{A} \otimes_{\mathsf{F}} \mathsf{E}$ over E obtained from $\mathfrak{A}$ by extension of scalars. It is easy to show that the dimension of the vector space spanned by the elements $1, \alpha, \alpha^2, \ldots$ of $\mathfrak{A}$ over F is the same as the dimension of the vector space spanned by the elements $1 \otimes_{\mathsf{F}} 1, \alpha \otimes_{\mathsf{F}} 1, \alpha^2 \otimes_{\mathsf{F}} 1, \ldots$ of $\mathfrak{A}^{\mathsf{E}}$ over E, for any element α of $\mathfrak{A}$, so that the minimal polynomial of α over F is the same as that of $\alpha \otimes_{\mathsf{F}} 1$ over E. It is well-known that $\mathfrak{A}^{\mathsf{E}}$ is isomorphic to $\mathsf{E}^{n_1 \times n_1}$ as an algebra over E so that, once again, this minimal polynomial must have degree at most n_1.

Next suppose that $\mathfrak{A}$ is simple over F, so that $k = 1$ and $\mathfrak{A} = \mathfrak{A}_1$. In this case $\mathfrak{A}$ can be regarded as a central simple algebra of dimension n_1^2 over its centre E_1. Now, as argued above, the minimal polynomial of any element α of $\mathfrak{A}$ over E_1 has degree at most n_1, and the elements $1, \alpha, \alpha^2, \ldots, \alpha^{n_1 - 1}$ span $\mathsf{E}_1[\alpha]$ over E_1. Since $[\mathsf{E}_1 : \mathsf{F}] = e_1$ there exists a basis $\beta_1, \beta_2, \ldots, \beta_{e_1}$ of E_1 over F, and it is easy to see that the elements $\beta_i \alpha^j$ such that $1 \leq i \leq e_1$ and $0 \leq j < n_1$ span $\mathsf{E}_1[\alpha]$ over F. Consequently $\mathsf{E}_1[\alpha]$ has dimension at most $e_1 n_1$, and since $\mathsf{F}[\alpha]$ is a subspace of $\mathsf{E}_1[\alpha]$, $\mathsf{F}[\alpha]$ has dimension at most $e_1 n_1$ over F as well. Since the degree of the minimal polynomial of α over F is the same as the dimension of $\mathsf{F}[\alpha]$ over F, the result follows for the case that $\mathfrak{A}$ is simple.

Finally, suppose that $\mathfrak{A}$ is semi-simple over F, and let $\alpha \in \mathsf{F}$. Since $\mathfrak{A}$ is a direct sum of its simple components α can be written (uniquely) as a sum $\alpha = \alpha_1 + \alpha_2 + \cdots + \alpha_k$ where $\alpha_i \in \mathfrak{A}_i$ for $1 \leq i \leq k$. Now, since $\mathfrak{A}_i$ is a simple algebra with dimension $e_i n_i^2$ over F and has a centre with dimension e_i over F, the above argument implies that the minimal polynomial f_i of α_i has degree at most $e_i n_i$ over F, for all i. However, the minimal polynomial of α is clearly just the least common multiple of $f_1, f_2, \ldots, f_k$ and is a divisor of the product of $f_1, f_2, \ldots, f_k$. It follows immediately that f has degree at most $d = e_1 n_1 + e_2 n_2 + \cdots + e_k n_k$, as desired. □

2.4 Matrix Normal Forms

Consider the *Frobenius decomposition* of a matrix $\alpha \in \mathsf{F}^{m \times m}$:

$$\alpha = U^{-1} S U \qquad \text{for } S = \begin{bmatrix} C_{g_1} & & & 0 \\ & C_{g_2} & & \\ & & \ddots & \\ 0 & & & C_{g_l} \end{bmatrix}, \tag{2.6}$$

where $U \in \mathsf{F}^{m \times m}$ is a nonsingular matrix, S is a block diagonal matrix with matrices $C_{g_1}, C_{g_2}, \ldots, C_{g_l}$ on the diagonal, and where C_{g_i} is the companion matrix of a polynomial $g_i \in \mathsf{F}[x]$ of positive degree such that g_{i+1} divides g_i for $1 \leq i \leq l$. While the transition matrix U is not unique, the matrix S is, and is called the *Frobenius form* of the matrix α. We will call a matrix U a *Frobenius transition matrix* for α if it satisfies equation (2.6) above.

Since the Frobenius form S is unique the polynomials $g_1, g_2, \ldots, g_l$ are unique as well, and are called the *elementary divisors* of α. As equation (2.6) should suggest, g_1 is the minimal polynomial of α and the characteristic polynomial of α is the product $g_1 g_2 \ldots g_l$.

Giesbrecht (1995) has provided a Las Vegas algorithm for computation of the Frobenius form and a Frobenius transition matrix for an arbitrary matrix $\alpha \in \mathsf{F}^{m \times m}$ over a sufficiently large field, and contributes an analysis of the algorithm for the case that field elements are chosen uniformly and independently from a finite subset of the ground field of size m^2 when computing the Frobenius form of an $m \times m$ matrix. It will be useful to apply this algorithm when elements are chosen from a larger set.

Lemma 2.4. *Let ϵ be a constant such that $0 < \epsilon < 1$ and let F be any field with at least m^2/ϵ elements. Given a matrix $T \in \mathsf{F}^{m \times m}$, a Las Vegas algorithm can be used to find the Frobenius form and a Frobenius transition matrix for T or to report failure — the latter with probability at most ϵ. The algorithm requires $O(\mathcal{MM}(m) \log m)$ operations in F, or $O(m^3)$ operations using standard arithmetic.*

Proof. See the presentation of Giesbrecht's algorithm and the proof of Theorem 4.1 given by Giesbrecht (1995); the complexity analysis does not need to be changed. The algorithm can fail at only one point — an application of the subroutine "FindModCycl" — and the fact that this fails with probability at most ϵ follows by an application of the Schwartz-Zippel lemma (Schwartz 1980, Zippel 1979). □

Recall that a polynomial in $\mathsf{F}[x]$ is *separable* over F if all its roots in any field extension of F are simple. Equivalently, a polynomial is separable over F if and only if the polynomial and its first derivative are relatively prime. If a polynomial is separable over a field then it is also separable over every extension of the field.

A different matrix normal form for a matrix $\alpha \in \mathsf{F}^{m\times m}$ will be of use when the minimal polynomial g_1 of α is separable over F. Consider polynomials $h_1, h_2, \ldots, h_l \in \mathsf{F}[x]$ such that

$$h_l = g_l \quad \text{and} \quad h_i = g_i/g_{i+1} \quad \text{for } 1 \le i \le l-1; \tag{2.7}$$

then $g_1 = h_1 h_2, \ldots, h_l$, so that $h_1, h_2, \ldots, h_l$ are pairwise relatively prime and separable over F. We will call these polynomials (which are clearly well defined from α, since the elementary divisors are) the *distinct power divisors* of α. It is easily checked that

$$g_i = h_i h_{i+1} \ldots h_l \quad \text{for } 1 \le i \le l,$$

and that the characteristic polynomial of α is $f_1 f_2^2 \ldots f_l^l$. Let $\delta_i = \deg(h_i)$ for $1 \le i \le l$ and, when $\delta_i > 0$, let

$$\alpha^{(i)} = \begin{bmatrix} C_{h_i} & & & 0 \\ & C_{h_i} & & \\ & & \ddots & \\ 0 & & & C_{h_i} \end{bmatrix}$$

be a block matrix with i matrices of order δ_i on the diagonal, so that $\alpha^{(i)} \in \mathsf{F}^{i\delta_i \times i\delta_i}$ whenever this matrix is defined. Finally, let

$$\hat{\alpha} = \begin{bmatrix} \alpha^{(1)} & & & 0 \\ & \alpha^{(2)} & & \\ & & \ddots & \\ 0 & & & \alpha^{(l)} \end{bmatrix} \in \mathsf{F}^{m\times m}$$

be a block diagonal matrix whose diagonal blocks are all the matrices $\alpha^{(i)}$ such that $\delta_i > 0$. It is easily checked that α and $\hat{\alpha}$ have the same elementary divisors and hence the same Frobenius form T. Consequently, if U is a Frobenius transition matrix for α and $\hat{U}$ is a Frobenius transition matrix for $\hat{\alpha}$ then

$$U^{-1}\alpha U = \hat{U}^{-1}\hat{\alpha}\hat{U} = T,$$

so that

$$\alpha = V^{-1}\hat{\alpha}V \quad \text{for } V = \hat{U}U^{-1}. \tag{2.8}$$

The matrix $\hat{\alpha}$ is clearly uniquely determined from α whenever the minimal polynomial of α is separable. Kaltofen *et al.* (1990) call this the “rational Jordan form” and investigate its properties in a more general setting. However, since this name has been used for several different matrix forms in the literature, we shall call this the *distinct power form* of α. Any nonsingular matrix $V \in \mathsf{F}^{m\times m}$ such that $\alpha = V^{-1}\hat{\alpha}V$ as above will be called a *distinct power transition matrix* for α. We define a *distinct power decomposition* of α to include a distinct power transition matrix for α, the distinct power form of α, and the orders of the matrices on the diagonal of the distinct power form.

Theorem 2.5. *Let ϵ be a constant such that $0 < \epsilon < 1$ and let F be any field with at least $2m^2/\epsilon$ elements. Given a matrix $\alpha \in \mathsf{F}^{m\times m}$ whose minimal polynomial is separable over F, a Las Vegas algorithm can be used to find a distinct power decomposition of α or to report failure — the latter with probability at most ϵ. The algorithm requires $O(\mathcal{MM}(m)\log m)$ operations over F, or $O(m^3)$ operations using standard arithmetic.*

Proof. The desired Las Vegas algorithm and its analysis are easily described as follows.

One first computes the Frobenius form and a Frobenius transition matrix U for α, at the cost stated in Lemma 2.4. Since F includes at least $2m^2/\epsilon = m^2/(\epsilon/2)$ elements, this computation can be implemented to fail with probability at most $\epsilon/2$.

Since the elementary divisors of α are now available, the distinct power divisors are easily computed using equation (2.7). Since exact division of polynomials can be performed at asymptotically the same cost as polynomial multiplication, h_i can be computed from g_i and g_{i+1} using $O(\mathcal{M}(\deg(g_i)))$ operations over F for $1 \le i \le l-1$ and, since $g_1 g_2 \ldots g_l$ is the characteristic polynomial of α and has degree m, all of the distinct power divisors can be computed from the elementary divisors using $O(\mathcal{M}(m))$ operations in total.

At this point one can simply write down the distinct power form of α by inspecting the distinct power divisors, using $O(m^2)$ operations. The Frobenius form of this matrix and, more importantly, a Frobenius transition matrix $\hat{U}$ for it, can be computed at the cost stated in Lemma 2.4, failing again with probability at most $\epsilon/2$.

Finally, a distinct power transition matrix $V = \hat{U}U^{-1}$ can be generated from the above transition matrices U and $\hat{U}$ using $O(\mathcal{M}(m))$ additional operations.

□

3 Self-Centralizing Elements and Their Properties

Once again let d be as defined in equation (2.5).

Definition 3.1. An element α of $\mathfrak{A}$ is a *self-centralizing* element of $\mathfrak{A}$ if the minimal polynomial of α is separable with (maximal) degree d over F.

3.1 Centralizers of Self-Centralizing Elements

Recall that $C_{\mathfrak{A}}(\alpha)$ is the centralizer of α in $\mathfrak{A}$. Clearly $\mathsf{F}[\alpha] \subseteq C_{\mathfrak{A}}(\alpha)$ for all α. The next result therefore explains the choice of name for "self-centralizing elements."

Theorem 3.2. *If α is a self-centralizing element of $\mathfrak{A}$ then $C_{\mathfrak{A}}(\alpha) = \mathsf{F}[\alpha]$.*

Proof. As in the proof of Lemma 2.3 it will be useful to consider the cases that $\mathfrak{A}$ is isomorphic to a full matrix ring over F, central simple over F, simple, and an arbitrary semi-simple algebra.

In the first case $\mathfrak{A}$ is isomorphic to $\mathsf{F}^{n_1\times n_1} = \mathsf{F}^{d\times d}$. Let $\psi : \mathfrak{A} \to \mathsf{F}^{d\times d}$ be an algebra isomorphism; then $\psi(\alpha)$ is a $d \times d$ matrix whose minimal polynomial over F (the same as the minimal polynomial of α) is separable with degree d. Consequently, $\psi(\alpha)$ is similar to a diagonal matrix with distinct entries on its diagonal and it is easily proved that $\mathsf{F}[\psi(\alpha)] =$

$\psi(\mathsf{F}[\alpha])$ and $C_{\psi(\mathfrak{A})}(\psi(\alpha)) = \psi(C_{\mathfrak{A}}(\alpha))$ are both equal to the set of diagonal matrices in $\mathsf{F}^{d\times d}$. Since $\psi(\mathsf{F}[\alpha]) = \psi(C_{\mathfrak{A}}(\alpha))$ and ψ is an isomorphism, $\mathsf{F}[\alpha] = C_{\mathfrak{A}}(\alpha)$.

In the case that $\mathfrak{A}$ is central simple over F it is useful (again, as in the proof of Lemma 2.3) to consider the algebra $\mathfrak{A}^{\mathsf{E}}$ over E, where E is an algebraic closure of F. Once again, $\mathfrak{A}^{\mathsf{E}}$ is isomorphic to $\mathsf{E}^{n_1\times n_1} = \mathsf{E}^{d\times d}$ as an algebra over E.

If α is self-centralizing in $\mathfrak{A}$ then, by definition, the minimal polynomial of α is separable with degree d. Since this is also the minimal polynomial of $\alpha \otimes_{\mathsf{F}} 1 \in \mathfrak{A}^{\mathsf{E}}$ over E, and since this polynomial is separable over E as well as over F, $\alpha \otimes_{\mathsf{F}} 1$ is self-centralizing in $\mathfrak{A}^{\mathsf{E}}$. The centralizer of $\alpha \otimes_{\mathsf{F}} 1$ is therefore equal to $\mathsf{E}[\alpha \otimes_{\mathsf{F}} 1]$ in $\mathfrak{A}^{\mathsf{E}}$ by the argument given above.

Now, since α has the same minimal polynomial over F as $\alpha \otimes_{\mathsf{F}} 1$ has over E, the dimension of $\mathsf{F}[\alpha]$ over F is the same as that of $\mathsf{E}[\alpha \otimes_{\mathsf{F}} 1]$ over E. The dimension of $C_{\mathfrak{A}}(\alpha)$ over F is the same as the dimension of $C_{\mathfrak{A}^{\mathsf{E}}}(\alpha \otimes_{\mathsf{F}} 1)$ over E as well, since the elements of either set can be obtained as linear combinations of elements of a basis by solving essentially the same homogeneous system of linear equations. Therefore $\mathsf{F}[\alpha]$ has the same dimension as $C_{\mathfrak{A}}(\alpha)$ over F and, since $\mathsf{F}[\alpha] \subseteq C_{\mathfrak{A}}(\alpha)$, $\mathsf{F}[\alpha] = C_{\mathfrak{A}}(\alpha)$.

In the third case, that $\mathfrak{A}$ is simple, $\mathfrak{A}$ may regarded as a central simple algebra over its centre E_1. If α is self-centralizing in $\mathfrak{A}$ then $\mathsf{F}[\alpha] \subseteq \mathsf{E}_1[\alpha]$ and $\mathsf{F}[\alpha]$ has dimension e_1n_1 over F. $\mathsf{E}_1[\alpha]$ therefore has dimension at least e_1n_1 over F as well. On the other hand, Lemma 2.3 implies that the minimal polynomial of α over E_1 has degree at most n_1. Suppose therefore that the degree is $r \le n_1$. Then $\mathsf{E}_1[\alpha]$ has dimension r over E_1 and, since $[\mathsf{E}_1 : \mathsf{F}] = e_1$, $\mathsf{E}_1[\alpha]$ has dimension at most $e_1r \le e_1n_1$ over F. Consequently $\mathsf{E}_1[\alpha]$ has dimension exactly $e_1r = e_1n_1$ over F, so $r = n_1$. Therefore $\mathsf{F}[\alpha] = \mathsf{E}_1[\alpha]$, again since one of these is a subspace of the other and both have the same dimension over F.

Now, the minimal polynomial of α over E_1 has full degree n_1 and is separable, since it is a divisor of the minimal polynomial of α over F. The element α is therefore self-centralizing in $\mathfrak{A}$ when $\mathfrak{A}$ is regarded as a central simple algebra over E_1. Since the centralizer $C_{\mathfrak{A}}(\alpha)$ is the same regardless of whether $\mathfrak{A}$ is considered as an algebra over F or over E_1, we now have that $C_{\mathfrak{A}}(\alpha) = \mathsf{E}_1[\alpha] = \mathsf{F}[\alpha]$ as desired.

In the final case it suffices to observe, again, that an element $\alpha \in \mathfrak{A}$ can be written uniquely as $\alpha = \alpha_1 + \alpha_2 + \cdots + \alpha_k$, where $\alpha_i \in \mathfrak{A}_i$ for $1 \le i \le k$. Let f be the minimal polynomial of α over F, let f_i be the minimal polynomial of α_i over F, and let δ_i be the degree of f_i over F for all i. Then, since f is the least common multiple of $f_1, f_2, \dots, f_k$ and has degree $d = e_1n_1 + e_2n_2 + \cdots + e_kn_k$ (if α is self-centralizing in $\mathfrak{A}$),

$$\delta_1 + \delta_2 + \cdots + \delta_k = \deg(f_1 f_2 \cdots f_k) \ge \deg(f) = e_1n_1 + e_2n_2 + \cdots + e_kn_k.$$

On the other hand, it follows by Lemma 2.3 that $\delta_i \le e_in_i$ as well for all i, so clearly $\deg(f_i) = \delta_i = e_in_i$ for each i. Since f_i is a divisor of f and f is separable, f_i is separable as well. Thus α_i is self-centralizing in $\mathfrak{A}_i$ and, since $\mathfrak{A}_i$ is simple,

$$C_{\mathfrak{A}_i}(\alpha_i) = \mathsf{F}[\alpha_i] \tag{3.1}$$

for $1 \le i \le k$.

The above inequalities imply that the product and least common multiple of $f_1, f_2, \dots, f_k$ have the same degree. Since the latter polynomial is always a factor of the former, this implies that these are the same. Therefore $f_1, f_2, \dots, f_k$ are pairwise relatively prime and

$$F[\alpha] = \mathsf{F}[\alpha_1] \oplus \mathsf{F}[\alpha_2] \oplus \cdots \oplus \mathsf{F}[\alpha_k]. \tag{3.2}$$

On the other hand, since $\mathfrak{A}$ is the direct sum of its simple components,

$$C_{\mathfrak{A}}(\alpha) = C_{\mathfrak{A}_1}(\alpha) \oplus C_{\mathfrak{A}_2}(\alpha) \oplus \cdots \oplus C_{\mathfrak{A}_k}(\alpha) = C_{\mathfrak{A}_1}(\alpha_1) \oplus C_{\mathfrak{A}_2}(\alpha_2) \oplus \cdots \oplus C_{\mathfrak{A}_k}(\alpha_k). \tag{3.3}$$

Equations (3.1), (3.2), and (3.3) clearly imply that $\mathsf{F}[\alpha] = C_{\mathfrak{A}}(\alpha)$ as desired. □

The next result follows from the above discussion.

Theorem 3.3. *Let $\alpha = \alpha_1 + \alpha_2 + \cdots + \alpha_k$, where $\alpha_i \in \mathfrak{A}_i$ for $1 \leq i \leq k$. Then α is self-centralizing in $\mathfrak{A}$ if and only if α_i is self-centralizing in $\mathfrak{A}_i$ for all i and the minimal polynomials of $\alpha_1, \alpha_2, \dots, \alpha_k$ over F are pairwise relatively prime.*

Proof. As argued above, if α is self-centralizing then, by inspection of the degrees of the minimal polynomials of α and of $\alpha_1, \alpha_2, \dots, \alpha_k$, the minimal polynomials of $\alpha_1, \alpha_2, \dots, \alpha_k$ must be pairwise relatively prime and of maximal degree. They are also separable since they each divide the minimal polynomial of α. Conversely, if the minimal polynomials of $\alpha_1, \alpha_2, \dots, \alpha_k$ are pairwise relatively prime and separable then the least common multiple of these polynomials is also their product, so that if each of these polynomials also has maximal degree then the minimal polynomial of α is separable with maximal degree as well. □

Suppose next that α is self-centralizing in $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ and consider the distinct power divisors $h_1, h_2, \dots, h_l$ of α as defined in Section 2.4. Let f_i be the minimal polynomial of α_i for $1 \leq i \leq k$.

Lemma 3.4. *If α is self-centralizing in $\mathfrak{A}$ and $f_1, f_2, \dots, f_k$ are as above then each polynomial f_i is a divisor of exactly one of the distinct power divisors of α and is relatively prime with each of the rest. In particular,*

$$h_i = \prod_{m_j = i} f_j. \tag{3.4}$$

Proof. Since α is self-centralizing, the polynomials $f_1, f_2, \dots, f_k$ are separable and pairwise relatively prime. It therefore suffices to prove that every irreducible factor of f_j is a divisor with the same multiplicity m_j of the characteristic polynomial of α, for equation (3.4) then follows from the definition of $h_1, h_2, \dots, h_l$ as the distinct power divisors of α.

Since $\mathsf{F}^{m\times 1}$ is a direct sum of simple $\mathfrak{A}$-modules, as shown in equation (2.2), and all simple $\mathfrak{A}$-modules that are faithful $\mathfrak{A}_i$-modules are isomorphic, there exists a nonsingular matrix X, whose columns are elements of carefully chosen bases for the simple modules $M_1^{(1)}, M_2^{(1)}, \dots, M_{s_k}^{(k)}$ shown in equation (2.2), such that

$$X^{-1}\alpha X = \begin{bmatrix} \alpha_1^{(1)} & & & 0 \\ & \alpha_2^{(1)} & & \\ & & \ddots & \\ 0 & & & \alpha_{s_k}^{(k)} \end{bmatrix},$$

where

$$\alpha_1^{(i)} = \alpha_2^{(i)} = \cdots = \alpha_{s_i}^{(i)} \in \mathsf{F}^{e_i n_i d_i \times e_i n_i d_i}$$

is a matrix that expresses the action of α on the simple module $M_j^{(i)}$ with respect to the basis of this module that is included as columns of X. Consequently, since each $M_j^{(i)}$ is a faithful and simple $\mathfrak{A}_i$-module, the minimal polynomial of $\alpha_j^{(i)}$ is the polynomial f_i, for $1 \leq i \leq k$ and $1 \leq j \leq s_i$. Now it is necessary and sufficient to establish that each matrix $\alpha_j^{(i)}$ has Frobenius form

$$\begin{bmatrix} C_{f_i} & & & 0 \\ & C_{f_i} & & \\ & & \ddots & \\ 0 & & & C_{f_i} \end{bmatrix}$$

with d_i elementary divisors that are all equal to f_i. Indeed, it will be sufficient to prove that $\alpha_j^{(i)}$ is similar to a matrix

$$\begin{bmatrix} C_{i,j} & & & 0 \\ & C_{i,j} & & \\ & & \ddots & \\ 0 & & & C_{i,j} \end{bmatrix} \in \mathsf{F}^{e_i n_i d_i \times e_i n_i d_i} \tag{3.5}$$

for any matrix $C_{i,j} \in \mathsf{F}^{e_i n_i \times e_i n_i}$ at all — for then it will be clear (by a comparison of degrees and taking advantage of the fact that f_i is separable) that $C_{i,j}$ has minimal polynomial f_i and is similar to C_{f_i} as needed.

With this in mind, let us consider $\mathfrak{A}_i$ as a central simple algebra over its centre, E_i, and consider $M_j^{(i)}$ as a simple module of dimension $d_i n_i$ over this extension of F. Recall that the minimal polynomial of α_i over E_i is a separable polynomial g_i of degree n_i over E_i such that f_i is divisible by g_i in $\mathsf{E}_i[x]$ — this was established and exploited in the proof of Theorem 3.2, above.

Now let K_i be an algebraic closure of E_i and consider the simple algebra $\mathfrak{A}_i^{\mathsf{K}_i} = \mathfrak{A}_i \otimes_{\mathsf{E}_i} \mathsf{K}_i \cong \mathsf{K}_i^{n_i \times n_i}$, and its module $M_j^{(i)} \otimes_{\mathsf{E}_i} \mathsf{K}_i$, over K_i. The latter module is a direct sum of d_i simple $\mathfrak{A}_i^{\mathsf{K}_i}$-modules that each have dimension n_i over K_i and these modules are isomorphic, since they are simple modules over the same simple algebra. Consequently there exists a basis

$$v_{1,1}^{\mathsf{K}_i}, v_{1,2}^{\mathsf{K}_i}, \ldots, v_{1,n_i}^{\mathsf{K}_i}, \ldots, v_{d_i,1}^{\mathsf{K}_i}, v_{d_i,2}^{\mathsf{K}_i}, \ldots, v_{d_i,n_i}^{\mathsf{K}_i} \in M_j^{(i)} \otimes_{\mathsf{E}_i} \mathsf{K}_i$$

for $M_j^{(i)} \otimes_{\mathsf{E}_i} K_i$, consisting of carefully chosen bases for each of the above d_i simple modules, such that the action of $\alpha_i \otimes_{\mathsf{E}_i} 1 \in \mathfrak{A}_i^{\mathsf{K}_i}$ with respect to this basis is given by a block diagonal matrix

$$\begin{bmatrix} C_{i,j}^{\mathsf{K}} & & & 0 \\ & C_{i,j}^{\mathsf{K}} & & \\ & & \ddots & \\ 0 & & & C_{i,j}^{\mathsf{K}} \end{bmatrix} \in \mathsf{K}^{d_i n_i \times d_i n_i}$$

with d_i copies of a matrix $C_{i,j}^{\mathsf{K}} \in \mathsf{K}^{n_i \times n_i}$ on its diagonal. Since the minimal polynomial of $\alpha_i \otimes_{\mathsf{E}_i} \mathsf{K}_i$ over K_i is the same as that of α_i over E_i, namely $g_i \in \mathsf{E}_i[x]$, and this polynomial has degree n_i, the matrix $C_{i,j}^{\mathsf{K}}$ is similar to the companion matrix C_{g_i} in $\mathsf{K}^{n_i \times n_i}$. Therefore there is also a basis for $M_j^{(i)} \otimes_{\mathsf{E}_i} \mathsf{K}_i$ such that the action of $\alpha_i \otimes_{\mathsf{E}_i} 1$ on this module with respect to this basis is given by the matrix

$$\widehat{M} = \begin{bmatrix} C_{g_i} & & & 0 \\ & C_{g_i} & & \\ & & \ddots & \\ 0 & & & C_{g_i} \end{bmatrix} \in \mathsf{E}_i^{n_i d_i \times n_i d_i} \subseteq \mathsf{K}_i^{n_i d_i \times n_i d_i}. \tag{3.6}$$

Happily, this implies that there exists a basis

$$v_{1,1}, v_{1,2}, \ldots, v_{1,n_i}, \ldots, v_{d_i,1}, v_{d_i,2}, \ldots, v_{d_i,n_i} \in M_j^{(i)} \tag{3.7}$$

for the module $M_j^{(i)}$ over E_i such that the action on α_i over $M_j^{(i)}$ with respect to this basis is given by the matrix $\widehat{M}$ as well: The action of α_i on $M_j^{(i)}$ over E_i with respect to an arbitrary basis is necessarily represented by some matrix $\overline{M}$ in $\mathsf{E}^{n_i d_i \times n_i d_i}$ that is similar to $\widehat{M}$ in $\mathsf{K}^{n_i d_i \times n_i d_i}$. Since $\widehat{M}$ and $\overline{M}$ both belong to $\mathsf{E}_i^{n_i d_i \times n_i d_i}$ they must be similar as matrices in this ring as well, so that a change of basis for $M_j^{(i)}$ over E_i will bring the matrix into the desired form.

Now consider the E_i-linear map $\phi_i : M_j^{(i)} \to M_j^{(i)}$ such that

$$\phi_i(v_{r,s}) = \begin{cases} v_{r+1,s} & \text{if } 1 \le r < n_i \text{ and } 1 \le s \le d_i, \\ v_{1,s} & \text{if } r = n_i \text{ and } 1 \le s \le d_i. \end{cases}$$

Since the action of this map with respect to basis in equation (3.7) is given by the (permutation) matrix

$$\begin{bmatrix} 0_{n_i} & & & & I_{n_i} \\ I_{n_i} & 0_{n_i} & & & \\ & I_{n_i} & & & \\ & & \ddots & & \\ 0_{n_i} & & & I_{n_i} & 0_{n_i} \end{bmatrix} \in \mathsf{E}_i^{n_i d_i \times n_i d_i}$$

where 0_{n_i} and I_{n_i} are the zero and identity matrices in $\mathsf{E}_i^{n_i \times n_i}$ respectively, it is clear that the actions of α_i and ϕ_i on $M_j^{(i)}$ commute.

Now let $u_1, u_2, \ldots, u_{e_i} \in \mathsf{E}_i \subseteq \mathfrak{A}_i$ be a basis for E_i over F and consider the action of α_i on $M_j^{(i)}$, as a module over F, with respect to the basis

$$u_1 v_{1,1}, u_2 v_{1,1}, \ldots, u_{e_i} v_{1,1}, u_1 v_{1,2}, u_2 v_{1,2}, \ldots, u_{e_i} v_{1,2}, \ldots, u_1 v_{d_i,n_i}, u_2 v_{d_i,n_i}, \ldots, u_{e_i} v_{d_i,n_i} \tag{3.8}$$

obtained by replacing each element $v_{i,j}$ of the basis in equation (3.7), above, by the block of vectors $u_1 v_{i,j}, u_2 v_{i,j}, \ldots, u_{e_i} v_{i,j}$. Since the subspace of $M_j^{(i)}$ over E_i spanned by the vectors

$v_{h,1}, v_{h,2}, \ldots, v_{h,n_i}$ is invariant under α_i for $1 \le h \le d_i$ (see, again, the matrix form in equation (3.6)), the subspace of $M_j^{(i)}$ over F spanned by the vectors

$$u_1 v_{h,1}, u_2 v_{h,1}, \ldots, u_{e_i} v_{h,1}, \ldots, u_1 v_{h,n_i}, u_2 v_{h,n_i}, \ldots, u_{e_i} v_{h,n_i}$$

is invariant under α_i as well. Thus, the action of α_i on $M_j^{(i)}$ with respect to the basis in equation (3.8) is given by a block-diagonal matrix

$$\begin{bmatrix} C_{i,j}^{(1)} & & & 0 \\ & C_{i,j}^{(2)} & & \\ & & \ddots & \\ 0 & & & C_{i,j}^{(d_i)} \end{bmatrix}$$

for matrices $C_{i,j}^{(1)}, C_{i,j}^{(2)}, \ldots, C_{i,j}^{(d_i)} \in \mathsf{F}^{e_i n_i \times e_i n_i}$. Furthermore, the above map ϕ_i commutes with α_i as an F-linear map. Since the action of this map with respect to the above basis is given by a (permutation) matrix

$$M_\phi = \begin{bmatrix} 0_{e_i n_i} & & & & I_{e_i n_i} \\ I_{e_i n_i} & 0_{e_i n_i} & & & \\ & I_{e_i n_i} & & & \\ & & \ddots & & \\ 0_{e_i n_i} & & & I_{e_i n_i} & 0_{e_i n_i} \end{bmatrix} \in \mathsf{F}^{e_i d_i n_i \times e_i d_i n_i},$$

where $0_{e_i n_i}$ and $I_{e_i n_i}$ are the zero and identity matrices in $\mathsf{F}^{e_i n_i \times e_i n_i}$ respectively, it follows that

$$\begin{bmatrix} C_{i,j}^{(1)} & & & 0 \\ & C_{i,j}^{(2)} & & \\ & & \ddots & \\ 0 & & & C_{i,j}^{(d_i)} \end{bmatrix} = M_\phi^{-1} \begin{bmatrix} C_{i,j}^{(1)} & & & 0 \\ & C_{i,j}^{(2)} & & \\ & & \ddots & \\ 0 & & & C_{i,j}^{(d_i)} \end{bmatrix} M_\phi = \begin{bmatrix} C_{i,j}^{(2)} & & & 0 \\ & C_{i,j}^{(3)} & & \\ & & \ddots & \\ 0 & & & C_{i,j}^{(1)} \end{bmatrix},$$

so that $C_{i,j}^{(1)} = C_{i,j}^{(2)} = \cdots = C_{i,j}^{(d_i)} = C_{i,j}$ for some matrix $C_{i,j} \in \mathsf{F}^{e_i n_i \times e_i n_i}$.

Since the matrix $\alpha_j^{(i)}$ also expresses the action of α_i on the module $M_j^{(i)}$ with respect to a basis over F, it now follows that $\alpha_j^{(i)}$ is similar to a matrix with the form given in equation (3.5) above, as desired to complete the proof. $\square$

Suppose yet again that α is self-centralizing in $\mathfrak{A}$ with distinct power divisors $h_1, h_2, \ldots, h_l$ and that the distinct power form $\hat{\alpha}$ of α is as shown in Section 2.4,

$$\hat{\alpha} = \begin{bmatrix} \alpha^{(1)} & & & 0 \\ & \alpha^{(2)} & & \\ & & \ddots & \\ 0 & & & \alpha^{(l)} \end{bmatrix},$$

where each matrix $\alpha^{(j)} \in \mathsf{F}^{j\delta_j \times j\delta_j}$ for $\delta_j = \deg(h_j)$ and $\alpha^{(j)}$ has minimal polynomial h_j. Let V be any distinct power transition matrix for α, so that $\alpha = V^{-1}\hat{\alpha}V$, and let

$$\tau_i = V^{-1} \begin{bmatrix} \Delta_{i,1} & & & 0 \\ & \Delta_{i,2} & & \\ & & \ddots & \\ 0 & & & \Delta_{i,l} \end{bmatrix} V \in \mathsf{F}^{m\times m}, \tag{3.9}$$

where $\Delta_{i,j} \in \mathsf{F}^{j\delta_j \times j\delta_j}$ is the identity matrix if $i = j$ and is the zero matrix otherwise, for $1 \le i, j \le l$. Clearly $\tau_1, \tau_2, \ldots, \tau_l$ are pairwise orthogonal idempotents in $\mathsf{F}^{m\times m}$ whose sum is the identity matrix.

Theorem 3.5. *If α is self-centralizing in $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ and the idempotents $\tau_1, \tau_2, \ldots, \tau_l$ are formed from α as above, then these are central idempotents in $\mathfrak{A}$.*

Proof. Since the polynomials $h_1, h_2, \ldots, h_l$ are pairwise relatively prime there exist polynomials $g_1, g_2, \ldots, g_l$ such that

$$g_i \equiv \begin{cases} 1 \pmod{h_j} & \text{if } j = i, \\ 0 \pmod{h_j} & \text{if } j \ne i, \end{cases}$$

for $1 \le i, j \le l$. If $\alpha^{(1)}, \alpha^{(2)}, \ldots, \alpha^{(l)}$ are on the diagonal of the distinct power form $\hat{\alpha}$ of α, as above, then $\alpha^{(i)}$ has minimal polynomial h_i for all i, and $g_i(\alpha^{(j)}) = \Delta_{i,j}$ for $1 \le i, j \le l$. Thus

$$g_i(\hat{\alpha}) = \begin{bmatrix} \Delta_{i,1} & & & 0 \\ & \Delta_{i,2} & & \\ & & \ddots & \\ 0 & & & \Delta_{i,l} \end{bmatrix}$$

and $g_i(\alpha) = g_i(V^{-1}\hat{\alpha}V) = V^{-1}g_i(\hat{\alpha})V = \tau_i$. On the other hand, it follows by Lemma 3.4 that

$$g_i \equiv \begin{cases} 1 \pmod{f_j} & \text{if } 1 \le j \le k \text{ and } m_j = i, \\ 0 \pmod{f_j} & \text{if } 1 \le j \le k \text{ and } m_j \ne i. \end{cases}$$

Since $\alpha = \alpha_1 + \alpha_2 + \cdots + \alpha_k$, where $\alpha_j \in \mathfrak{A}_j$ with minimal polynomial f_j for $1 \le j \le k$, $g_i(\alpha_j)$ is the identity element of $\mathfrak{A}_j$ (and a central primitive idempotent in $\mathfrak{A}$) if $m_j = i$, and $g_i(\alpha_j) = 0$ otherwise. Now since

$$g_i(\alpha) = g_i(\alpha_1) + g_i(\alpha_2) + \cdots + g_i(\alpha_k),$$

it follows that $\tau_i = g_i(\alpha)$ is the sum of (distinct) central primitive idempotents in $\mathfrak{A}$, so that τ_i is a central idempotent of $\mathfrak{A}$ as claimed. $\square$

This implies that if $\omega_1, \omega_2, \ldots, \omega_k$ are the central primitive idempotents of $\mathfrak{A}$, and the identity elements of $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$, respectively, then

$$\tau_i = \sum_{m_j = i} \omega_j. \tag{3.10}$$

Thus $\tau_1, \tau_2, \ldots, \tau_l$ do not depend on the choice of the self-centralizing element α or the distinct power transition matrix V used to define them.

3.2 Existence and Density of Self-Centralizing Elements

Theorem 3.6. *If $\mathfrak{A}$ is a separable matrix algebra over an infinite field F then $\mathfrak{A}$ contains a self-centralizing element.*

Proof. It will be useful once again to consider several cases.

Suppose first that $\mathfrak{A}$ is simple, so that $k = 1$. In this case $\mathfrak{A} = \mathfrak{A}_1 \cong \mathsf{D}_1^{t_1 \times t_1}$, where D_1 is a division algebra that is central simple over the centre E_1 of $\mathfrak{A}$ and where the dimension of D_1 over E_1 is a perfect square. Once again let this dimension be d_1^2, so that $n_1 = d_1 t_1$ and $n = e_1 n_1^2 = e_1 d_1^2 t_1^2$.

As shown, for example, by Pierce (1982), D includes a subfield K that is separable over E_1 such that $[\mathsf{K} : \mathsf{E}_1] = d_1$. Since $\mathfrak{A}$ is a separable algebra, the field E_1 is separable over F. It follows, for example, by Proposition 2.5.8 of Bastida (1984) that K is also separable over F. Furthermore, $[\mathsf{K} : \mathsf{F}] = [\mathsf{K} : \mathsf{E}_1][\mathsf{E}_1 : \mathsf{F}] = e_1 d_1$. Consequently there exists an element a of $\mathsf{K} \subseteq \mathsf{D}$ such that $\mathsf{F}[a] = \mathsf{K}$ and such that the minimal polynomial of a is separable with degree $e_1 d_1$ over F.

Now let $\psi : \mathfrak{A} \to \mathsf{D}_1^{t_1 \times t_1}$ be an isomorphism of algebras over F. It suffices to choose α as

$$\alpha = \psi^{-1}\left(\begin{bmatrix} a_1 & & & 0 \\ & a_2 & & \\ & & \ddots & \\ 0 & & & a_{t_1} \end{bmatrix}\right) \tag{3.11}$$

where $a_1, a_2, \ldots, a_{t_1} \in \mathsf{K} \subseteq \mathsf{D}_1$ such that $\mathsf{F}[a_1] = \mathsf{F}[a_2] = \cdots = \mathsf{F}[a_{t_1}] = \mathsf{K}$ and the minimal polynomials of $a_1, a_2, \ldots, a_t$ over F are distinct. Then, since K is a separable extension of F these minimal polynomials will be separable and irreducible over F, and the minimal polynomial of α over F will be their product, a separable polynomial with degree $e_1 d_1 t_1 = e_1 n_1$.

If $\mathsf{K} = \mathsf{F}$ then it suffices to choose $a_1, a_2, \ldots, a_{t_1}$ as distinct elements from F. On the other hand, if $\mathsf{K} \neq \mathsf{F}$, then we can set $a_1 = a$ for the element a described above such that $\mathsf{F}[a] = \mathsf{K}$. If $b \in \mathsf{F}$ then $\mathsf{F}[a+b] = \mathsf{F}[a]$, since clearly $a+b \in \mathsf{F}[a]$ and $a = (a+b) - b \in \mathsf{F}[a+b]$. Furthermore, if $g(x) \in \mathsf{F}[x]$ is the minimal polynomial of a over F and g has distinct roots $c_1, c_2, \ldots, c_{e_1 d_1}$ in an extension of F, then the minimal polynomial of $a + b$ over F is $g(x - b)$ and this polynomial has distinct roots $c_1 + b, c_2 + b, \ldots, c_{e_1 d_1} + b$ in the same extension. Thus the minimal polynomial of $a + b$ is also separable over F. It is therefore sufficient to set $a_i = a + b_i$, for $2 \leq i \leq t_1$, where $b_2, b_3, \ldots, b_{t_1}$ are chosen from F in such a way that the minimal polynomials of $a_1, a_2, \ldots, a_{t_1}$ over F are pairwise relatively prime. Since these polynomials are each irreducible in $\mathsf{F}[x]$, this will be the case as long as each polynomial has a root in an extension of F that is not also a root of any of the rest. Now, since F is infinite, it is clear that suitable elements $b_2, b_3, \ldots, b_{t_1}$ of F can be found. Thus a self-centralizing element of $\mathfrak{A}$ exists if $\mathfrak{A}$ is simple.

Suppose that $\mathfrak{A}$ is separable but not simple over an infinite field F. The above argument implies that a self-centralizing element β_i of $\mathfrak{A}_i$ exists for each of the simple components $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$. It now suffices to set $\alpha_1 = \beta_1$ and to set $\alpha_i = \beta_i + b_i 1_{\mathfrak{A}_i}$, for $2 \leq i \leq k$, where $b_2, b_3, \ldots, b_k$ are chosen from F to ensure that the minimal polynomials of $\alpha_1, \alpha_2, \ldots, \alpha_k$ are pairwise relatively prime. Each element α_i will be self-centralizing in $\mathfrak{A}_i$ by essentially the argument used in the construction of α in the case that $\mathfrak{A}$ is simple above, and Theorem 3.3 will then be applicable. Since F is infinite it is easy to prove that suitable elements $b_2, b_3, \ldots, b_k$ can be found. □

The next result establishes that self-centralizing elements of separable algebras are also easy to find. Once again, let d be as given in equation (2.5).

Theorem 3.7. *Let $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ be a separable algebra of dimension n over a field F, and suppose a self-centralizing element is included in the F-linear span of elements $\gamma_1, \gamma_2, \ldots, \gamma_h$ of F. Let S be a finite subset of F with size at least $3d^3/2\epsilon$, for $\epsilon > 0$. If the elements $s_1, s_2, \ldots, s_h$ are chosen uniformly and independently from S then the element*

$$s_1\gamma_1 + s_2\gamma_2 + \cdots + s_h\gamma_h$$

is self-centralizing in $\mathfrak{A}$ with probability at least $1 - \epsilon$.

Proof. A polynomial $f \in \mathsf{F}[x_1, x_2, \ldots, x_h]\backslash\{0\}$ with total degree at most $\frac{3}{2}d^3$ will be produced such that, for all elements $s_1, s_2, \ldots, s_h$ of F, if $f(s_1, s_2, \ldots, s_h) \neq 0$ then $s_1\gamma_1+s_2\gamma_2+\cdots+s_h\gamma_h$ is a self-centralizing element of $\mathfrak{A}$. The result will then follow by an application of the Schwartz-Zippel lemma (see Schwartz 1980, Zippel 1979).

Since the F-linear span of $\gamma_1, \gamma_2, \ldots, \gamma_h$ includes a self-centralizing element, there exist elements $\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_h$ of F such that the element

$$\hat{s} = \hat{s}_1\gamma_1 + \hat{s}_2\gamma_2 + \cdots + \hat{s}_h\gamma_h$$

is self-centralizing in $\mathfrak{A}$. Let $y_1, y_2, \ldots, y_h$ be indeterminates over F and let

$$\sigma = \gamma_1 y_1 + \gamma_2 y_2 + \cdots + \gamma_h y_h \in \mathsf{F}[y_1, y_2, \ldots, y_h]^{m\times m},$$

so that $\hat{s} = \sigma(\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_h)$. Now consider the system of polynomial equations

$$\sigma^d + z_{d-1}\sigma^{d-1} + \cdots + z_1\sigma + z_0 1 = 0 \tag{3.12}$$

in the indeterminates $y_1, y_2, \ldots, y_h, z_0, z_1, \ldots, z_{d-1}$. This system includes m^2 equations (since σ^i is an $m \times m$ matrix) which are linear in the indeterminates $z_0, z_1, \ldots, z_{d-1}$. Replacing each indeterminate y_i by the field element $\hat{s}_i$, for $1 \leq i \leq h$, we obtain a system of linear equations

$$\hat{s}^d + z_{d-1}\hat{s}^{d-1} + \cdots + z_1\hat{s} + z_0 1 = 0 \tag{3.13}$$

in the indeterminates $z_0, z_1, \ldots, z_{d-1}$. This system has a unique solution whose entries (and the leading term, 1) are the coefficients of the minimal polynomial of $\hat{s}$ over F. Therefore there is a subset of d of these equations with full rank d. The corresponding equations in the system (3.12) form a system

$$M \begin{bmatrix} z_0 \\ z_1 \\ \vdots \\ z_{d-1} \end{bmatrix} = v \tag{3.14}$$

where $M \in \mathsf{F}[y_1, y_2, \ldots, y_h]^{d\times d}$ and $v \in \mathsf{F}[y_1, y_2, \ldots, y_h]^{d\times 1}$.

Let $g = \det M \in \mathsf{F}[y_1, y_2, \ldots, y_h]$; then g is not identically zero, since a nonsingular matrix in $\mathsf{F}^{d\times d}$ is obtained from M by replacing y_i with $\hat{s}_i$ for all i. Furthermore if $s_1, s_2, \ldots, s_h \in \mathsf{F}$ such that $g(s_1, s_2, \ldots, s_h) \neq 0$ then it follows by the definition of g that the elements $1, \sigma(s_1, s_2, \ldots, s_h), \sigma(s_1, s_2, \ldots, s_h)^2, \ldots, \sigma(s_1, s_2, \ldots, s_h)^{d-1}$ of $\mathfrak{A}$ are linearly independent over F. In this case, Lemma 2.3 implies that the the minimal polynomial of $\sigma(s_1, s_2, \ldots, s_h) = s_1\gamma_1 + s_2\gamma_2 + \ldots s_h\gamma_h$ over F has degree exactly d.

Cramer's rule can now be applied to the system shown in (3.14) to obtain polynomials $h_0, h_1, \ldots, h_{d-1} \in \mathsf{F}[y_1, y_2, \ldots, y_h]$ such that the minimal polynomial of $s_1y_1+s_2y_2+\cdots+s_hy_h$ over F is

$$x^d + \frac{h_{d-1}(s_1, s_2, \ldots, s_h)}{g(s_1, s_2, \ldots, s_h)}x^{d-1} + \cdots + \frac{h_1(s_1, s_2, \ldots, s_h)}{g(s_1, s_2, \ldots, s_h)}x + \frac{h_0(s_1, s_2, \ldots, s_h)}{g(s_1, s_2, \ldots, s_h)}$$

whenever $s_1, s_2, \ldots, s_h \in \mathsf{F}$ such that $g(s_1, s_2, \ldots, s_h) \neq 0$. Consider the polynomials

$$h = gx^d + h_{d-1}x^{d-1} + \cdots + h_1x + h_0 \in \mathsf{F}[x, y_1, y_2, \ldots, y_h]$$

and

$$f = \begin{cases} \mathrm{Res}_x\left(h, \frac{\partial h}{\partial x}\right) \in \mathsf{F}[y_1, y_2, \ldots, y_h] & \text{if } d \neq 0 \text{ in } \mathsf{F}, \\ g \cdot \mathrm{Res}_x\left(h, \frac{\partial h}{\partial x}\right) \in \mathsf{F}[y_1, y_2, \ldots, y_h] & \text{otherwise.} \end{cases}$$

Since $h(\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_h) \in \mathsf{F}[x]$ is the product of $g(\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_h)$ and the minimal polynomial of $\hat{s}_1\gamma_1 + \hat{s}_2\gamma_2 + \cdots + \hat{s}_h\gamma_h$, $h(\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_h)$ is a separable polynomial in $\mathsf{F}[x]$. Therefore $f(\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_h) \neq 0$.

Conversely, let $s_1, s_2, \ldots, s_h \in \mathsf{F}$ such that $f(s_1, s_2, \ldots, s_h)$ is nonzero. Since g divides f (because f is the determinant of a Sylvester matrix of polynomials whose entries in one row are all divisible by g when $d \neq 0$, and by definition otherwise), $g(s_1, s_2, \ldots, s_h)$ is nonzero as well, so the minimal polynomial of $s_1\gamma_1 + s_2\gamma_2 + \cdots + s_h\gamma_h$ has maximal degree d, and $h(s_1, s_2, \ldots, s_h)$ is the product of this minimal polynomial and $g(s_1, s_2, \ldots, s_h)$. Since $f(s_1, s_2, \ldots, s_h) \neq 0$, $h(s_1, s_2, \ldots, s_h)$ is a separable polynomial in $\mathsf{F}[x]$ and the minimal polynomial of $s_1\gamma_1+s_2\gamma_2+\cdots+s_h\gamma_h$ is therefore separable as well. Thus, $s_1\gamma_1+s_2\gamma_2+\cdots+s_h\gamma_h$ is self-centralizing in $\mathfrak{A}$, as desired.

It remains only to bound the total degree of f. The entries of the matrix σ^i each have total degree at most i in $y_1, y_2, \ldots, y_h$, for $0 \leq i \leq d$. Therefore each entry in the i^{th} column of the matrix M shown in equation (3.14) has total degree at most $i-1$ in these indeterminates, and the entries of the vector v have total degree at most d. The determinant g of M, and the polynomials $h_0, h_1, \ldots, h_{d-1}$ obtained by an application of Cramer's rule to this system, therefore each have total degree at most $\binom{d}{2}$ in $y_1, y_2, \ldots, y_h$. Since f is a factor of the determinant of a $(2d-1)\times(2d-1)$ Sylvester matrix whose nonzero entries are scalar multiples of these polynomials*, it follows as required that f has total degree at most

$$(2d-1)\binom{d}{2} = \frac{2d^3+d^2-d}{2} \leq \frac{3d^3}{2}. \quad \square$$

*Indeed, if the characteristic of F does not divide d, so that $\partial h/\partial x$ has degree $d-1$, then f is equal to this determinant. Otherwise this matrix is block triangular and its determinant is the product of f and a nonnegative power of g.

3.3 Certification of Self-Centralizing Elements

Theorem 3.7 yields a simple Monte Carlo algorithm to generate a self-centralizing element: Choose a random linear combination of a set of elements of $\mathfrak{A}$ whose F-linear span is known to include such an element.

In this section we describe a method to either certify that a given element α of $\mathfrak{A}$ is self-centralizing or reject the element, assuming that a basis for $\mathfrak{A}$ over F is available. This method is also randomized and may only fail by rejecting an element that is, indeed, self-centralizing. Another method that is somewhat slower, but guaranteed never to give an incorrect answer, is mentioned at the end of the section.

Once again consider an element α of $\mathfrak{A}$. The minimal polynomial f of α over F is easily computed by generating the Frobenius form of α and, since f is separable if and only if f and f' are relatively prime, one can efficiently detect and reject any element whose minimal polynomial is not separable over F.

If α's minimal polynomial is separable, and the degree bound d is known, then it is easy to complete our procedure — we simply compare the degree of f to d, accepting α if the degree equals d and rejecting α otherwise. We will therefore continue by giving a method that can be used when d is unknown, noting that it is never necessary to use this again after a self-centralizing element of $\mathfrak{A}$ has been found and certified, since d is available after that.

The following (partial) converse of Theorem 3.2 will serve as the basis for our test.

Theorem 3.8. *If α is an element of $\mathfrak{A}$ whose minimal polynomial over F is separable but has degree less than d then $\mathsf{F}[\alpha]$ is a proper subset of $C_{\mathfrak{A}}(\alpha)$.*

Proof. Suppose the centre of $\mathfrak{A}$ is contained in $\mathsf{F}[\alpha]$ (the result is trivial otherwise). As usual, let $\alpha = \alpha_1 + \alpha_2 + \cdots + \alpha_k$, where α_i is a member of the simple component $\mathfrak{A}_i$ of $\mathfrak{A}$, and let f_i be the minimal polynomial of α_i over F for $1 \le i \le k$.

Suppose $f_1, f_2, \ldots, f_k$ are not pairwise relatively prime; then there exist distinct integers i and j between 1 and k such that the greatest common divisor $g_{i,j}$ of f_i and f_j has positive degree. However, since the identity element ω_i of $\mathfrak{A}_i$ is in the centre of $\mathfrak{A}$, and this is contained in $\mathsf{F}[\alpha]$ by assumption, $\omega_i = h(\alpha)$ for some polynomial $h \in \mathsf{F}[x]$. Since $i \neq j$ and $h(\alpha) \in \mathfrak{A}_i$, $h(\alpha_j) = 0$ in $\mathfrak{A}_j$, implying that h is divisible by f_j and therefore by its factor $g_{i,j}$. On the other hand, since $h(\alpha_i) = \omega_i$ in $\mathfrak{A}_i$, $h \equiv 1 \pmod{f_i}$, implying that h is relatively prime with f_i and therefore with its factor $g_{i,j}$. This clearly contradicts the fact that $g_{i,j}$ has positive degree. Thus $f_1, f_2, \ldots, f_k$ are pairwise relatively prime, the minimal polynomial of α over F is their product, and

$$\mathsf{F}[\alpha] = \mathsf{F}[\alpha_1] \oplus \mathsf{F}[\alpha_2] \oplus \cdots \oplus \mathsf{F}[\alpha_k].$$

Since the centre E_i of $\mathfrak{A}_i$ is contained in $\mathsf{F}[\alpha_i]$ (within $\mathfrak{A}_i$), $\mathsf{F}[\alpha_i] = \mathsf{E}_i[\alpha_i]$. Suppose the minimal polynomial of α_i over E_i has degree $\widehat{n}_i$; then this is also the dimension of $\mathsf{E}_i[\alpha_i]$ over E_i. Since E_i is a field extension with degree e_i over F, $\mathsf{F}[\alpha_i]$ clearly has dimension $e_i\widehat{n}_i$ over F, so that the minimal polynomial f_i of α_i over F has degree $e_i\widehat{n}_i$. Since the minimal polynomial of α over F is the product of $f_1, f_2, \ldots, f_k$, this minimal polynomial has degree

$$e_1\widehat{n}_1 + e_2\widehat{n}_2 + \cdots + e_k\widehat{n}_k < d = e_1n_1 + e_2n_2 + \ldots e_kn_k.$$

It follows that $\widehat{n}_i < n_i$ for at least one integer i; fix any such i.

It now remains only to prove that there is an element β_i of $\mathfrak{A}_i$ such that $\alpha_i\beta_i = \beta_i\alpha_i$ but $\beta_i \notin \mathsf{F}[\alpha_i]$. For the remainder of the proof, let us consider $\mathfrak{A}_i$ as a central simple algebra over its centre E_i; it now suffices to show that the dimension of the centralizer of α_i in $\mathfrak{A}_i$ over E_i is strictly greater than $\widehat{n}_i$. We will show that the dimension is greater than or equal to n_i.

Since the dimensions are invariant under extension of scalars, it suffices to show that the dimension of the centralizer of $\alpha_i \otimes_{\mathsf{E}_i} 1$ in $\mathfrak{A}_i^{\mathsf{K}} = \mathfrak{A}_i \otimes_{\mathsf{E}_i} \mathsf{K}$ over K is at least n_i, for some field extension K of E_i. In particular, it is sufficient to prove this when K is an algebraic closure of E_i, so that

$$\mathfrak{A}_i^{\mathsf{K}} = \mathfrak{A}_i \otimes_{\mathsf{E}_i} \mathsf{K} \cong \mathsf{K}^{n_i \times n_i}.$$

Let

$$\psi : \mathfrak{A}_i^{\mathsf{K}} \to \mathsf{K}^{n_i \times n_i}$$

be an isomorphism of algebras over K and consider the matrix $\psi(\alpha_i \otimes_{\mathsf{E}_i} 1) \in \mathsf{K}^{n_i \times n_i}$. The minimal polynomial of this matrix over K is the same as the minimal polynomial of α_i over E_i and, since this is a factor of the minimal polynomial f_i of α_i over F, this polynomial is separable over both E_i and K. Since its degree is strictly less than n_i, the matrix $\psi(\alpha_i \otimes_{\mathsf{E}_i} 1)$ is diagonalizable in $\mathsf{K}^{n_i \times n_i}$ but is similar to a diagonal matrix D_i whose diagonal entries are not distinct. Now

$$D_i = X^{-1}\psi(\alpha_i \otimes_{\mathsf{E}_i} 1)X$$

for some nonsingular matrix $X \in \mathsf{K}^{n_i \times n_i}$. The matrix D_i commutes with all diagonal matrices, so that its centralizer has dimension at least n_i over K. Since a matrix β commutes with D_i if and only $X\beta X^{-1}$ commutes with $\psi(\alpha_i \otimes_{\mathsf{E}_i} 1)$, and ψ is an algebra isomorphism, the dimension of the centralizer of $\alpha_i \otimes_{\mathsf{E}_i} 1$ over K is also at least n_i, and the dimension of the centralizer of α_i over E_i is at least n_i as well. $\square$

If a basis $\gamma_1, \gamma_2, \ldots, \gamma_n$ for $\mathfrak{A}$ over F is available, then we may complete the process of deciding whether α is self-centralizing by checking whether the dimension of the space of solutions of the homogeneous system of linear equations

$$\alpha\left(\sum_{i=1}^{n} x_i\gamma_i\right) - \left(\sum_{i=1}^{n} x_i\gamma_i\right)\alpha = 0,$$

in unknowns $x_1, x_2, \ldots, x_n$, is the same as the degree of the minimal polynomial of α over F. It therefore suffices to consider a system with m^2 equations in n unknowns. However, as suggested in Section 2.2, it may be possible to improve on this by inspecting matrix-vector products instead of the entries of matrices in $\mathfrak{A}$. Consider the algorithm shown in Figure 1 on page 21.

Lemma 3.9. *If $\alpha \in \mathfrak{A}$ is not self-centralizing, and the algorithm in Figure 1 is executed with α and a basis for $\mathfrak{A}$ as input, then the algorithm returns the answer* No.

Input: An element α of a separable algebra $A \subseteq \mathsf{F}^{m\times m}$ whose minimal polynomial f has degree $\widehat{d}$ over the field F, and a basis $\gamma_1, \gamma_2, \ldots, \gamma_n$ for $\mathfrak{A}$ over F

Question: Is α self-centralizing in $\mathfrak{A}$?

1. **if** $\gcd(f, f') = 1$ **then**
2. $\quad i := 0$; $\delta_i := n$
$\quad$ **loop**
3. $\quad\quad i := i + 1$
4. $\quad\quad$ Randomly choose a vector $v_i \in \mathsf{F}^{m\times 1}$
5. $\quad\quad$ Compute the dimension δ_i over F of the space of solutions for the homogeneous system of linear equations

$$\sum_{j=1}^{n} x_j(\alpha\gamma_j v_l - \gamma_j\alpha v_l) = 0 \qquad \text{for } 1 \le l \le i$$

$\quad\quad$ in the n unknowns $x_1, x_2, \ldots, x_n$
6. $\quad\quad$ **if** $\delta_i = \widehat{d}$ **then**
7. $\quad\quad\quad$ **answer** Yes
$\quad\quad$ **else**
8. $\quad\quad\quad$ **if** $(\delta_i = \delta_{i-1}$ **or** $i > m)$ **then**
9. $\quad\quad\quad\quad$ **answer** No
$\quad\quad\quad$ **end if**
$\quad\quad$ **end if**
$\quad$ **end loop**
else
10. $\quad$ **answer** No
end if

Figure 1: Certification of a Self-Centralizing Element

Proof. Since α is not self-centralizing, either its minimal polynomial f is not separable, or it is separable but the degree $\widehat{d}$ of f is less than d. In the former case $\gcd(f, f')$ has positive degree, so the test in step 1 will fail and step 10 will be executed to reject α. In the latter case Theorem 3.8 implies that $\mathsf{F}[\alpha]$ is a proper subset of the centralizer of α in $\mathfrak{A}$. It follows that the dimension of the solution space of the homogeneous system of linear equations considered at line 5 will never be less than $\widehat{d} + 1$, and the test at line 6 will always fail. Therefore the test at line 8 will eventually succeed, either because two dimensions δ_{i-1} and δ_i coincide, or because $m + 1$ vectors have been considered (so that $i = m + 1 > m$). Thus the algorithm will eventually return the answer No (by executing line 9) in this case as well. $\square$

For $i \ge 1$, let R_i be the maximum (over all choices of the vectors $v_1, v_2, \ldots, v_i \in \mathsf{F}^{m\times 1}$) of the rank of the coefficient matrix of the system of linear equations shown at line 5 on the i^{th} execution of the loop body. Clearly $R_i \le R_{i+1}$ for $i \ge 1$. Furthermore, since the vector $[s_1, s_2, \ldots, s_n]^t$ is a solution for this system whenever $s_1\gamma_1 + s_2\gamma_2 + \cdots + s_n\gamma_n$ belongs to the centralizer of $\alpha \in \mathfrak{A}$, $R_i \leq n - \delta$ for all $i \geq 1$ where δ is the dimension of this centralizer.

Let N be as defined in equation (2.4) on page 4.

Lemma 3.10. $R_N = R_{N+1} = n - \delta$.

Proof. Since $R_N \leq R_{N+1} \leq n - \delta$, it suffices to show that $R_N \geq n - \delta$.

Consider the given system when $i = N$ and suppose $v_1, v_2, \ldots, v_N$ is a distinguishing set for $\mathfrak{A}$. In this case, for every element β of $\mathfrak{A}$, $(\beta\alpha - \alpha\beta)v_i$ for $1 \leq i \leq N$ if and only if β commutes with α, so that $[s_1, s_2, \ldots, s_n]^t$ is a solution for the given system if and only if $s_1\gamma_1 + s_2\gamma_2 + \cdots + s_n\gamma_n$ is in the centralizer. Thus the rank of the coefficient matrix of the system is $n - \delta$. This clearly implies that $R_N \geq n - \delta$, as needed. $\square$

Lemma 3.11. *If α is not in the centre of $\mathfrak{A}$ then $R_1 > 0$ and, in general, if $1 \leq i < N$ such that $R_i < n - \delta$ then $R_{i+1} \geq R_i + 1$.*

Proof. Consider the second claim first, suppose to the contrary that $R_i = R_{i+1} < n - \delta$, and let $v_1, v_2, \ldots, v_i$ be vectors such that the system given in line 5 (on the i^{th} execution of the loop body) has rank R_i when these vectors are used. Then, since $R_{i+1} = R_i$, the additional equations obtained by considering any other vector v must be linear combinations of the equations that have already been obtained, implying that $R_i = R_{i+1} = R_{i+2} = \cdots = R_N$, and contradicting Lemma 3.10.

The first claim follows by essentially the same argument, since it can be used to show that if $R_1 = 0$ then $R_i = 0$ as well for all $i \geq 1$, contradicting Lemma 3.10 and the fact that $\delta < n$ when α is not in the centre of $\mathfrak{A}$. $\square$

Now let N_α be the smallest positive integer such that $R_{N_\alpha} = n - \delta$, so that $N_\alpha \leq N$ by Lemma 3.10.

Lemma 3.12. *Let ϵ be a real number such that $0 < \epsilon < 1$, and suppose S is a finite subset of F that includes at least n/ϵ distinct elements. If the algorithm shown in Figure 1 is executed with inputs α and a basis for $\mathfrak{A}$, and the entries of the vectors $v_1, v_2, \ldots$ used by this algorithm are selected uniformly and independently from S, then all three of the following conditions are satisfied with probability at least $1 - \epsilon$.*

- *If α is self-centralizing in $\mathfrak{A}$ then the loop body of the algorithm is executed exactly $\ell = N_\alpha$ times, and the algorithm returns the answer* Yes.
- *If α is not self-centralizing in $\mathfrak{A}$ then the loop body of the algorithm is executed exactly $\ell = 1 + N_\alpha$ times, and the algorithm returns the answer* No.
- *If ℓ is defined as in the above two statements, then the linear system considered on the i^{th} execution of the loop body has rank R_i, for $1 \leq i \leq \ell$.*

Proof. The claim is trivial if α is in the centre of $\mathfrak{A}$, because the coefficient matrix of every system that can be considered has rank zero in this case. If α is also self-centralizing then $\widehat{d} = d = n = \delta_1$, regardless of the choice of v_1, and the test at line 6 will succeed on the first execution of the loop body. If α is not self-centralizing then $\delta_2 = \delta_1 = n = d \neq \widehat{d}$, so that the test at line 8 will succeed on the second execution. All three conditions are satisfied in either case. Suppose, therefore, that α is not in the centre.

Now, $\delta \neq n$ and $R_1 > 0$. Since the centralizer of α in $\mathfrak{A}$ has dimension δ, the set of elements $\beta\alpha - \alpha\beta$ such that $\beta \in \mathfrak{A}$ has dimension $n - \delta$ over F. Let $\beta_1, \beta_2, \ldots, \beta_{n-\delta} \in \mathsf{F}$ such that $\beta_1\alpha - \alpha\beta_1, \beta_2\alpha - \alpha\beta_2, \ldots, \beta_{n-\delta}\alpha - \alpha\beta_{n-\delta}$ are linearly independent and therefore form a basis for this set.

Let $\overline{v}$ be an m-dimensional vector whose entries are distinct indeterminates over F. To prove that the coefficient matrix for the system considered on the first execution of the loop body has rank R_1 with high probability, consider the $m \times (n - \delta)$ matrix of polynomials

$$\begin{bmatrix}(\beta_1\alpha - \alpha\beta_1)\overline{v} & (\beta_2\alpha - \alpha\beta_2)\overline{v} & \ldots & (\beta_{n-\delta}\alpha - \alpha\beta_{n-\delta})\overline{v}\end{bmatrix}.$$

The definition of R_1 implies that there is a vector $v \in \mathsf{F}^{m\times 1}$ such that, if $\overline{v}$ were replaced by v in the above matrix, then the resulting matrix would have rank R_1. This matrix would therefore have a nonsingular $R_1 \times R_1$ submatrix. The corresponding submatrix of the above matrix of polynomials is thus an $R_1 \times R_1$ matrix whose determinant is a nonzero polynomial g_1 with total degree at most R_1 in the entries of $\overline{v}$. Furthermore, it is clear by the definitions of g_1 and R_1 that if $\widehat{v} \in \mathsf{F}^{m\times 1}$ such that $g_1(\widehat{v}) \neq 0$, then the matrix obtained from the above by replacing $\overline{v}$ with $\widehat{v}$ has rank R_1, as does the coefficient matrix of the system obtained on the first execution of the loop body if $\widehat{v}$ is the first vector selected. It follows by an application of the Schwartz-Zippel lemma that if v_1 is randomly selected as described in the claim, then the probability that the first system has rank less than R_1 is at most $\epsilon R_1/n$.

Suppose next that $1 \leq i < N_\alpha$ and that vectors $v_1, v_2, \ldots, v_i$ have been chosen so that the coefficient matrix of the system considered at line 5 on the j^{th} execution of the loop body (involving vectors $v_1, v_2, \ldots, v_j$) has rank R_j for $1 \leq j \leq i$. Now, the tests at lines 6 and 8 will both fail on the i^{th} execution of the loop body since $\delta_0 = n$, $1 \leq R_1 < R_2 < \cdots < R_i < R_N = \delta$, and $\delta_j = n - R_j$ for $1 \leq j \leq i$. An $i + 1^{\text{st}}$ execution will therefore be performed. Let $\overline{v}$ be a vector of indeterminates as before, and consider the matrix

$$\begin{bmatrix}(\beta_1\alpha - \alpha\beta_1)v_1 & (\beta_2\alpha - \alpha\beta_2)v_1 & \ldots & (\beta_{n-\delta}\alpha - \alpha\beta_{n-\delta})v_1 \\ (\beta_1\alpha - \alpha\beta_1)v_2 & (\beta_2\alpha - \alpha\beta_2)v_2 & \ldots & (\beta_{n-\delta}\alpha - \alpha\beta_{n-\delta})v_2 \\ \vdots & \vdots & \ddots & \vdots \\ (\beta_1\alpha - \alpha\beta_1)v_i & (\beta_2\alpha - \alpha\beta_2)v_i & \ldots & (\beta_{n-\delta}\alpha - \alpha\beta_{n-\delta})v_i \\ (\beta_1\alpha - \alpha\beta_1)\overline{v} & (\beta_2\alpha - \alpha\beta_2)\overline{v} & \ldots & (\beta_{n-\delta}\alpha - \alpha\beta_{n-\delta})\overline{v}\end{bmatrix}.$$

The submatrix including all columns and the top mi rows has rank R_i by the choice of $v_1, v_2, \ldots, v_i$, and it follows by the definition of R_{i+1} that there exists a vector v_{i+1} such that the matrix obtained from the above by replacing $\overline{v}$ with v_{i+1} has rank R_{i+1}. This matrix would have a nonsingular $R_{i+1} \times R_{i+1}$ submatrix such that the top R_i rows of this submatrix are selected from the top mi rows of the entire matrix. A consideration of the corresponding submatrix of the above matrix of polynomials and another application of the Schwartz-Zippel lemma establish that if a matrix $\widehat{v}_{i+1}$ is randomly selected as described in the claim, and $\widehat{v}_{i+1}$ replaces $\overline{v}$, then the resulting matrix has rank less than R_{i+1} with probability at most $\epsilon(R_{i+1} - R_i)/n$. This also bounds the probability that the system generated on the $i + 1^{\text{st}}$ execution of the loop body has rank less than R_{i+1} if the system obtained on the i^{th} execution had full rank R_i.

It follows by induction on i that if $1 \leq i \leq N_\alpha$ and $v_1, v_2, \ldots, v_i$ are chosen as described then the probability that the j^{th} coefficient matrix has rank R_j for all j between 1 and i is at least $1 - \epsilon(R_i/n)$. In particular, the system obtained after N_α executions of the loop body has maximal rank $R_{N_\alpha} = n - \delta$ with probability at least $1 - \epsilon(R_{N_\alpha}/n) \geq 1 - \epsilon$. Suppose for the remainder of the argument that this system does have maximal rank.

Now, if α is self-centralizing then the algorithm will terminate on the N_α^{th} execution of the loop body, returning the answer Yes, because the test at line 6 will succeed. If α is not self-centralizing, then it will terminate on the $N_\alpha + 1^{\text{st}}$ execution of the loop body instead, returning the answer No, because the ranks of the last two systems considered must be the same, but must also be less than $n - \delta$.

Therefore all three conditions are satisfied with probability at least $1 - \epsilon$, as claimed. $\square$

A final lemma concerns the cost of implementing this algorithm.

Lemma 3.13. *The algorithm shown in Figure 1 can be implemented in such a way that each execution of the loop body can be performed using $O(nm^2 + \frac{n^2}{m^2}\mathcal{MM}(m))$ operations, or $O(nm^2 + n^2 m)$ operations if standard arithmetic is used.*

Proof. Consider the i^{th} execution of the loop body. If $i = 1$ this requires that a homogeneous system of m equations in n unknowns $x_1, x_2, \ldots, x_n$ be formed and examined, while if $i > 1$ then it involves the addition of another m equations in these unknowns to a system that has been constructed in previous executions of the loop body. The loop body can be implemented to have the above complexity, provided that information about the previous system is maintained and used.

Suppose, in particular, that the coefficient matrix of this system has rank r. It will be assumed that r linearly independent rows of the coefficient matrix, the indices of r linearly independent columns specifying a nonsingular $r \times r$ submatrix X, and the inverse of this submatrix are maintained.

Since $r = 0$ before the first execution of the loop body, this information can be initialized in constant time before this first execution begins.

The beginning of the i^{th} execution of the loop body involves the incrementing of a variable and the selection of a vector v_i from $\mathsf{F}^{m \times 1}$, and this can clearly be performed at the stated cost. The equations to be added to the system at this point have the form

$$\sum_{j=1}^{n} x_j(\alpha\gamma_j v_i - \gamma_j \alpha v_i) = 0,$$

where $\gamma_1, \gamma_2, \ldots, \gamma_n$ is a basis for $\mathfrak{A}$, and these can be formed using at most $4n$ multiplications of $m \times m$ matrices (in $\mathfrak{A}$) by the vector v_i, at cost $O(nm^2)$.

Now it remains only to compute the rank δ_i of the current system and to generate the data that will be needed for the next execution of the loop — for, once δ_i is known (and δ_{i-1} is recalled), the remaining steps of the loop body can be executed using a constant number of operations.

Suppose $m \geq n$; then the new equations can be split into $\lceil m/n \rceil$ sets of at most n equations each and added to the previous system in $\lceil m/n \rceil$ stages, one set at a time. Since each intermediate system has rank at most n, the system will include at most $2n$ equations in n unknowns at each stage. Therefore the process of computing the rank of each intermediate system, and selecting and inverting a nonsingular submatrix of maximal size, can be implemented using $O(\mathcal{MM}(n))$ operations. Since $O(m/n)$ stages are required, the entire process can be completed using $O(\frac{m}{n}\mathcal{MM}(n)) = O(mn^2)$ operations.

Suppose instead that $m < n$. In this case, one should begin if $i > 1$ by eliminating the entries in the new rows of the current system's coefficient matrix that lie in the columns that were used to form the nonsingular matrix X currently in use. Since X^{-1} is available, this elimination can be performed using $O(\frac{m}{n}\mathcal{MM}(n))$ operations. The resulting m equations can be inspected to determine which new equations should be added to the set that will be used in the next execution of the loop, as well as the rows and columns that should be added to the nonsingular submatrix X.

Suppose that a matrix

$$\widehat{X} = \begin{bmatrix} X & C \\ R & Y \end{bmatrix}$$

has been selected. Since X is a nonsingular $\delta_{i-1} \times \delta_{i-1}$ matrix and $\widehat{X}$ is a nonsingular $\delta_i \times \delta_i$ matrix, $C \in \mathsf{F}^{\delta_{i-1}\times(\delta_i-\delta_{i-1})}$, $R \in \mathsf{F}^{(\delta_i-\delta_{i-1})\times\delta_{i-1}}$, and $Y \in \mathsf{F}^{(\delta_i-\delta_{i-1})\times(\delta_i-\delta_{i-1})}$. Furthermore $\delta_{i-1} \le n$ and $\delta_i - \delta_{i-1} \le m$, because the entire system has rank at most n and the new system has been obtained by adding only m new equations to the previous one. It is well known (and easily verified) that

$$\widehat{X} = \begin{bmatrix} I_{\delta_{i-1}} & 0 \\ RX^{-1} & I_{(\delta_i-\delta_{i-1})} \end{bmatrix} \begin{bmatrix} X & 0 \\ 0 & S \end{bmatrix} \begin{bmatrix} I_{\delta_{i-1}} & X^{-1}C \\ 0 & I_{(\delta_i-\delta_{i-1})} \end{bmatrix}$$

and

$$\begin{aligned} \widehat{X}^{-1} &= \begin{bmatrix} I_{\delta_{i-1}} & -X^{-1}C \\ 0 & I_{(\delta_i-\delta_{i-1})} \end{bmatrix} \begin{bmatrix} X^{-1} & 0 \\ 0 & S^{-1} \end{bmatrix} \begin{bmatrix} I_{\delta_{i-1}} & 0 \\ -RX^{-1} & I_{(\delta_i-\delta_{i-1})} \end{bmatrix} \\ &= \begin{bmatrix} X^{-1} + X^{-1}CS^{-1}RX^{-1} & -X^{-1}CS^{-1} \\ -S^{-1}RX^{-1} & S^{-1} \end{bmatrix} \end{aligned}$$

for $S = Y - RX^{-1}C \in \mathsf{F}^{(\delta_i-\delta_{i-1})\times(\delta_i-\delta_{i-1})}$. A careful scheduling of operations will permit $\widehat{X}^{-1}$ to be computed from X' and X^{-1} using $O(\frac{n^2}{m^2}\mathcal{MM}(m)) = O(n^2m)$ operations, as required. $\square$

Now we can bound the cost to certify a self-centralizing element.

Theorem 3.14. *Suppose as usual that $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ is a separable algebra with dimension n over a field F, and let $\gamma_1, \gamma_2, \dots, \gamma_n$ be a basis for $\mathfrak{A}$ over F. Suppose as well that ϵ is a real number such that $0 < \epsilon < 1$ and that S is a finite subset of F including at least n/ϵ distinct elements.*

Let $\alpha \in \mathfrak{A}$, and suppose that the algorithm shown in Figure 1 is executed on inputs α and $\gamma_1, \gamma_2, \dots, \gamma_n$, in such a way that the entries of the vectors $v_1, v_2, \dots$ used by this algorithm are chosen uniformly and independently from S. Then each of the following conditions is satisfied.

- *The algorithm always terminates and returns either* Yes *or* No *as output, after performing* $O(nm^3 + \frac{n^2}{m}\mathcal{MM}(m))$ *operations, or* $O(nm^3 + n^2m^2)$ *operations using standard arithmetic.*

- *If* α *is not self-centralizing in* $\mathfrak{A}$ *then the algorithm's output is always* No.

- *If* α *is self-centralizing in* $\mathfrak{A}$ *then the algorithm's output is* Yes *with probability at least* $1 - \epsilon$.

- *The algorithm will terminate after* $O(Nnm^2 + N\frac{n^2}{m^2}\mathcal{MM}(m))$ *operations, or* $O(Nnm^2 + Nn^2m)$ *operations using standard arithmetic, with probability at least* $1 - \epsilon$.

Proof. It is clear by inspection of the algorithm that, if it terminates at all, then it does so by returning either Yes or No (but not both). Furthermore, since the parameter i is incremented on each execution of the loop, a glance at line 8 will confirm that the loop is never executed more than $m + 1$ times. This, and Lemma 3.13, are sufficient to establish the first claim — for the cost of executing the loop clearly dominates the cost of executing the other steps.

The second claim is a consequence of Lemma 3.9, and the third is a consequence of Lemma 3.12.

Finally, the last claim follows from Lemma 3.12, which implies that with high probability the loop body will be executed at most $N_\alpha + 1 \leq N + 1$ times, and Lemma 3.13, which bounds the cost of each execution of this loop. □

As noted above, the algorithm may return No with small probability when its input α is self-centralizing in $\mathfrak{A}$. This behaviour can be eliminated by checking the system of equations

$$\sum_{i=1}^{n}(x_i\gamma_i\alpha - x_i\alpha\gamma_i) = 0$$

at any point in the loop body when the original algorithm would return No; if the dimension of the solution space for this system equals the degree of the minimal polynomial of α then (since it has already been confirmed that this minimal polynomial is separable), the algorithm should return the answer Yes. On the other hand, No should be returned if the dimension and degree are different.

With this change, the worst case complexity of the algorithm will clearly increase, since a system with m^2 equations and n unknowns may be considered. However, Theorems 3.2 and 3.8 imply that this more expensive algorithm will always return a correct output, as desired.

4 Centering Pairs and Their Properties

4.1 Definitions

It turns out that certain pairs of self-centralizing elements are more useful in combination than any one such element.

Definition 4.1. A pair of elements α and β of $\mathfrak{A}$ is a *centering pair* if α and β are both self-centralizing in $\mathfrak{A}$ and

$$\mathrm{Centre}(\mathfrak{A}) = C_{\mathfrak{A}}(\alpha) \cap C_{\mathfrak{A}}(\beta) = \mathsf{F}[\alpha] \cap \mathsf{F}[\beta]. \tag{4.1}$$

Having a centering pair α and β for $\mathfrak{A}$ is clearly of great advantage in computing the centre of $\mathfrak{A}$, since a basis for the centre over F could be obtained by solving the homogeneous system of linear equations

$$(y_0 + y_1\alpha + \cdots + y_{d-1}\alpha^{d-1})\beta - \beta(y_0 + y_1\alpha + \cdots + y_{d-1}\alpha^{d-1}) = 0$$

for the unknowns $y_0, y_1, \ldots, y_{d-1}$ in F: Every solution $[s_0, s_1, \ldots, s_{d-1}]^t \in \mathsf{F}^d$ determines an element

$$s_0 + s_1\alpha + \cdots + s_{d-1}\alpha^{d-1}$$

of $\mathsf{F}[\alpha]$ that commutes with β. Since β is self-centralizing in $\mathfrak{A}$, this implies that the above element belongs to $\mathsf{F}[\beta]$ as well. It therefore belongs to $\mathsf{F}[\alpha] \cap \mathsf{F}[\beta]$ which is the centre of $\mathfrak{A}$ by definition. Conversely, every element of the centre belongs to the set $\{s_0 + s_1\alpha + \cdots + s^{d-1}\alpha^{d-1} : s_0, s_1, \ldots, s_{d-1} \in \mathsf{F}\}$ and specifies a solution for this system.

While it is plausible that this method is faster than previous general methods for computation of the centre, it requires that we form and solve a system of m^2 linear equations in d unknowns. We can do considerably better than this by projecting from the space of matrices to the space of vectors. It will be shown in the sequel that with high probability the desired relationships still hold, and this motivates the following definition.

Definition 4.2. A pair α and β of elements of a separable matrix algebra $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ is a *complemented centering pair* for $\mathfrak{A}$ if this pair is a centering pair for $\mathfrak{A}$ and, furthermore, there exists a pair of vectors u and v in $\mathsf{F}^{m\times 1}$ such that

$$(\mu u = \nu u \text{ and } \mu v = \nu v) \implies \mu = \nu \in \mathsf{F}[\alpha] \cap \mathsf{F}[\beta] \tag{4.2}$$

for all $\mu \in \mathsf{F}[\alpha]$ and all $\nu \in \mathsf{F}[\beta]$. Any pair of vectors u and v satisfying condition (4.2), above, is said to *complement* the centering pair α and β.

4.2 Existence and Density of Centering Pairs

Once again let d be as given in equation (2.5).

Theorem 4.3. *Let $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ be a separable matrix algebra over a field F. If F is infinite then $\mathfrak{A}$ includes a complemented centering pair of elements α and β.*

Theorem 4.4, below, will be used to prove Theorem 4.3 and will therefore be proved first.

Theorem 4.4. *Let $\mathfrak{A}$ be as above, and suppose $\gamma_1, \gamma_2, \ldots, \gamma_h \in \mathfrak{A}$ such that there is a complemented centering pair α and β in the F-linear span of $\gamma_1, \gamma_2, \ldots, \gamma_h$. Let ϵ be a real number such that $0 < \epsilon < 1$ and suppose S is a finite subset of F that includes at least $5d^3/\epsilon$ distinct elements. Then, if elements $a_1, a_2, \ldots, a_h, b_1, b_2, \ldots, b_h, c_1, c_2, \ldots, c_m, d_1, d_2, \ldots, d_m$ are chosen uniformly and independently from S, then the elements $a_1\gamma_1 + a_2\gamma_2 + \cdots + a_h\gamma_h$ and $b_1\gamma_1 + b_2\gamma_2 + \cdots + b_h\gamma_h$ form a complemented centering pair in $\mathfrak{A}$, complemented by the vectors $[c_1, c_2, \ldots, c_m]^t$ and $[d_1, d_2, \ldots, d_m]^t$ in $\mathsf{F}^{m\times 1}$, with probability at least $1 - \epsilon$.*

Proof of Theorem 4.4. Let $s_1, s_2, \ldots, s_h$, $t_1, t_2, \ldots, t_h$, $u_1, u_2, \ldots, u_m$, and $v_1, v_2, \ldots, v_m$ be indeterminates over the field F. It is given that there exist elements $\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_h, \hat{t}_1, \hat{t}_2, \ldots, \hat{t}_h \in \mathsf{F}$ such that the elements

$$\hat{s} = \hat{s}_1\gamma_1 + \hat{s}_2\gamma_2 + \ldots \hat{s}_h\gamma_h \quad \text{and} \quad \hat{t} = \hat{t}_1\gamma_1 + \hat{t}_2\gamma_2 + \ldots \hat{t}_h\gamma_h$$

of $\mathfrak{A}$ form a complemented centering pair. Consider matrices of polynomials

$$\sigma = \gamma_1 s_1 + \gamma_2 s_2 + \cdots + \gamma_h s_h \quad \text{and} \quad \tau = \gamma_1 t_1 + \gamma_2 t_2 + \cdots + \gamma_h t_h,$$

so that $\hat{s} = \sigma(\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_h)$ and $\hat{t} = \tau(\hat{t}_1, \hat{t}_2, \ldots, \hat{t}_h)$. Clearly

$$\sigma(r_1, r_2, \ldots, r_h) = \tau(r_1, r_2, \ldots, r_h) = r_1\gamma_1 + r_2\gamma_2 + \cdots + r_h\gamma_h \in \mathfrak{A}$$

for all $r_1, r_2, \ldots, r_h \in \mathsf{F}$.

It can be established as in the proof of Theorem 3.7 that there exist nonzero polynomials $f_\alpha \in \mathsf{F}[s_1, s_2, \ldots, s_h]$ and $f_\beta \in \mathsf{F}[t_1, t_2, \ldots, t_h]$ (formed using $\hat{s}$ and $\hat{t}$ respectively) such that $f_\alpha(\hat{s}_1, \hat{s}_2, \ldots, \hat{s}_h) \neq 0$, $f_\beta(\hat{t}_1, \hat{t}_2, \ldots, \hat{t}_h) \neq 0$, each polynomial has total degree at most $\frac{2d^3+d^2-d}{2}$ in its indeterminates, and such that for all $r_1, r_2, \ldots, r_h \in \mathsf{F}$, if either $f_\alpha(r_1, r_2, \ldots, r_h)$ or $f_\beta(r_1, r_2, \ldots, r_h)$ is nonzero then $r_1\gamma_1 + r_2\gamma_2 + \cdots + r_h\gamma_h$ is self-centralizing in $\mathfrak{A}$.

Since $\hat{s}$ and $\hat{t}$ form a complemented centering pair, there also exist vectors

$$\hat{u} = \begin{bmatrix} \hat{u}_1 \\ \hat{u}_2 \\ \vdots \\ \hat{u}_m \end{bmatrix} \in \mathsf{F}^m \quad \text{and} \quad \hat{v} = \begin{bmatrix} \hat{v}_1 \\ \hat{v}_2 \\ \vdots \\ \hat{v}_m \end{bmatrix} \in \mathsf{F}^m$$

that complement $\hat{s}$ and $\hat{t}$. Thus there exists a homogeneous system of $2m$ linear equations

$$\begin{aligned} (x_0 1 + x_1\hat{s} + x_2\hat{s}^2 + \cdots + x_{d-1}\hat{s}^{d-1})\hat{u} - (y_0 1 + y_1\hat{t} + y_2\hat{t}^2 + \cdots + y_{d-1}\hat{t}^{d-1})\hat{u} = 0, \\ (x_0 1 + x_1\hat{s} + x_2\hat{s}^2 + \cdots + x_{d-1}\hat{s}^{d-1})\hat{v} - (y_0 1 + y_1\hat{t} + y_2\hat{t}^2 + \cdots + y_{d-1}\hat{t}^{d-1})\hat{v} = 0 \end{aligned} \tag{4.3}$$

in $2d$ indeterminates $x_0, x_1, \ldots, x_{d-1}, y_0, y_1, \ldots, y_{d-1}$, such that

$$a_0 1 + a_1\hat{s} + \cdots + a_{d-1}\hat{s}^{d-1} = b_0 1 + b_1\hat{t} + \cdots + b_{d-1}\hat{t}^{d-1}$$

for each solution $[a_0, a_1, \ldots, a_{d-1}, b_0, b_1, \ldots, b_{d-1}]^t \in \mathsf{F}^{2d}$ of this system, with the above element $a_0 1 + a_1\hat{s} + \cdots + a_{d-1}\hat{s}^{d-1}$ of $\mathfrak{A}$ in the centre of $\mathfrak{A}$. Conversely, every element of the centre is equal to both $a_0 1 + a_1\hat{s} + \cdots + a_{d-1}\hat{s}^{d-1}$ and $b_0 1 + b_1\hat{t} + \cdots + b_{d-1}\hat{t}^{d-1}$ for some solution $[a_0, a_1, \ldots, a_{d-1}, b_0, b_1, \ldots, b_{d-1}]^t$. Writing

$$\vec{x} = \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{d-1} \end{bmatrix} \quad \text{and} \quad \vec{y} = \begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_{d-1} \end{bmatrix},$$

the system of linear equations shown in (4.3), above, can be expressed as

$$\hat{A}\begin{bmatrix}\vec{x}\\\vec{y}\end{bmatrix}=0$$

where $\hat{A}\in\mathsf{F}^{2m\times 2d}$. Since the space of solutions of this system has the same dimension e as the centre of $\mathfrak{A}$, $\hat{A}$ has rank $2d-e$ and has a nonsingular $(2d-e)\times(2d-e)$ submatrix $\hat{B}$.

Now set

$$\chi=\begin{bmatrix}u_1\\u_2\\\vdots\\u_m\end{bmatrix}\in\mathsf{F}[u_1,u_2,\dots,u_m]^{m\times 1}\quad\text{and}\quad\psi=\begin{bmatrix}v_1\\v_2\\\vdots\\v_m\end{bmatrix}\in\mathsf{F}[v_1,v_2,\dots,v_m]^{m\times 1},$$

so that $\hat{u}=\chi(\hat{u}_1,\hat{u}_2,\dots,\hat{u}_m)$ and $\hat{v}=\psi(\hat{v}_1,\hat{v}_2,\dots,\hat{v}_m)$, and consider the system of equations

$$\begin{aligned}(x_01+x_1\sigma+x_2\sigma^2+\cdots+x_{d-1}\sigma^{d-1})\chi-(y_01+y_1\tau+y_2\tau^2+\cdots+y_{d-1}\tau^{d-1})\chi=0,\\(x_01+x_1\sigma+x_2\sigma^2+\cdots+x_{d-1}\sigma^{d-1})\psi-(y_01+y_1\tau+y_2\tau^2+\cdots+y_{d-1}\tau^{d-1})\psi=0;\end{aligned}\tag{4.4}$$

this can be written as

$$A\begin{bmatrix}\vec{x}\\\vec{y}\end{bmatrix}$$

where $A\in\mathsf{F}[s_1,s_2,\dots,s_h,t_1,t_2,\dots,t_h,u_1,u_2,\dots,u_m,v_1,v_2,\dots,v_m]^{2m\times 2d}$ such that

$$A(\hat{s}_1,\hat{s}_2,\dots,\hat{s}_h,\hat{t}_1,\hat{t}_2,\dots,\hat{t}_h,\hat{u}_1,\hat{u}_2,\dots,\hat{u}_m,\hat{v}_1,\hat{v}_2,\dots,\hat{v}_m)=\hat{A}\in\mathsf{F}^{2m\times 2d}.$$

Choosing the same rows and columns as were used to define $\hat{B}$ from $\hat{A}$, one can also define a matrix $B\in\mathsf{F}[s_1,s_2,\dots,s_h,t_1,t_2,\dots,t_h,u_1,u_2,\dots,u_m,v_1,v_2,\dots,v_m]^{(2d-e)\times(2d-e)}$ such that

$$B(\hat{s}_1,\hat{s}_2,\dots,\hat{s}_h,\hat{t}_1,\hat{t}_2,\dots,\hat{t}_h,\hat{u}_1,\hat{u}_2,\dots,\hat{u}_m,\hat{v}_1,\hat{v}_2,\dots,\hat{v}_m)=\hat{B}\in\mathsf{F}^{d\times d}.$$

Consider now the polynomial

$$f=f_\alpha f_\beta\det B\in\mathsf{F}[s_1,s_2,\dots,s_h,t_1,t_2,\dots,t_h,u_1,u_2,\dots,u_m,v_1,v_2,\dots,v_m].$$

By construction, $f(\hat{s}_1,\hat{s}_2,\dots,\hat{s}_h,\hat{t}_1,\hat{t}_2,\dots,\hat{t}_h,\hat{u}_1,\hat{u}_2,\dots,\hat{u}_m,\hat{v}_1,\hat{v}_2,\dots,\hat{v}_m)\neq 0$, so this polynomial is nonzero. On the other hand, if $\overline{s}_1,\overline{s}_2,\dots,\overline{s}_h$, $\overline{t}_1,\overline{t}_2,\dots,\overline{t}_h$, $\overline{u}_1,\overline{u}_2,\dots,\overline{u}_m$, and $\overline{v}_1,\overline{v}_2,\dots,\overline{v}_m$ are elements of F such that

$$f(\overline{s}_1,\overline{s}_2,\dots,\overline{s}_h,\overline{t}_1,\overline{t}_2,\dots,\overline{t}_h,\overline{u}_1,\overline{u}_2,\dots,\overline{u}_m,\overline{v}_1,\overline{v}_2,\dots,\overline{v}_m)\neq 0,\tag{4.5}$$

then clearly $f_\alpha(\overline{s}_1,\overline{s}_2,\dots,\overline{s}_h)$ and $f_\beta(\overline{t}_1,\overline{t}_2,\dots,\overline{t}_h)$ are both nonzero, so that the elements

$$\overline{\alpha}=\overline{s}_1\gamma_1+\overline{s}_2\gamma_2+\cdots+\overline{s}_h\gamma_h\quad\text{and}\quad\overline{\beta}=\overline{t}_1\gamma_1+\overline{t}_2\gamma_2+\cdots+\overline{t}_h\gamma_h$$

of $\mathfrak{A}$ are both self-centralizing. Furthermore, the determinant of the matrix

$$B(\overline{s}_1,\overline{s}_2,\dots,\overline{s}_h,\overline{t}_1,\overline{t}_2,\dots,\overline{t}_h,\overline{u}_1,\overline{u}_2,\dots,\overline{u}_m,\overline{v}_1,\overline{v}_2,\dots,\overline{v}_m)$$

is nonzero. If we set

$$\overline{u} = \begin{bmatrix} \overline{u}_1 \\ \overline{u}_2 \\ \vdots \\ \overline{u}_m \end{bmatrix} \quad \text{and} \quad \overline{v} = \begin{bmatrix} \overline{v}_1 \\ \overline{v}_2 \\ \vdots \\ \overline{v}_m \end{bmatrix},$$

then this implies that the coefficient matrix of the homogeneous system of $2m$ linear equations

$$\begin{bmatrix} \overline{u} & \overline{\alpha}\,\overline{u} & \dots & \overline{\alpha}^{d-1}\overline{u} & -\overline{u} & -\overline{\beta}\overline{u} & \dots & -\overline{\beta}^{d-1}\overline{u} \\ \overline{v} & \overline{\alpha}\,\overline{v} & \dots & \overline{\alpha}^{d-1}\overline{v} & -\overline{v} & -\overline{\beta}\overline{v} & \dots & -\overline{\beta}^{d-1}\overline{v} \end{bmatrix} \begin{bmatrix} x_0 \\ \vdots \\ x_{d-1} \\ y_0 \\ \vdots \\ y_{d-1} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

in $2d$ unknowns $x_0, x_1, \dots, x_{d-1}, y_0, y_1, \dots, y_{d-1}$ has (maximal) rank $2d - e$, and that the space of solutions for this system has dimension e over F.

It now follows that $\overline{\alpha}$ and $\overline{\beta}$ form a centering pair in $\mathfrak{A}$: Since $\overline{\alpha}$ and $\overline{\beta}$ are both self-centralizing, the centre of $\mathfrak{A}$ is contained in $\mathsf{F}[\overline{\alpha}] \cap \mathsf{F}[\overline{\beta}]$, and is only a proper subset of this vector space if the dimension of $\mathsf{F}[\overline{\alpha}] \cap \mathsf{F}[\overline{\beta}]$ exceeds e. However, for every element

$$a_0 1 + a_1\overline{\alpha} + \cdots + a_{d-1}\overline{\alpha}^{d-1} = b_0 1 + b_1\overline{\beta} + \cdots + b_{d-1}\beta^{d-1}$$

of $\mathsf{F}[\overline{\alpha}] \cap \mathsf{F}[\overline{\beta}]$ there is a (distinct) solution $[a_0, \dots, a_{d-1}, b_0, \dots, b_{d-1}]^t$ for the above system, so the fact that the space of solutions for the system has dimension e implies that $\mathsf{F}[\overline{\alpha}] \cap \mathsf{F}[\overline{\beta}]$ also has dimension at most e. Thus $\mathsf{F}[\overline{\alpha}] \cap \mathsf{F}[\overline{\beta}] = \mathrm{Centre}(\mathfrak{A})$ as needed.

The fact that the solution space for the system has dimension e also implies that, for all $\mu \in \mathsf{F}[\overline{\alpha}]$ and $\nu \in \mathsf{F}[\overline{\beta}]$,

$$(\mu\overline{u} = \nu\overline{u} \text{ and } \mu\overline{v} = \nu\overline{v}) \implies \mu = \nu \in \mathsf{F}[\overline{\alpha}] \cap \mathsf{F}[\overline{\beta}],$$

for, otherwise, the dimension of the solution space would exceed that of $\mathsf{F}[\overline{\alpha}] \cap \mathsf{F}[\overline{\beta}]$. Thus the vectors $\overline{u}$ and $\overline{v}$ complement the centering pair $\overline{\alpha}$ and $\overline{\beta}$.

It remains only to bound the degree of the above polynomial f and to apply the Schwartz-Zippel lemma (Schwartz 1980, Zippel 1979) in order to establish the result. An inspection of the above system confirms that each entry of the matrix A, and its submatrix B, has total degree at most d in the indeterminates $s_0, \dots, s_h$, $t_0, \dots, t_h$, $u_1, \dots, u_m$, and $v_1, \dots, v_m$. Since B is a matrix with order $2d - e < d$, its determinant is a polynomial with total degree at most $(2d - e)d < 2d^2$ in these indeterminates. Since $f = f_\alpha f_\beta \det B$, the degree bounds given above for f_α and f_β imply that f has total degree less than $5d^3$, as required. $\square$

It remains for us to prove Theorem 4.3.

Lemma 4.5. *Let $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ be a separable algebra with simple components $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$ over F, and let $\omega_1, \omega_2, \ldots, \omega_k$ be the central primitive idempotents of $\mathfrak{A}$ and the identity elements of algebras $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$ respectively. Suppose*

$$\alpha = \alpha_1 + \alpha_2 + \cdots + \alpha_k \quad \text{and} \quad \beta = \beta_1 + \beta_2 + \cdots + \beta_k$$

where as usual $\alpha_i, \beta_i \in \mathfrak{A}_i$ for all i, and suppose α and β are both self-centralizing in $\mathfrak{A}$.

Consider α_i and β_i as elements of $\mathfrak{A}_i$ (so that $\mathsf{F}[\alpha_i]$ has spanning set $\omega_i, \alpha_i, \alpha_i^2, \ldots$ and $\mathsf{F}[\beta_i]$ is spanned by $\omega_i, \beta_i, \beta_i^2, \ldots$). If

$$\mathsf{F}[\alpha_i] \cap \mathsf{F}[\beta_i] = \mathrm{Centre}(\mathfrak{A}_i)$$

for all i, so that α_i and β_i form a centering pair in $\mathfrak{A}_i$ for all i, then α and β form a centering pair in $\mathfrak{A}$.

Furthermore, if for all i there exist vectors $\vec{u}_i$ and $\vec{v}_i$ such that, for all $\mu_i \in \mathsf{F}[\alpha_i] \subseteq \mathfrak{A}_i$ and for all $\nu_i \in \mathsf{F}[\beta_i] \subseteq \mathfrak{A}_i$,

$$(\mu_i\omega_i\vec{u}_i = \nu_i\omega_i\vec{u}_i \quad \text{and} \quad \mu_i\omega_i\vec{v}_i = \nu_i\omega_i\vec{v}_i) \implies \mu_i = \nu_i \in \mathsf{F}[\alpha_i] \cap \mathsf{F}[\beta_i],$$

then α and β form a complemented centering pair that is complemented by the vectors

$$u = \omega_1\vec{u}_1 + \omega_2\vec{u}_2 + \cdots + \omega_k\vec{u}_k \quad \text{and} \quad v = \omega_1\vec{v}_1 + \omega_2\vec{v}_2 + \cdots + \omega_k\vec{v}_k.$$

Proof. Since α and β are self-centralizing,

$$\mathsf{F}[\alpha] = \mathsf{F}[\alpha_1] \oplus \mathsf{F}[\alpha_2] \oplus \cdots \oplus \mathsf{F}[\alpha_k] \quad \text{and} \quad \mathsf{F}[\beta] = \mathsf{F}[\beta_1] \oplus \mathsf{F}[\beta_2] \oplus \cdots \oplus \mathsf{F}[\beta_k].$$

Since $\mathfrak{A}$ has simple components $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$,

$$\mathrm{Centre}(\mathfrak{A}) = \mathrm{Centre}(\mathfrak{A}_1) \oplus \mathrm{Centre}(\mathfrak{A}_2) \oplus \cdots \oplus \mathrm{Centre}(\mathfrak{A}_k)$$

as well. It follows immediately that, if $\mathsf{F}[\alpha_i] \cap \mathsf{F}[\beta_i] = \mathrm{Centre}(\mathfrak{A}_i)$ in $\mathfrak{A}_i$ for all i, then (in $\mathfrak{A}$)

$$\begin{aligned} \mathsf{F}[\alpha] \cap \mathsf{F}[\beta] &= (\mathsf{F}[\alpha_1] \oplus \mathsf{F}[\alpha_2] \oplus \cdots \oplus \mathsf{F}[\alpha_k]) \cap (\mathsf{F}[\beta_1] \oplus \mathsf{F}[\beta_2] \oplus \cdots \oplus \mathsf{F}[\beta_k]) \\ &= (\mathsf{F}[\alpha_1] \cap \mathsf{F}[\beta_1]) \oplus (\mathsf{F}[\alpha_2] \oplus \mathsf{F}[\beta_2]) \oplus \cdots \oplus (\mathsf{F}[\alpha_k] \oplus \mathsf{F}[\beta_k]) \\ &= \mathrm{Centre}(\mathfrak{A}_1) \oplus \mathrm{Centre}(\mathfrak{A}_2) \oplus \cdots \oplus \mathrm{Centre}(\mathfrak{A}_k) = \mathrm{Centre}(\mathfrak{A}), \end{aligned}$$

establishing the first part of the claim.

Suppose next that there exist vectors $\vec{u}_i$ and $\vec{v}_i$ for all i with the stated property, and let u and v be as above. Suppose as well that $\mu \in \mathsf{F}[\alpha]$ and $\nu \in \mathsf{F}[\beta]$, and write

$$\mu = \mu_1 + \mu_2 + \cdots + \mu_k \quad \text{and} \quad \nu = \nu_1 + \nu_2 + \cdots + \nu_k$$

where as usual $\mu_i, \nu_i \in \mathfrak{A}_i$ for all i. If $\mu u = \nu u$ and $\mu v = \nu v$ then $\omega_i\mu u = \omega_i\nu u$ and $\omega_i\mu v = \omega_i\nu v$ for all i and, since $\omega_i\mu_j = \omega_i\nu_j = 0$ whenever $i \neq j$, this implies that $\omega_i\mu_i\omega_i\vec{u}_i = \omega_i\nu_i\omega_i\vec{u}_i$ and $\omega_i\mu_i\omega_i\vec{v}_i = \omega_i\nu_i\omega_i\vec{v}_i$. Now, since ω_i is central in $\mathfrak{A}$ and is an idempotent, it follows that $\mu_i\omega_i\vec{u}_i = \nu_i\omega_i\vec{u}_i$ and $\mu_i\omega_i\vec{v}_i = \nu_i\omega_i\vec{v}_i$, so that $\mu_i = \nu_i \in \mathsf{F}[\alpha_i] \cap \mathsf{F}[\beta_i] = \mathrm{Centre}(\mathfrak{A}_i)$ in $\mathfrak{A}_i$ for each i. Therefore

$$\mu = \nu \in \mathrm{Centre}(\mathfrak{A}_1) \oplus \mathrm{Centre}(\mathfrak{A}_2) \oplus \cdots \oplus \mathrm{Centre}(\mathfrak{A}_k) = \mathrm{Centre}(\mathfrak{A}) = \mathsf{F}[\alpha] \cap \mathsf{F}[\beta],$$

as required. $\square$

Proof of Theorem 4.3. Suppose first that $\mathfrak{A}$ is simple and isomorphic to $\mathsf{F}^{n_1\times n_1}$ over F. Then there exist distinct elements $\lambda_1, \lambda_2, \dots, \lambda_{n_1}$ of F and an element α of $\mathfrak{A}$ whose minimal polynomial is

$$f = \prod_{i=1}^{n_1}(x-\lambda_i) \in \mathsf{F}[x].$$

Furthermore any simple $\mathfrak{A}$-module contains elements $x_1, x_2, \dots, x_{n_1}$ such that $\alpha x_i = \lambda_i x_i$ for $1 \le i \le n_1$. These elements are linearly independent since they are eigenvectors corresponding to distinct eigenvalues of α and, since any simple $\mathfrak{A}$-module has dimension n_1 over F, they form a basis for the module containing them. The action of α on the module with respect to the basis $x_1, x_2, \dots, x_{n_1}$ is clearly given by the matrix

$$\phi(\alpha) = \begin{bmatrix} \lambda_1 & & & 0 \\ & \lambda_2 & & \\ & & \ddots & \\ 0 & & & \lambda_{n_1} \end{bmatrix} \in \mathsf{F}^{n_1\times n_1}.$$

Suppose $f = x^{n_1} + f_{n_1-1}x^{n_1-1} + \cdots + f_1 x + f_0$, so that $1, f_{n_1-1}, \dots, f_1, f_0$ are the coefficients of f. Since $\mathfrak{A} \cong \mathsf{F}^{n_1\times n_1}$, there exists an element β of $\mathfrak{A}$ whose action on the module with respect to the basis $x_1, x_2, \dots, x_{n_1}$ is given by the companion matrix of f:

$$\phi(\beta) = C_f = \begin{bmatrix} 0 & & & & & -f_0 \\ 1 & 0 & & & & -f_1 \\ & 1 & & & & -f_2 \\ & & \ddots & & & \vdots \\ & & & 1 & 0 & -f_{n-2} \\ 0 & & & & 1 & -f_{n-1} \end{bmatrix}.$$

In this case, $\beta x_i = x_{i+1}$ for $1 \le i \le n-1$, so that if $0 \le j \le n-1$ then $\beta^j x_1 = x_{j+1}$. Now let $u = x_1$ and $v = x_1 + x_2 + \cdots + x_{n_1}$, and suppose $f_1, f_2 \in \mathsf{F}[x]$ such that $f_1(\alpha)u = f_2(\beta)u$ and $f_1(\alpha)v = f_2(\beta)v$. It suffices to consider the case that f_1 and f_2 both have degree less than n_1, since $f_1(\alpha) = \widehat{f_1}(\alpha)$ and $f_2(\beta) = \widehat{f_2}(\beta)$ for $\widehat{f_1} \equiv f_1 \bmod f$ and $\widehat{f_2} \equiv f_2 \bmod f$. Therefore, let

$$f_1 = f_{1,n_1-1}x^{n_1-1} + f_{1,n_1-2}x^{n_1-2} + \cdots + f_{1,1}x + f_{1,0}$$

and let

$$f_2 = f_{2,n_1-1}x^{n_1-1} + f_{2,n_1-2}x^{n_1-2} + \cdots + f_{2,1}x + f_{2,0}.$$

Since $u = x_1$ is an eigenvector of α for eigenvalue λ_1, $f_1(\alpha)u = f_1(\lambda_1)x_1$. On the other hand, it follows by the above equations that

$$f_2(\beta)u = \sum_{i=0}^{n_1-1} f_{2,i}\beta^i x_1 = \sum_{i=0}^{n_1-1} f_{2,i}x_{i+1}.$$

Since $f_1(\alpha)u = f_2(\beta)u$ and $x_1, x_2, \ldots, x_{n_1}$ are linearly independent over F, this implies that $f_1(\lambda_1) = f_{2,0}$ and that $f_{2,i} = 0$ for $1 \le i \le n_1 - 1$, so $f_2(x) = f_1(\lambda_1) \in \mathsf{F}$ and $f_2(\beta) = f_1(\lambda_1)I_{n_1}$ is in the centre of $\mathfrak{A}$.

On the other hand, since $v = x_1 + x_2 + \cdots + x_{n_1}$,

$$f_1(\alpha)v = f_1(\lambda_1)x_1 + f_1(\lambda_2)x_2 + \cdots + f_1(\lambda_{n_1})x_{n_1},$$

by the choice of $x_1, x_2, \ldots, x_{n_1}$, while

$$f_2(\beta)v = f_1(\lambda_1)I_n v = f_1(\lambda_1)x_1 + f_1(\lambda_1)x_2 + \cdots + f_1(\lambda_1)x_{n_1}.$$

The linear independence of $x_1, x_2, \ldots, x_{n_1}$ and the condition that $f_1(\alpha)v = f_2(\beta)v$ imply (by a comparison of the coefficients of $x_1, x_2, \ldots, x_{n_1}$ in the above expressions) that

$$f_1(\lambda_1) = f_1(\lambda_2) = \cdots = f_1(\lambda_{n_1}).$$

Since f_1 has degree less than n_1 and $\lambda_1, \lambda_2, \ldots, \lambda_{n_1}$ are distinct, it follows that $f_{1,0} = f_1(\lambda_1)$ and $f_{1,i} = 0$ for $1 \le i \le n_1 - 1$ as well, so that $f_1(x) = f_1(\lambda_1) = f_2(x)$, and

$$f_1(\alpha) = f_1(\lambda_1)I_n = f_2(\beta)$$

with both in the centre of $\mathfrak{A}$. Thus α and β form a complemented centering pair that is complemented by the vectors u and v in this case.

Lemma 4.5 can now be applied to establish the result for the case that $\mathfrak{A}$ is separable over an infinite field F, such that each simple component is isomorphic to a full matrix ring over F. In particular, this can be used to prove the result for the case that $\mathfrak{A}$ is separable over F and F is algebraically closed.

It remains to consider the case that $\mathfrak{A}$ is separable over an arbitrary infinite field F. Let K be an algebraic closure of F and consider the algebra $\mathfrak{A}^{\mathsf{K}} = \mathfrak{A} \otimes_{\mathsf{F}} \mathsf{K}$ obtained from $\mathfrak{A}$ by extension of scalars. Let $\gamma_1, \gamma_2, \ldots, \gamma_n$ be a basis for $\mathfrak{A}$ over F, so that the elements $\gamma_1 \otimes_{\mathsf{F}} 1, \gamma_2 \otimes_{\mathsf{F}} 1, \ldots, \gamma_n \otimes_{\mathsf{F}} 1$ form a basis for $\mathfrak{A}^{\mathsf{K}}$ over K. Let S be a finite subset of F with size at least $10d^3$; since F is infinite some such set exists.

Now, suppose elements $s_1, s_2, \ldots, s_n$, $t_1, t_2, \ldots, t_n$, $u_1, u_2, \ldots, u_m$, and $v_1, v_2, \ldots, v_m$ are chosen uniformly and independently from S. Let

$$\alpha = s_1\gamma_1 + s_2\gamma_2 + \cdots + s_n\gamma_n \quad \text{and} \quad \beta = t_1\gamma_1 + t_2\gamma_2 + \cdots + t_n\gamma_n,$$

and note that

$$\alpha \otimes_{\mathsf{F}} 1 = s_1(\gamma_1 \otimes_{\mathsf{F}} 1) + s_2(\gamma_2 \otimes_{\mathsf{F}} 1) + \cdots + s_n(\gamma_n \otimes_{\mathsf{F}} 1)$$

and

$$\beta \otimes_{\mathsf{F}} 1 = t_1(\gamma_1 \otimes_{\mathsf{F}} 1) + t_2(\gamma_2 \otimes_{\mathsf{F}} 1) + \cdots + t_n(\gamma_n \otimes_{\mathsf{F}} 1)$$

as well.

Since $\mathfrak{A}$ is separable over F and F is infinite, Theorem 3.6 implies that $\mathfrak{A}$ contains a self-centralizing element, so there must be such an element in the F-linear span of the basis

$\gamma_1, \gamma_2, \dots, \gamma_n$. It therefore follows by Theorem 3.7 that the probability that α (respectively, β) is *not* self-centralizing in $\mathfrak{A}$ is at most $\frac{3}{20}$, so that the probability that α and β are not both self-centralizing in $\mathfrak{A}$ is at most $\frac{3}{10}$.

It follows by the argument given above that $\mathfrak{A}^{\mathsf{K}}$ has a complemented centering pair. Theorem 4.4 therefore implies that $\alpha \otimes_{\mathsf{F}} 1$ and $\beta \otimes_{\mathsf{F}} 1$ form a complemented centering pair of $\mathfrak{A}^{\mathsf{K}}$, that is complemented by the vectors

$$\vec{u} \otimes_{\mathsf{F}} 1 = \begin{bmatrix} u_1 \otimes_{\mathsf{F}} 1 \\ u_2 \otimes_{\mathsf{F}} 1 \\ \vdots \\ u_m \otimes_{\mathsf{F}} 1 \end{bmatrix} \quad \text{and} \quad \vec{v} \otimes_{\mathsf{F}} 1 = \begin{bmatrix} v_1 \otimes_{\mathsf{F}} 1 \\ v_2 \otimes_{\mathsf{F}} 1 \\ \vdots \\ v_m \otimes_{\mathsf{F}} 1 \end{bmatrix} \in \mathsf{K}^{m \times 1},$$

with probability at least $\frac{1}{2}$. Now, if $\vec{u} = [u_1, u_2, \dots, u_m]^t$ and $\vec{v} = [v_1, v_2, \dots, v_m]^t$ in $\mathsf{F}^{m \times 1}$, then $\mathsf{F}[\alpha] \cap \mathsf{F}[\beta] = \text{Centre}(\mathfrak{A})$ and furthermore that, for all $f, g \in \mathsf{F}[x]$,

$$(f(\alpha)\vec{u} = g(\beta)\vec{u} \ \text{ and } f(\alpha)\vec{v} = g(\beta)\vec{v}) \implies (f(\alpha) = g(\beta) \in \mathsf{F}[\alpha] \cap \mathsf{F}[\beta])$$

with probability at least $\frac{1}{2}$ as well. Thus the probability that α and β do *not* form a complemented centering pair, complemented by the vectors $\vec{u}$ and $\vec{v}$, is at most $\frac{3}{10} + \frac{1}{2} = \frac{4}{5} < 1$.

Since a complemented centering pair can be randomly chosen with positive probability, a complemented centering pair must clearly exist. $\square$

4.3 A Monte Carlo Algorithm for a Complemented Centering Pair and Generator for the Centre

A randomized (Monte Carlo) algorithm to compute a complemented centering pair α and β, vectors u and v that complement this pair, and a generator γ for the centre of a separable algebra $\mathfrak{A}$ over F is shown in Figure 2 on page 35. Its analysis yields the following result.

Theorem 4.6. *Let $\mathfrak{A} \subseteq \mathsf{F}^{m \times m}$ be a separable algebra with dimension n over an infinite field F, let ϵ be a real number such that $0 < \epsilon < 1$, and suppose S is a finite subset of F that includes at least $8m^3/\epsilon$ distinct elements. Then a randomized (Monte Carlo) algorithm can be used to compute elements α, β and γ of $\mathfrak{A}$ and vectors $u, v \in \mathsf{F}^{m \times 1}$ such that α and β form a complemented centering pair for $\mathfrak{A}$ complemented by the vectors u and v, and γ generates the centre of $\mathfrak{A}$, with probability at least $1 - \epsilon$, using $O(\mathcal{MM}(m) \log m + \mathcal{R}(\mathfrak{A}))$ operations, or $O(m^3 + \mathcal{R}(\mathfrak{A}))$ operations if standard arithmetic is used. Here $\mathcal{R}(\mathfrak{A})$ is the cost to compute an S-linear combination of a set of elements of $\mathfrak{A}$ whose F-linear span includes a complemented centering pair.*

Recall that Theorem 4.3 implies that a complemented centering pair exists. Thus if a basis for $\mathfrak{A}$ is available we can set $\mathcal{R}(\mathfrak{A}) = nm^2$.

Proof of Theorem 4.6. Consider the algorithm shown in Figure 2. Theorem 4.3 implies that a complemented centering pair for $\mathfrak{A}$ exists. Theorem 4.4 implies that the elements α and β chosen in step 1 form a complemented centering pair for $\mathfrak{A}$, complemented by the vectors u and v chosen in step 3, with probability at least $1 - \frac{5\epsilon}{8}$, when α and β are chosen as S-linear

Input: A separable matrix algebra $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ over an infinite field F

Output: Elements α, β and γ of $\mathfrak{A}$, vectors u and v in $\mathsf{F}^{m\times 1}$, and a positive integer e such that α and β form a complemented centering pair for $\mathfrak{A}$ complemented by the vectors u and v, γ is a generator for the centre of $\mathfrak{A}$, and e is the dimension of the centre with high probability.

1. Randomly choose elements α and β from $\mathfrak{A}$.
2. Compute the degree d of the minimal polynomial of α over F.
3. Randomly choose vectors $u, v \in \mathsf{F}^{m\times 1}$.
4. Compute the dimension e and a basis

$$\begin{bmatrix} a_{1,0} \\ \vdots \\ a_{1,d-1} \\ b_{1,0} \\ \vdots \\ b_{1,d-1} \end{bmatrix}, \begin{bmatrix} a_{2,0} \\ \vdots \\ a_{2,d-1} \\ b_{2,0} \\ \vdots \\ b_{2,d-1} \end{bmatrix}, \ldots, \begin{bmatrix} a_{e,0} \\ \vdots \\ a_{e,d-1} \\ b_{e,0} \\ \vdots \\ b_{e,d-1} \end{bmatrix} \in \mathsf{F}^{2d\times 1}$$

for the set of solutions of the homogeneous system of linear equations

$$\begin{bmatrix} u & \alpha u & \ldots & \alpha^{d-1}u & -u & -\beta u & \ldots & -\beta^{d-1}u \\ v & \alpha v & \ldots & \alpha^{d-1}v & -v & -\beta v & \ldots & -\beta^{d-1}v \end{bmatrix} \begin{bmatrix} y_0 \\ \vdots \\ y_{d-1} \\ z_0 \\ \vdots \\ z_{d-1} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

in the indeterminates $y_0, \ldots, y_{d-1}, z_0, \ldots, z_{d-1}$.
5. Randomly choose elements $c_1, c_2, \ldots, c_e$ from (a finite subset of) F.
6. Set $s_i = \sum_{j=1}^{e} c_j a_{j,i}$ for $0 \le i \le d-1$ and set $\gamma = \sum_{i=0}^{d-1} s_i \alpha^i$.
7. Return the above elements α, β and γ of $\mathfrak{A}$, vectors u and v, and integer e.

Figure 2: A Monte Carlo Algorithm for a Centering Pair and the Centre

combinations of elements of $\mathfrak{A}$ as described above and the entries of the vectors u and v are chosen uniformly and independently from S. Thus the probability of failure to find a complemented centering pair and complementing vectors is at most $5\epsilon/8$. The cost of steps 1 and 3 is clearly at most $O(\mathcal{R}(\mathfrak{A}) + m)$.

The degree d of the minimal polynomial of α is readily available if the Frobenius form of α can be computed. It therefore follows by Lemma 2.4 that step 2 of the algorithm can be performed using $O(\mathcal{MM}(m) \log m)$ operations in F, or $O(m^3)$ operations using standard arithmetic, by a Las Vegas algorithm that fails with probability at most $\epsilon/8m \le \epsilon/8$.

Now consider the homogeneous system of linear equations that is formed and solved in step 4. The cost of forming this system is dominated by the cost of computing matrix-vector products $v, \alpha v, \alpha^2 v, \ldots, \alpha^{d-1}v$ for a given element α of $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ and a given vector $v \in \mathsf{F}^{m\times 1}$, and thus the system can be formed using $O(\mathcal{M}(m)\log m)$ operations (see, for example, Keller-Gehrig 1985), or at cost $O(m^3)$ using standard arithmetic by forming fewer than m matrix-vector products. The system includes $2m$ equations in $2d$ unknowns and, since $m \geq d$, this system can be solved using $O(\mathcal{MM}(m))$ operations. It follows by the definition of a complemented centering pair that if α and β form such a pair that is complemented by the vectors u and v, and if the set of vectors

$$\begin{bmatrix} a_{1,0} \\ \vdots \\ a_{1,d-1} \\ b_{1,0} \\ \vdots \\ b_{1,d-1} \end{bmatrix}, \begin{bmatrix} a_{2,0} \\ \vdots \\ a_{2,d-1} \\ b_{2,0} \\ \vdots \\ b_{2,d-1} \end{bmatrix}, \ldots, \begin{bmatrix} a_{e,0} \\ \vdots \\ a_{e,d-1} \\ b_{e,0} \\ \vdots \\ b_{e,d-1} \end{bmatrix}$$

is a basis for the set of solutions for this system (as in step 4), then the set

$$\sum_{j=0}^{d-1} a_{1,j}\alpha^j, \sum_{j=0}^{d-1} a_{2,j}\alpha^j, \ldots, \sum_{j=0}^{d-1} a_{e,j}\alpha^j$$

of elements of $\mathfrak{A}$ forms a basis for the centre of $\mathfrak{A}$ over F. In this case, the element γ that is generated in step 6 is a random linear combination of the elements of such a basis, so that $\gamma \in \mathrm{Centre}(\mathfrak{A})$ and, furthermore, γ is a self-centralizing element in $\mathrm{Centre}(\mathfrak{A})$ with probability at least $1 - 3\epsilon/16 > 1 - \epsilon/4$. That is, the probability that γ is not self-centralizing in the centre is less than $\epsilon/4$. Now, since any self-centralizing element of a commutative algebra is a generator for the algebra, this implies that the probability that $\mathsf{F}[\gamma] \neq \mathrm{Centre}(\mathfrak{A})$ is at most $\epsilon/4$, if steps 1–4 of the algorithm succeeded.

Finally, note that $\gamma = g(\alpha)$ where $g(x) = s_{d-1}x^{d-1} + s_{d-2}x^{d-2} + \cdots + s_0$ and where the coefficients $s_{d-1}, s_{d-2}, \ldots, s_0$ are as computed in step 6 of the algorithm. These coefficients can be computed from the values generated in earlier steps using $O(ed) = O(m^2)$ operations. Since a Frobenius form and transition matrix for α have been computed in earlier steps, γ can be computed by evaluating the polynomial g at the matrix α deterministically using $O(\mathcal{MM}(m)\log m)$ steps, or $O(m^3)$ operations using standard arithmetic, if the earlier steps succeeded (see Section 6 of Giesbrecht 1995).

Thus the entire algorithm can be implemented at the cost that has been claimed, and the probability of failure is at most $5\epsilon/8 + \epsilon/8 + \epsilon/4 = \epsilon$, as required. □

4.4 A Las Vegas Algorithm for a Complemented Centering Pair and Generator for the Centre

A Las Vegas algorithm to compute these values is shown in Figure 3 on page 37. In this case, both a basis and a set of generators for the algebra $\mathfrak{A}$ are specified as input. Of course, one could use the elements of the basis as the generators and execute the algorithm using

Input: A basis $\gamma_1, \gamma_2, \ldots, \gamma_n$ for a separable matrix algebra $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ over an infinite field F, and a set of generators $\zeta_1, \zeta_2, \ldots, \zeta_s$ for $\mathfrak{A}$ over F

Output: Either

- Elements α, β, and γ of $\mathfrak{A}$ and vectors u and v in $\mathsf{F}^{m\times 1}$ such that α and β form a complemented centering pair for $\mathfrak{A}$ complemented by the vectors u and v, and such that the centre of $\mathfrak{A}$ is $\mathsf{F}[\gamma]$

or

- failure

1. Apply the algorithm shown in Figure 2, choosing elements of $\mathfrak{A}$ by forming linear combinations of $\gamma_1, \gamma_2, \ldots, \gamma_n$, to generate α, β, γ, u, v, and an estimate e for the dimension of the centre of $\mathfrak{A}$.
2. Apply the algorithm shown in Figure 1 on inputs α and $\gamma_1, \gamma_2, \ldots, \gamma_n$ to try to certify α as self-centralizing in $\mathfrak{A}$.
3. Return α, β, γ, u and v as output if all five of the following conditions are satisfied; return failure otherwise.
   (a) The executions of algorithms in steps 1 and 2 completed successfully (that is, no application of a Las Vegas algorithm failed).
   (b) The execution of the algorithm in step 2 resulted in the answer Yes.
   (c) The minimal polynomial of β is separable over F and has the same degree as the minimal polynomial of α over F.
   (d) The minimal polynomial of γ is separable with degree e over F.
   (e) $\zeta_i\gamma = \gamma\zeta_i$ for $1 \le i \le s$.

Figure 3: A Las Vegas Algorithm for a Centering Pair and the Centre

the basis alone as input. However, the complexity of the algorithm improves substantially if a smaller set of generators is supplied. The analysis of the algorithm yields the following result.

Theorem 4.7. *Let $\mathfrak{A} \subseteq \mathsf{F}^{m\times m}$ be a separable algebra with dimension n over an infinite field F. Let ϵ be a real number such that $0 < \epsilon < 1$, and suppose that S is a finite subset of F with size at least $10m^3/\epsilon$. Then a complemented centering pair for $\mathfrak{A}$, complementing vectors, and a single generator γ of the centre of $\mathfrak{A}$ can be computed from a basis and a set of s generators for $\mathfrak{A}$, by a Las Vegas algorithm that samples the algebra $\mathfrak{A}$ by computing S-linear combinations of the given basis, and that either returns the desired values or with probability at most ϵ reports failure.*

This computation can be performed using $O(nm^3 + \frac{n^2}{m}\mathcal{MM}(m))$ operations, or $O(nm^3 + n^2m^2)$ operations using standard arithmetic, in the worst case. However, with probability at least $1-\epsilon$, the number of operations used is $O(Nnm^2 + (N\frac{n^2}{m^2} + s + \log m)\mathcal{MM}(m))$, or $O(N(nm^2 + n^2m) + sm^3)$ using standard arithmetic.

Proof. Consider the above algorithm, and suppose that all five of the conditions listed in step 3 are satisfied, so that values α, β, and γ of $\mathfrak{A}$ and vectors u and v are returned.

Since conditions 3(a) and 3(b) are satisfied, it follows by Theorem 3.14 that α is self-centralizing in $\mathfrak{A}$.

Since condition 3(c) is satisfied, β is self-centralizing in $\mathfrak{A}$ as well, so that the centre of $\mathfrak{A}$ is contained in $\mathsf{F}[\alpha] \cap \mathsf{F}[\beta]$.

Condition 3(d) implies that $\mathsf{F}[\alpha] \cap \mathsf{F}[\beta] \subseteq \mathsf{F}[\gamma]$, so that $\mathsf{F}[\gamma]$ includes the centre of the algebra.

Finally, condition 3(e) confirms that γ is in the centre, so that $\mathsf{F}[\gamma] = \text{Centre}(\mathfrak{A})$. Since the vectors u and v were used with α and β to compute γ in step 1, this confirms that α and β form a complemented centering pair complemented by the vectors u and v. Thus, either the algorithm reports $\mathsf{failure}$ or its outputs are correct.

Since condition 3(e) can be checked deterministically using $O(s)$ matrix multiplications, the error probability and complexity results stated in the claim are consequences of Lemma 2.4 and Theorems 3.14 and 4.6, which can be used to bound the failure probability and complexity of each of the remaining steps. $\square$

5 Wedderburn Decomposition of Separable Algebras

Suppose once again that $\mathfrak{A} \subseteq \mathsf{F}^{m \times m}$ is a separable algebra over F, with simple components $\mathfrak{A}_1, \mathfrak{A}_2, \ldots, \mathfrak{A}_k$, and that γ is a generator for the centre of $\mathfrak{A}$. Then γ is a "splitting element" for the algebra $\mathfrak{A}$, as defined by Eberly (1991), and the simple components of $\mathfrak{A}$ can be generated from γ in polynomial time if a factorization of the minimal polynomial of γ in $\mathsf{F}[x]$ is available. Indeed, the algorithm for the Wedderburn decomposition of semi-simple algebras over large perfect fields in Section 3 of Eberly (1991) can also be applied to separable algebras over arbitrary large fields, since the centre of the algebra is a direct sum of simple extensions of F in this case. Using this process one can obtain bases for each of the simple components.

A rather different data structure to identify the simple components of a matrix algebra is discussed by Eberly & Giesbrecht (1999). In particular a *semi-simple transition matrix* is considered, that is, a matrix $X \in \mathsf{F}^{m \times m}$ whose columns include the elements of bases for $\mathfrak{A}_1\mathsf{F}^{m \times 1}, \mathfrak{A}_2\mathsf{F}^{m \times 1}, \ldots, \mathfrak{A}_k\mathsf{F}^{m \times 1}$, and a *semi-simple transition*, which includes this matrix and the dimensions of the above subspaces $\mathfrak{A}_1\mathsf{F}^{m \times 1}, \mathfrak{A}_2\mathsf{F}^{m \times 1}, \ldots, \mathfrak{A}_k\mathsf{F}^{m \times 1}$ of $\mathsf{F}^{m \times 1}$ (see Definition 3.1 of Eberly & Giesbrecht 1999). This can be computed quite efficiently if γ and a factorization of the minimal polynomial of γ are available.

Theorem 5.1. *Suppose ϵ is a real number such that $0 < \epsilon < 1$ and that F is a field including at least $2m^2/\epsilon$ distinct elements. Given a generator γ for the centre of a separable algebra $\mathfrak{A} \subseteq \mathsf{F}^{m \times m}$ and a factorization of the minimal polynomial of γ in $\mathsf{F}[x]$, a semi-simple transition for $\mathfrak{A}$ can be computed using a Las Vegas algorithm that fails with probability less than ϵ, using $O(\mathcal{MM}(m) \log m)$ operations, or $O(m^3)$ operations using standard arithmetic.*

Proof. By Lemma 2.4, a Frobenius decomposition for γ can be generated at the above cost using a Las Vegas algorithm that fails with probability at most $\epsilon/2$. The characteristic polynomial of γ can computed from the Frobenius form of this matrix using $O(m\mathcal{M}(m))$ operations, and since the factorization of the minimal polynomial of γ is available, a factorization of the characteristic polynomial of γ can be computed using a divide and conquer strategy with $O(m\mathcal{M}(m)) \subseteq O(\mathcal{MM}(m))$ operations as well.

Now, since γ generates the centre of $\mathfrak{A}$, $\gamma = \gamma_1 + \gamma_2 + \cdots + \gamma_k$ where $\gamma_i \in \mathfrak{A}_i$ for $1 \le i \le k$ and where the minimal polynomials of $\gamma_1, \gamma_2, \ldots, \gamma_k$ are each irreducible in $\mathsf{F}[x]$ and are pairwise relatively prime. Thus, these are the irreducible factors of the minimal polynomials of γ, and γ is similar to a matrix

$$\widehat{\gamma} = \begin{bmatrix} \widehat{\gamma}_1 & & & 0 \\ & \widehat{\gamma}_2 & & \\ & & \ddots & \\ 0 & & & \widehat{\gamma}_k \end{bmatrix},$$

where $\widehat{\gamma}_i$ is a block diagonal matrix whose diagonal blocks are copies of the companion matrix of the minimal polynomial of γ_i. The order of the matrix $\widehat{\gamma}_i$ can be deduced from the factorization of the characteristic polynomial of γ.

A Frobenius decomposition of $\widehat{\gamma}$ can now be computed by a Las Vegas algorithm failing with probability $\epsilon/2$. At this point, matrices X_1 and X_2 are known such that $X_1\gamma X_1^{-1}$ and $X_2\widehat{\gamma}X_2^{-1}$ are both equal to the common Frobenius form of γ and $\widehat{\gamma}$, and it is easily confirmed that $X_2^{-1}X_1$ is a semi-simple transition matrix for $\mathfrak{A}$, and that the orders of the matrices $\widehat{\gamma}_1, \widehat{\gamma}_2, \ldots, \widehat{\gamma}_k$ are the dimensions of $\mathfrak{A}_1\mathsf{F}^{m\times 1}, \mathfrak{A}_2\mathsf{F}^{m\times 1}, \ldots, \mathfrak{A}_k\mathsf{F}^{m\times 1}$ as needed. $\square$

Of course, the factorization of the minimal polynomial of γ is required above, and the cost to factor this polynomial may dominate the cost of the other operations. However, a self-centralizing element may help to reduce the cost of this factorization as well.

Suppose in particular that g_i is the minimal polynomial of γ_i for $1 \le i \le k$, for $\gamma_1, \gamma_2, \ldots \gamma_k$ as above, so that the minimal polynomial g of γ is the product of $g_1, g_2, \ldots, g_k$. Let $n_1, n_2, \ldots, n_k$ and $m_1, m_2, \ldots, m_k$ be as defined in Section 2, and let

$$\widehat{g}_{i,j} = \prod_{\substack{n_h = i\\ m_h = j}} g_h \tag{5.1}$$

for $i, j \ge 1$. Clearly,

$$g = \prod_{i,j\ge 1} \widehat{g}_{i,j}.$$

Theorem 5.2. *Let ϵ be a real number such that $0 < \epsilon < 1$ and suppose F is a field including at most $4m^2/\epsilon$ distinct elements. If $\mathfrak{A}$, α, γ, and g are as above, then the above factors $\widehat{g_{i,j}}$ of g of positive degree can be computed by a Las Vegas algorithm that fails with probability at most ϵ, using $(\mathcal{MM}(m)\log m)$ operations over F, or using $O(m^3)$ operations using standard arithmetic.*

Proof. Let X be a distinct power transition matrix for the self-centralizing element α. Then, as noted in Section 3.1,

$$X^{-1}\gamma X = \begin{bmatrix} \gamma^{(1)} & & & 0 \\ & \gamma^{(2)} & & \\ & & \ddots & \\ 0 & & & \gamma^{(l)} \end{bmatrix}$$

for matrices $\gamma^{(1)}, \gamma^{(2)}, \ldots, \gamma^{(l)}$ — for, otherwise, the idempotents $\tau_1, \tau_2, \ldots, \tau_l$ considered in Theorem 3.5 would not be central in $\mathfrak{A}$. Furthermore, a comparison of the above equation with equation (3.9) confirms that $\gamma^{(j)}$ has minimal polynomial

$$\prod_{m_h=j} g_h,$$

order $j\delta_j$ (for δ_j as defined in Section 3.1), and characteristic polynomial

$$\prod_{m_h=j} g_h^{jn_h} = \prod_{i\geq 1} \widehat{g}_{i,j}^{ij}.$$

Since the polynomials $\widehat{g}_{i,j}$ are separable and pairwise relatively prime, it is clear that the distinct power divisors of $\gamma^{(j)}$ with positive degree are exactly the polynomials $\widehat{g}_{i,j}$ with positive degree.

These polynomials can therefore be obtained by computing a distinct power decomposition for α, applying the distinct power transition matrix X to γ to generate the matrices $\gamma^{(1)}, \gamma^{(2)}, \ldots, \gamma^{(l)}$, and then computing the distinct power decompositions of each of these matrices. Since the sum of the orders of the matrices $\gamma^{(1)}, \gamma^{(2)}, \ldots, \gamma^{(l)}$ is m and F contains at least $4m^2/\epsilon$ elements, the complexity and failure bounds in the claim now follow from Theorem 2.5. $\square$

References

J. R. Bastida. *Field Extensions and Galois Theory*, vol. 22 of *Encyclopedia of Mathematics and its Applications.* Addison-Wesley, 1984.

A. M. Cohen, G. Ivanyos, and D. B. Wales. Finding the radical of an algebra of linear transformations. *J. Pure Appl. Algebra* **117&118**, pp. 177–193, 1997.

D. Coppersmith and S. Winograd. On the asymptotic complexity of matrix multiplication. *SIAM J. Comput.* **11**, pp. 472–492, 1982.

C. W. Curtis and I. Reiner. *Representation Theory of Finite Groups and Associative Algebras.* New York Interscience (New York), 1962.

W. Eberly. Decompositions of algebras over finite fields and number fields. *Computational Complexity* **1**, pp. 179–206, 1991.

W. Eberly and M. Giesbrecht. Efficient decomposition of associative algebras over finite fields. *J. Symb. Comp.*, to appear, 1999.

K. Friedl and L. Rónyai. Polynomial time solutions of some problems in computational algebra. In *Proceedings, 7th ACM Symposium on Theory of Computing*, pp. 153–162, Providence, RI, USA, 1985.

M. Giesbrecht. Nearly optimal algorithms for canonical matrix forms. *SIAM J. Comput.* **24**, pp. 948–969, 1995.

D. F. Holt and S. Rees. Testing modules for irreducibility. *J. Austral. Math. Soc.* **57**, pp. 1–16, 1994.

G. Ivanyos. Finding the radical of matrix algebras using Fitting decompositions. *J. Pure Appl. Algebra* **139**, pp. 159–182, 1999.

G. Ivanyos and L. Rónyai. Finding maximal orders in semisimple algebras over $\mathbb{Q}$. *Computational Complexity* **3**, pp. 245–261, 1993.

E. Kaltofen, M. S. Krishnamoorthy, and B. D. Saunders. Parallel algorithms for matrix normal forms. *Linear Algebra and its Applications* **136**, pp. 189–208, 1990.

W. Keller-Gehrig. Fast algorithms for the characteristic polynomial. *Theoret. Comput. Sci.* **36**, pp. 309–317, 1985.

R. A. Parker. The computer calculation of modular characters (the meat-axe). In *Computational Group Theory: Proceedings of the London Mathematical Society Symposium on Computational Group Theory*, pp. 267–274, London, 1984. Academic Press.

R. Pierce. *Associative Algebras.* Springer-Verlag (Heidelberg), 1982.

L. Rónyai. Simple algebras are difficult. In *Proceedings, 19th ACM Symposium on Theory of Computing*, pp. 398–408, New York, 1987.

L. Rónyai. Computing the structure of finite algebras. *J. Symbolic Comput.* **9**, pp. 355–373, 1990.

L. Rónyai. Algorithmic properties of maximal orders in simple algebras over $\mathbb{Q}$. *Computational Complexity* **2**, pp. 225–243, 1992.

G. J. A. Schneider. Computing with endomorphism rings of modular representations. *J. Symbolic Comput.* **9**, pp. 607–636, 1990.

A. Schönhage. Schnelle Multiplikation von Polynomen über Körpern der Charakteristik 2. *Acta Informatica* **7**, pp. 395–398, 1977.

A. Schönhage and V. Strassen. Schnelle Multiplikation großer Zahlen. *Computing* **7**, pp. 281–292, 1971.

J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. Assoc. Comput. Mach.* **27**, pp. 701–717, 1980.

J. H. M. Wedderburn. On hypercomplex numbers. *Proc. London Math. Soc.* **6**(2), pp. 77–118, 1907.

R. Zippel. Probabilistic algorithms for sparse polynomials. In *Proc. EUROSAM 79*, pp. 216–226, Marseille, 1979.