November, 1976 ESL-R-700

EFFICIENT MULTIPLICATION IN SEMISIMPLE ALGEBRAS

by

Michael Conrad Loui

B.S., Yale University
(1975)

SUBMITTED IN PARTIAL FULFILLMENT OF THE

REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

November 1976

Signature of Author......Michael C. Loui......
Department of Electrical Engineering and Computer Science,
24 November 1976

Certified by......Alan S. Willsky......
Thesis Supervisor

Accepted by..............................
Chairman, Departmental Committee on Graduate Students

EFFICIENT MULTIPLICATION IN SEMISIMPLE ALGEBRAS

by

Michael Conrad Loui

Submitted to the Department of Electrical Engineering and
Computer Science on 24 November 1976 in partial fulfillment
of the requirements for the degree of Master of Science

ABSTRACT

Let $\mathcal{A}$ be a finite-dimensional, associative, semisimple algebra over a field F. A generalization of the discrete Fourier transform is presented; this transform permits efficient multiplication of elements of $\mathcal{A}$: one multiplies corresponding projections of each element in the simple subalgebras into which $\mathcal{A}$ decomposes. It is demonstrated that when F contains at least $[\mathcal{A}:F]$ elements, the number of these simple subalgebras equals the number of distinct irreducible factors of the characteristic polynomial of the regular representation of the general element of $\mathcal{A}$. A new explicit calculation of this polynomial is given for an arbitrary algebra of a finite abelian group over a field with a primitive root of unity. Several examples are considered in detail: quotient polynomial rings, abelian group algebras, dihedral group algebras, and generalized quaternion group algebras.

This thesis emphasizes the minimization of nonscalar multiplications. A theorem previously promulgated is proved correctly: if $\pi(z)$, a polynomial of degree n, has k distinct irreducible factors in $F[z]$, then computing products in $F[z]/(\pi(z))$ requires at least 2n-k nonscalar multiplications.

THESIS SUPERVISOR: Alan S. Willsky

TITLE: Associate Professor of Electrical Engineering

ACKNOWLEDGEMENTS

First, I wish to thank Alan Willsky, my thesis supervisor, for his sagacious guidance and trenchant wit. This thesis has benefited from his astute, perspicacious insights.

It is a pleasure for me to acknowledge the assistance of three other colleagues. Charles Fiduccia, Assistant Professor of Computer Science at SUNY / Stony Brook, generously spent several hours trading novel ideas with me. Loren Platzman inspired one of the proofs in one of several enjoyable conversations. Robert Washburn, fellow Yale alumnus, patiently tolerated my innumerable incoherent monologues on my preliminary results and other mathematical trivia.

Finally, I am grateful for the financial support provided by a Fannie and John Hertz Foundation graduate fellowship.

TABLE OF CONTENTS

1. INTRODUCTION

Introduced by Cooley and Tukey in 1965 [C2], the fast Fourier transform (FFT) algorithm has influenced many researchers--particularly those in digital signal processing and in the theory of computation. Using the FFT, academicians have discovered efficient algorithms for multiplying polynomials and multiplying large integers; in both cases the concomitant convolutions are computed with a FFT [A1] [B3].

Multiplying polynomials $\sum_{i=0}^{n-1} a_i z^i$ and $\sum_{j=0}^{n-1} b_j z^j$ yields an acyclic convolution of the coefficient vectors $[a_0\ a_1\ \ldots\ a_{n-1}]$ and $[b_0\ b_1\ \ldots\ b_{n-1}]$:

$$\left(\sum_{i=0}^{n-1} a_i z^i\right)\left(\sum_{j=0}^{n-1} b_j z^j\right) = \sum_{k=0}^{2n-2} c_k z^k\ ,\quad c_k = \sum_{i+j=k} a_i b_j .$$

Suppose one multiplies these polynomials modulo z^n-1. This operation is then a (convolutional) multiplication of elements of the group algebra $F[Z_n]$, where F is a field and Z_n is the cyclic group of order n. We obtain $(\sum_{i=0}^{n-1} a_i z^i)(\sum_{j=0}^{n-1} b_j z^j) \equiv \sum_{k=0}^{n-1} c_k z^k \bmod (z^n-1)$, where

$$c_k = \sum_{i+j \equiv k \bmod n} a_i b_j = \sum_{i+j=k} a_i b_j + \sum_{i+j=k+n} a_i b_j$$

for $k = 0, \ldots, n-1$. In this case the coefficient vector $[c_0\ c_1\ \ldots\ c_{n-1}]$ is the cyclic convolution of $[a_0 \ldots a_{n-1}]$ and $[b_0 \ldots b_{n-1}]$. One can use the discrete Fourier transform to calculate a cyclic convolution: if $\alpha(z)$, $\beta(z)$, and $\gamma(z)$ are polynomials of degree $n-1$, then $\gamma(z) = \alpha(z)\beta(z) \bmod (z^n-1)$ if, and only if, $\alpha(\omega^k)\beta(\omega^k) = \gamma(\omega^k)$ for every $k = 0,1,\ldots,n-1$, where ω is a primitive nth root of unity. The vector whose components are $\alpha(\omega^0)$, $\alpha(\omega^1)$, ..., $\alpha(\omega^{n-1})$ is the discrete Fourier transform of the coefficient vector $[a_0 \ldots a_{n-1}]$ of $\alpha(z) = a_0 + a_1 z + \ldots a_{n-1}z^{n-1}$.

More generally, consider multiplication of elements of the group algebra $\mathbb{C}[G]$, where G is a finite group. If $\alpha = \sum_{g\in G} a_g g$ and $\beta = \sum_{g\in G} b_g g$, then

$$\alpha\beta = \gamma = \sum_{h\in G} c_h h, \quad \text{where} \quad c_h = \sum_{g\in G} a_g b_{g^{-1}h} \quad \text{for } h\in G.$$

These products arise in filtering problems for random finite group homomorphic sequential systems, studied by Willsky [W1]. A probability distribution on the state variable is expressed as an element of the group algebra $\mathbb{C}[G]$ in which G is the state space; distributions for input and output variables are treated analogously. Employing the irreducible, inequivalent matrix representations $\{T^1, \ldots, T^m\}$ of G, Willsky defines for the probability distributions a transform yielding a set of matrices. Let μ_i be the dimension of T^i for $i = 1,\ldots,m$. Then the transform pair for $\alpha = \sum_{g\in G} a_g g$ is

$$C^i(\alpha) = \frac{\mu_i}{n} \sum_{g\in G} a_g T^i(g^{-1})^{t} \text{ for } i = 1,\ldots,m;$$

$$a_g = \sum_{i=1}^{m} \sum_{j=1}^{\mu_i} \sum_{k=1}^{\mu_i} C^i(\alpha)_{jk} \; T^i(g)_{jk}, \tag{1.1}$$

where $C^i(\alpha)_{jk}$ is the j,k entry of $C^i(\alpha)$, and $T^i(g)_{jk}$ is the j,k entry of $T^i(g)$, and superscript t denotes the transpose of the matrix. To obtain the transform of the convolution of two pobability distributions, one multiplies corresponding matrices of the transforms of the distributions: $\alpha\beta = \gamma$ if, and only if, $\mathcal{C}^i(\alpha)\mathcal{C}^i(\beta) = \mathcal{C}^i(\gamma)$ for $i = 1,\dots,m$.

Fiduccia [F1] noticed that the FFT algorithm computes a discrete Fourier transform by successive division of polynomials. For instance, if the polynomial $\alpha(z)$ has degree 7 and ω is a primitive eighth root of unity, then one evaluates α at $\omega^0, \omega^1, \dots, \omega^7$ as follows:

$$\begin{aligned}
\alpha(z) &= q_0(z)[(z-\omega^0)(z-\omega^4)(z-\omega^2)(z-\omega^6)] + \alpha_0(z)\\
&= q_1(z)[(z-\omega^1)(z-\omega^5)(z-\omega^3)(z-\omega^7)] + \alpha_1(z)\\
\alpha_0(z) &= q_{00}(z)[(z-\omega^0)(z-\omega^4)] + \alpha_{00}(z)\\
&= q_{01}(z)[(z-\omega^2)(z-\omega^6)] + \alpha_{01}(z)\\
\alpha_{00}(z) &= q_{000}(z)(z-\omega^0) + \alpha(\omega^0)\\
&= q_{001}(z)(z-\omega^4) + \alpha(\omega^4)\\
\alpha_{01}(z) &= q_{010}(z)(z-\omega^2) + \alpha(\omega^2)\\
&= q_{011}(z)(z-\omega^6) + \alpha(\omega^6)\\
\alpha_1(z) &= q_{10}(z)[(z-\omega^1)(z-\omega^5)] + \alpha_{10}(z)\\
&= q_{11}(z)[(z-\omega^3)(z-\omega^7)] + \alpha_{11}(z)\\
\alpha_{10}(z) &= q_{100}(z)(z-\omega^1) + \alpha(\omega^1)\\
&= q_{101}(z)(z-\omega^5) + \alpha(\omega^5)\\
\alpha_{11}(z) &= q_{110}(z)(z-\omega^3) + \alpha(\omega^3)\\
&= q_{111}(z)(z-\omega^7) + \alpha(\omega^7).
\end{aligned}$$

Each polynomial division can be performed rapidly because every divisor has the form $z^i - \omega^j$; for example, $(z-\omega^1)(z-\omega^5)(z-\omega^3)(z-\omega^7) = z^4 - \omega^4$. Write $\phi(z) \bmod \phi_1(z)$ for the remainder (residue) when ϕ is divided by ϕ_1; for instance, $\alpha(\omega^3) = \alpha(z) \bmod (z-\omega^3)$. For any polynomials ϕ, ϕ_1, and ϕ_2, $(\phi \bmod (\phi_1\phi_2)) \bmod \phi_1 = \phi \bmod \phi_1$. This fact justifies the FFT.

The inverse FFT interpolates a polynomial by the reverse process--successive polynomial multiplication. For the example above the inverse transform reconstructs $\alpha(z)$ from the values $\alpha(\omega^0)$, $\alpha(\omega^1)$, ..., $\alpha(\omega^7)$. The FFT and its inverse each require $\mathcal{O}(n \log n)$ total arithmetic operations for a polynomial of degree n. In fact, Morgenstern [M2] established a $\mathcal{O}(n \log n)$ lower bound on the number of arithmetic operations used by any algorithm computing the discrete Fourier transform with additions and multiplications only by scalars of modulus 1.

Rader [R1] and Winograd [W6] discovered variants of the FFT. Rader's modification computes the discrete Fourier transform when n, the degree of the polynomial, is prime. Related to Rader's idea, Winograd's innovation is particularly fast for small n.

Generalizing the FFT, Moenck and Borodin [M1] [B2] devised algorithms for rapid modular representation of polynomials and Chinese remaindering. Let $\alpha(z)$ be a polynomial of degree n, and let $\psi_1(z)$, ..., $\psi_n(z)$ be relatively prime polynomials. Suppose we wish to calculate the residues $\alpha(z) \bmod \psi_1(z)$, ..., $\alpha(z) \bmod \psi_n(z)$; if all ψ_i have the form $z - u_i$, then this calculation evaluates α at u_1, ..., u_n. Imitating the FFT, one divides successively by products of $\{\psi_1, \ldots, \psi_n\}$: viz., $\psi_1\psi_2\cdots\psi_n$; ...; $\psi_1\psi_2\psi_3\psi_4$, ..., $\psi_{n-3}\psi_{n-2}\psi_{n-1}\psi_n$; $\psi_1\psi_2$, $\psi_3\psi_4$, ..., $\psi_{n-1}\psi_n$; ψ_1, ψ_2, ..., ψ_n. This Moenck-Borodin algorithm uses $\mathcal{O}(n (\log n)^2)$ total arithmetic operations.

Conversely, let $\psi_1(z), \ldots, \psi_n(z)$ be relatively prime polynomials. For any polynomials $\alpha_1(z), \ldots, \alpha_n(z)$ such that $\deg \alpha_i < \deg \psi_i$ for each i, the Chinese Remainder Theorem asserts the existence of a polynomial $\alpha(z)$ such that $\alpha(z) \equiv \alpha_i(z) \pmod{\psi_i(z)}$ for $i = 1,\ldots,n$. Define $\Psi_i(z)$ and $\Phi_i(z)$ for $i = 1,\ldots,n$ so that

$$\Psi_i(z) = \left(\prod_{k=1}^{n} \psi_k(z)\right) / \psi_i(z) \quad \text{and}$$
$$\Phi_i(z)\Psi_i(z) \equiv 1 \pmod{\psi_i(z)}; \tag{1.2}$$

one can choose each Φ_i so that $\deg \Phi_i < \deg \psi_i$. Then

$$\alpha(z) = \sum_{i=1}^{n} \Phi_i(z)\Psi_i(z)\alpha_i(z) \tag{1.3}$$

has the desired property. When all ψ_i have the form $z - u_i$ for some scalars u_i, equations (1.2) and (1.3) are the Lagrange interpolation formula. Suitably reversing the Moenck-Borodin algorithm for calculating modular residues produces a rapid Chinese Remainder Algorithm that uses $\mathcal{O}(n (\log n)^2)$ total arithmetic operations on problems of size n. One may find complete expositions of modular algorithms and the FFT in [A1] and [B3].

* * *

In this thesis we consider efficient computations of products in finite-dimensional, associative algebras, particularly semisimple algebras. All algebras in subsequent chapters are associative and finite-dimensional over a field.

In Chapter 2 we present a model of computation--the straight-line program--and define the multiplicative complexity of an algebra to be the minimum number of nonscalar rudimentary multiplications necessary for computing products in the algebra. We then give a brief exegesis of recent results on quotient polynomial rings, including a correct proof of a theorem announced by Winograd [W5]: the multiplicative complexity of $F[z]/(\pi(z))$ is at least 2n-k, where n is the degree of $\pi(z)$ and k is the number of its distinct irreducible factors.

After summarizing definitions and properties of representations of semisimple algebras, we define the general element ξ of a semisimple algebra A and the characteristic polynomial χ of the regular representation of ξ. We establish a one-one correspondence between distinct, irreducible factors of χ and the irreducible, inequivalent representations of A when the base field is sufficiently large.

Using the structure of the semisimple algebra A over a field F, we devise a generalization of the discrete Fourier transform: to multiply elements α and β of A, one multiplies the projections of α and β in the simple subalgebras of A; each projection is a small matrix with entries in a division algebra over F. This technique uses fewer nonscalar multiplications than others.

Chapter 4 presents examples to illustrate the ideas of Chapter 3. Related to our general methods are known algorithms for quotient polynomial algebras and abelian group algebras. For quotient polynomial algebras and abelian group algebras in which the base field contains a primitive root of unity we calculate the characteristic polynomial of the regular representation of the general element. Furthermore, we give algorithms for computing products in algebras of dihedral groups and generalized quaternion groups. We consider algebras of these groups not only over the complex numbers $\mathbb{C}$, but also over the reals $\mathbb{R}$ and rationals $\mathbb{Q}$.

In Chapter 5 we suggest topics for further study.

2. BACKGROUND

2A. The Problem

A finite-dimensional, associative algebra A of dimension n over a field F is simultaneously a ring $(A,+,\cdot,0,1)$ (with multiplicative identity 1) and a vector space of finite dimension n over F. An algebra is semisimple if it contains no nilpotent left ideals: for no left ideal J in A and no positive integer m does $\{u_1u_2\cdots u_m |$ every $u_i \in J\}$ equal (0), the zero ideal. Let $\{v_1, v_2, \ldots, v_n\}$ be a basis for algebra A over field F. Elements of A have the form $a_1v_1 + a_2v_2 + \ldots + a_nv_n$, where the coefficients $a_1, a_2, \ldots, a_n$ are in F.

Example. The complex numbers $\mathbb{C}$ form an algebra of dimension 2 over the real numbers $\mathbb{R}$. Elements of $\mathbb{C}$ can be written $a_1\cdot 1 + a_2\cdot\sqrt{-1}$ with a_1, a_2 in $\mathbb{R}$. Any finite-dimensional extension field is an algebra over the base field.

Example. The noncommutative division ring of real quaternions $Q_{\mathbb{R}}$ is an algebra of dimension 4 over $\mathbb{R}$ with basis $1, \hat{i}, \hat{j}, \hat{k}$ and multiplication defined by

$$\hat{i}^2 = \hat{j}^2 = \hat{k}^2 = -1, \quad \hat{i}\hat{j} = -\hat{j}\hat{i} = \hat{k},$$
$$\hat{j}\hat{k} = -\hat{k}\hat{j} = \hat{i}, \quad \hat{k}\hat{i} = -\hat{i}\hat{k} = \hat{j}.$$

Each quaternion has the form $a_0 + a_1\hat{i} + a_2\hat{j} + a_3\hat{k}$, where a_0, a_1, a_2, and a_3 are in $\mathbb{R}$. Analogously, one can define quaternions over any field F: $Q_F = \{a_0 + a_1\hat{i} + a_2\hat{j} + a_3\hat{k} |\ a_0, a_1, a_2, a_3 \text{ in } F\}$. If char F = 2, then Q_F is not semisimple because $1 + \hat{i}$ generates a nilpotent left ideal.

Example. Let $\pi(z)$ be a polynomial in $F[z]$. The quotient ring $F[z]/(\pi(z))$ is an algebra over the field F of dimension $\deg \pi$, the degree of π. This algebra is semisimple if, and only if, the irreducible factors of $\pi(z)$ in $F[z]$ have multiplicity 1. If an irreducible factor $\pi_0(z)$ of $\pi(z)$ has multiplicity greater than 1, then the ideal generated by $\pi(z)/\pi_0(z)$ is nilpotent.

Example. The algebra $F[G]$ of a finite group G over a field F is the set of all formal sums $\sum_{g\in G} a_g \cdot g$ with every a_g in F. Multiplication of elements of $F[G]$ is defined by

$$\left(\sum_{g\in G} a_g g\right)\left(\sum_{g\in G} b_g g\right) = \sum_{h\in G} c_h h, \quad \text{where } c_h = \sum_{g\in G} a_g b_{g^{-1}h}.$$

Suppose G is abelian and the prime p divides the order of G; let e be the identity of G and let $\hat{g}$ in G have order p; if F has characteristic p, then the ideal in $F[G]$ generated by $\hat{g} - e$ is nilpotent, and hence $F[G]$ is not semisimple. For any algebra $F[G]$, Maschke's Theorem [C3] asserts that $F[G]$ is semisimple if the characteristic of F does not divide the cardinality of G. Much is known about the representations of semisimple algebras of finite groups over algebraically closed fields.

We seek efficient algorithms for computing products of elements in finite-dimensional, associative algebras, especially semisimple algebras. Our model of computation is the straight-line program. For each algebra A we devise a straight-line program P that computes the product of a pair of elements of A. Let $\{v_1, \ldots, v_n\}$ be a basis for A. To multiply $\alpha = a_1v_1 + \ldots + a_nv_n$ and $\beta = b_1v_1 + \ldots + b_nv_n$ yielding $\gamma = \alpha\beta = c_1v_1 + \ldots +$

$c_n v_n$, we input $\{a_1, \ldots, a_n, b_1, \ldots, b_n\}$ to P, which calculates $\{c_1, \ldots, c_n\}$. Think of $\{a_1, \ldots, a_n, b_1, \ldots, b_n\}$ as a set of indeterminates; the $c_1, \ldots, c_n$ are bilinear expressions in $\{a_1, \ldots, a_n\}$ and $\{b_1, \ldots, b_n\}$.

Formally, a straight-line program for our problem is a finite sequence of statements $\theta_1, \theta_2, \ldots, \theta_r$ of the form:

(a) $\theta_k \leftarrow a_k$ for $k = 1,\ldots,n$;

(b) $\theta_k \leftarrow b_{k-n}$ for $k = n+1,\ldots,2n$;

(c) for $k > 2n$ either

(i) $\theta_k \leftarrow u\theta_i$ for some scalar $u \epsilon F$, some $i < k$; or

(ii) $\theta_k \leftarrow \theta_i + \theta_j$ for some $i, j < k$; or

(iii) $\theta_k \leftarrow \theta_i \cdot \theta_j$ for some $i, j < k$.

Define an evaluation of a statement to be the map $\mathscr{E}\colon \{\theta_1, \ldots, \theta_r\} \rightarrow F[a_1,\ldots,a_n,b_1,\ldots b_n]$ such that

$$\begin{aligned}
\mathscr{E}(\theta_k) &= a_k && \text{if } 1 \le k \le n,\\
\mathscr{E}(\theta_k) &= b_{k-n} && \text{if } n+1 \le k \le 2n,\\
\mathscr{E}(\theta_k) &= u\,\mathscr{E}(\theta_i) && \text{for form (i)},\\
\mathscr{E}(\theta_k) &= \mathscr{E}(\theta_i) + \mathscr{E}(\theta_j) && \text{for form (ii)},\\
\mathscr{E}(\theta_k) &= \mathscr{E}(\theta_i)\,\mathscr{E}(\theta_j) && \text{for form (iii)}.
\end{aligned}$$

The program P computes $\gamma = \alpha\beta$ from $\{a_1, \ldots, a_n, b_1, \ldots, b_n\}$ if for some $\theta_{k_1}, \ldots, \theta_{k_n}$ of P we have $c_1 = \mathscr{E}(\theta_{k_1}), \ldots, c_n = \mathscr{E}(\theta_{k_n})$.

Example. Suppose we multiply complex numbers over the field of reals: $\alpha\beta = (a_1v_1 + a_2v_2)(b_1v_1 + b_2v_2) = (c_1v_1 + c_2v_2)$, where $c_1 = a_1b_1 - a_2b_2$ and $c_2 = a_1b_2 + a_2b_1$. A straight-line program for the product is:

$$\begin{array}{lll}
\theta_1 \leftarrow a_1 & \theta_5 \leftarrow \theta_1 \cdot \theta_3 & \theta_{10} \leftarrow \theta_1 + \theta_2 \\
\theta_2 \leftarrow a_2 & \theta_6 \leftarrow (-1)\theta_5 & \theta_{11} \leftarrow \theta_3 + \theta_4 \\
\theta_3 \leftarrow b_1 & \theta_7 \leftarrow \theta_2 \cdot \theta_4 & \theta_{12} \leftarrow \theta_{10} \cdot \theta_{11} \\
\theta_4 \leftarrow b_2 & \theta_8 \leftarrow (-1)\theta_7 & \theta_{13} \leftarrow \theta_{12} + \theta_6 \\
 & \theta_9 \leftarrow \theta_5 + \theta_8 & \theta_{14} \leftarrow \theta_{13} + \theta_8
\end{array}$$

This program computes $\alpha\beta$ because $c_1 = \mathscr{E}(\theta_9)$ and $c_2 = \mathscr{E}(\theta_{14})$.

A statement of type (iii) is a *nonscalar multiplication* if neither $\mathscr{E}(\theta_i)$ nor $\mathscr{E}(\theta_j)$ is an element of F. (Recall that $a_1, \ldots, a_n, b_1, \ldots, b_n$ are *indeterminates*.) For instance, the program in the example above contains three nonscalar multiplications. For every algebra $\mathcal{A}$ we desire a program for computing products in $\mathcal{A}$ having the fewest nonscalar multiplications. The *multiplicative complexity* of an algebra is the minimum number of nonscalar multiplications necessary for computing a product in the algebra. The multiplicative complexity of $\mathcal{A}$ over F is at most the square of the dimension $n = [\mathcal{A}:F]$: one could compute a set of n^2 products $\{a_ib_j \mid 1 \le i \le n, 1 \le j \le n\}$; the $c_1, \ldots, c_n$ are linear combinations of elements of this set.

Assessing the algebraic complexity of a problem, researchers often count only multiplications because on an electronic computer multiplications generally require more time than additions: the number of multiplications seems to characterize the speed of a program. We ignore scalar multiplications (cf. [F2], [W3]). When the base field is the field of rationals $\mathbb{Q}$, scalar multiplications by fixed rational numbers can be simulated by additions.

Because we compute bilinear forms, we may assume that all nonscalar multiplication steps, when evaluated, are products of linear combinations of the given indeterminates.

Theorem 2.1 ([W3], [B3, p. 35]). For every program computing a set of bilinear forms in $\{a_1, \ldots, a_n\}$ and $\{b_1, \ldots, b_n\}$ there exists another program computing the same set of bilinear forms with the same number of nonscalar multiplications, each of which, when evaluated by the map $\mathscr{E}$, has the form $L_1 \cdot L_2$, where L_1 and L_2 are linear combinations of $\{a_1, \ldots, a_n, b_1, \ldots, b_n\}$.

In a bilinear program all nonscalar multiplications have the form $L_1 \cdot L_2$, where L_1 is a linear combination of $\{a_1, \ldots, a_n\}$, L_2 a linear combination of $\{b_1, \ldots, b_n\}$.

Theorem 2.2 ([P1]). If the $\{a_1, \ldots, a_n\}$ do not commute with the $\{b_1, \ldots, b_n\}$, then for any program computing bilinear forms in $\{a_1, \ldots, a_n\}$ and $\{b_1, \ldots, b_n\}$ there exists a bilinear program computing the same forms having the same number of nonscalar multiplications.

Theorem 2.3 ([P1]). If the $\{a_1, \ldots, a_n\}$ do commute with the $\{b_1, \ldots, b_n\}$ [the usual case], then for any program P computing bilinear forms in $\{a_1, \ldots, a_n\}$ and $\{b_1, \ldots, b_n\}$ there exists a bilinear program computing the same forms having at most twice the number of nonscalar multiplications in $\mathcal{P}$.

Hopcroft and Kerr [H2] demonstrated that commutativity can reduce the number of nonscalar multiplications. If $m \geq 5$, the product of $m \times 2$ and 2×2 matrices over a commutative ring can be computed with fewer nonscalar multiplications than for the product over a noncommutative ring.

2B. Prior Results

Winograd [W3] first established a general framework for obtaining lower bounds on the number of nonscalar multiplications needed to compute functions. His technique and related techniques of Fiduccia are summarized in [A1] and [B3].

Fiduccia and Zalcstein [F2] derived a lower bound on the multiplicative complexity of division algebras.

Theorem 2.4 ([F3]). Let $\mathcal{D}$ be a division algebra of dimension $n = [\mathcal{D}:F]$ over a field F. The multiplicative complexity of $\mathcal{D}$ is at least $2n-1$.

Using the Chinese Remainder Algorithm extensively, they constructed an efficient algorithm for computing products in quotient polynomial rings. Let $\pi(z)$ be a fixed polynomial of degree n in $F[z]$ having prime factorization $\pi(z) = \pi_1(z)^{\nu_1}\pi_2(z)^{\nu_2}\cdots\pi_k(z)^{\nu_k}$; let $n_i = \deg \pi_i^{\nu_i}$ for $i = 1,\ldots,k$.

Algorithm 2.5. To compute the coefficients $\{c_0, \ldots, c_{n-1}\}$ of the product $\gamma(z) = \alpha(z)\beta(z) = c_0 + c_1 z + \ldots + c_{n-1}z^{n-1}$ of $\alpha(z) = a_0 + a_1 z + \ldots + a_{n-1}z^{n-1}$ and $\beta(z) = b_0 + b_1 z + \ldots + b_{n-1}z^{n-1}$ from $\{a_0, a_1, \ldots, a_{n-1}, b_0, \ldots, b_{n-1}\}$ in $F[z]/(\pi(z))$.

Step 1. Perform Step 2 through Step 4 for $i = 1,\ldots,k$.

Step 2. Compute $\hat{\alpha}_i(z) = \alpha(z) \bmod \pi_i(z)^{\nu_i}$ and $\hat{\beta}_i(z) = \beta(z) \bmod \pi_i(z)^{\nu_i}$; the degrees of the polynomials $\hat{\alpha}_i(z)$ and $\hat{\beta}_i(z)$ are at most n_i-1.

Step 3. For any $2n_i-1$ distinct elements $u_1, u_2, \ldots, u_{2n_i-1}$ in F multiply $\hat{\alpha}_i(u_j)\hat{\beta}_i(u_j)$ for $j = 1,2,\ldots,2n_i-1$.

Step 4. Using the Lagrange interpolation formula (a special case of the Chinese Remainder Algorithm), recover the polynomial $\hat{\gamma}_i(z) = [\hat{\alpha}_i(z)\hat{\beta}_i(z) \bmod ((z - u_1)\cdots(z - u_{2n_i-1}))] \bmod \pi_i(z)^{\nu_i}$. The product of $\hat{\alpha}_i(z)$ and $\hat{\beta}_i(z)$ has degree at most $2n_i-2$; thus, $\hat{\alpha}_i(z)\hat{\beta}_i(z) = \hat{\alpha}_i(z)\hat{\beta}_i(z) \bmod ((z - u_1)\cdots(z - u_{2n_i-1}))$.

Step 5. Use the Chinese Remainder Algorithm to obtain $\gamma(z) = \alpha(z)\beta(z) \bmod \pi(z)$ from all the $\gamma_i(z) = \alpha(z)\beta(z) \bmod \pi_i(z)^{\nu_i}$. □

The idea implemented in Steps 3 and 4 also appears in the Toom-Cook algorithm for multiplying large integers [K1]. Only Step 3 involves nonscalar multiplications. There are $\sum_{i=1}^{k}(2n_i-1) = 2n-k$ nonscalar multiplications.

Theorem 2.6 ([F2]). Let the fixed polynomial $\pi(z)$ have prime factorization $\pi_1(z)^{\nu_1}\cdots\pi_k(z)^{\nu_k}$ in $F[z]$, where F is a field. Let $n = \deg \pi$ and $n_i = \deg \pi_i^{\nu_i}$ for $i = 1,\ldots,k$. If F has at least $\max\{2n_1-1, \ldots,$

$2n_k$-1} distinct elements, then the multiplicative complexity of $F[z]/(\pi(z))$ is at most 2n-k.

Furthermore, Fiduccia and Zalcstein devised an algorithm for multiplying elements of commutative semisimple algebras over perfect fields. In Section 4A we shall show that this algorithm is a special case of our general methods.

Theorem 2.7 ([F2]). Let A be a commutative semisimple algebra of dimension n over a perfect field F. Suppose A decomposes into a direct sum of k simple algebras: $A \cong A_1 \oplus \ldots \oplus A_k$, where each A_i is simple. Let $n_{\max} = \max\{[A_1:F], \ldots, [A_k:F]\}$. If F has at least $2n_{\max}-1$ elements, then the multiplicative complexity of A over F is at most 2n-k.

Proof. Each A_i is a commutative division algebra over F--i.e., a finite-dimensional extension field of F. Since each A_i is algebraic over F, which is perfect, there exist irreducible polynomials π_i in F[z] of degrees $\deg \pi_i = [A_i:F]$ such that $A_i \cong F[z]/(\pi_i(z))$ for $i = 1,\ldots,k$. □

Winograd [W5] attempted to prove that Algorithm 2.5 is optimal, but his arguments contain a flaw that we now correct. Let $\{a_1, \ldots, a_n\}$ be a set of n distinct indeterminates. Let F be a field.

Definition. A set of r-vectors $\{\underline{v}_1, \ldots, \underline{v}_k\}$ with components in $F[a_1,\ldots,a_n]$ is linearly independent modulo F^r if whenever $\sum_{i=1}^k u_i\underline{v}_i \in F^r$ for some $u_1, \ldots, u_k$ in F we have $u_1 = \ldots = u_k = 0$.

Lemma 2.8. Let $\{\underline{v}_1, \ldots, \underline{v}_k\}$ be a set of r-vectors with components in $F[a_1,\ldots,a_n]$. Suppose $k \geq 2$ and this set has a subset of δ vectors linearly independent modulo F^r. Then for any $u_2, \ldots, u_k$ in F the set $\{\underline{v}_i + u_i\underline{v}_1 | i = 2,\ldots,k\}$ includes a subset of δ-1 vectors linearly independent modulo F^r.

Proof. [Al, p. 435]. □

Definition. Let P be a program computing $M\underline{x}$, where matrix M has entires in $F[a_1,\ldots,a_n]$ and $\underline{x}$ is a vector with components in $F[b_1,\ldots,b_n]$. A nonscalar multiplication in P is active if one of the operands involves an indeterminate b_i and the other is not an element of F.

The next lemma generalizes Winograd's technique for deriving lower bounds [W3].

Lemma 2.9. Let $\underline{y}$ be a r-vector with components in $F[a_1,\ldots,a_n]$. Let $\underline{x}$ be a vector of p linearly independent linear combinations of indeterminates $\{b_1, \ldots, b_n\}$. If matrix M with entries in $F[a_1,\ldots,a_n]$ has δ columns linearly independent modulo F^r, and if $\delta > 1$, then any program computing $M\underline{x} + \underline{y}$ requires at least δ active multiplications.

Proof. Modifying the proof of Theorem 12.2 in [Al], we proceed by induction on δ.

Case: $\delta = 1$. Some entry ε in M is in $F[a_1,\ldots,a_n]$ but not in F. Computing $M\underline{x} + \underline{y}$ requires a multiplication εb_j for some j.

Case: $\delta > 1$. Let $\underline{b}$ be the column vector $[b_1 \ldots b_p]^t$ and let $\underline{x} = Q\underline{b}$, where Q is a $p \times p$ invertible matrix with entries in F. By induction, any program P for $M\underline{x} + \underline{y}$ has at least $\delta - 1 \geq 1$ active multiplications. Let $\theta \leftarrow \theta' \cdot \theta''$ be the first active multiplication in $\mathcal{P}$:

$$\mathscr{E}(\theta') = \phi(a_1,\ldots,a_n) + \textstyle\sum_{i=1}^{p} c_i b_i,$$

where ϕ is a polynomial in $F[a_1,\ldots,a_n]$ and every $c_i \in F$. Without loss of generality, take $c_1 \neq 0$. Define $d_1, \ldots, d_p$ by $[d_1 \ldots d_p] = [c_1 \ldots c_p]Q^{-1}$. Then $\sum_{i=1}^{p} c_i b_i = \sum_{i=1}^{p} d_i x_i$. Assume $d_1 \neq 0$ without loss of generality.

From $\mathcal{P}$ we can construct a new program $\mathcal{P}'$ with one fewer active multiplication. This new program will compute $M'\underline{x}' + \underline{y}'$, where the $r\times(p-1)$ matrix M' will have $\delta-1$ columns linearly independent modulo F^r. By the inductive hypothesis, $\mathcal{P}'$ will have at least $\delta-1$ active multiplications; therefore, $\mathcal{P}$ has at least δ active multiplications.

To obtain $\mathcal{P}'$ from $\mathcal{P}$ we use a substitution argument [B3, Ch. 2]. Replace

$$\begin{aligned} b_1 &\text{ by } -c_1^{-1}[\phi(a_1,\ldots,a_n) + \textstyle\sum_{i=2}^{p} c_i b_i], \\ x_1 &\text{ by } -d_1^{-1}[\phi(a_1,\ldots,a_n) + \textstyle\sum_{i=2}^{p} d_i x_i]; \end{aligned} \tag{2.1}$$

these substitutions force $\mathscr{E}(\theta')$ to be zero. Compared with $\mathcal{P}$, the program $\mathcal{P}'$ has one fewer active multiplication.

We now define M', $\underline{x}'$, and $\underline{y}'$. Let the ith column of M be $\underline{m}_i$ and set the (i-1)th column of M' to be

$$\underline{m}'_{i-1} = \underline{m}_i - \frac{d_i}{d_1}\underline{m}_1 \qquad \text{for } i = 2,\ldots,p.$$

By Lemma 2.8, M' has at least $\delta-1$ linearly independent columns (modulo F^r). Let $\underline{q}_i$ be the ith column of Q. With the substitution (2.1) we have

$$\begin{bmatrix} x_1 \\ \cdot \\ \cdot \\ \cdot \\ x_p \end{bmatrix} = \tilde{Q} \begin{bmatrix} b_2 \\ \cdot \\ \cdot \\ \cdot \\ b_p \end{bmatrix} + Q \begin{bmatrix} -c_1^{-1}\phi(a_1,\ldots,a_n) \\ 0 \\ \vdots \\ 0 \end{bmatrix},$$

where

$$\tilde{Q} = [\underline{q}_2 - \frac{c_2}{c_1}\underline{q}_1 \mid \cdots \mid \underline{q}_p - \frac{c_p}{c_1}\underline{q}_1].$$

The matrix $\tilde{Q}$ has p-1 linearly independent columns, and hence its row rank is p-1. A simple calculation reveals that the first row of $\tilde{Q}$ is a linear combination of the others:

$$q_{1j} - \frac{c_j}{c_1}q_{11} = -\sum_{i=2}^{p} \frac{d_i}{d_1}(q_{ij} - \frac{c_j}{c_1}q_{i1}) \qquad \text{for } j = 2,\ldots,p.$$

Delete the first row of $\tilde{Q}$ to form a $(p-1)\times(p-1)$ matrix Q'. Since $\tilde{Q}$ has row rank p-1, Q' is nonsingular. Define $\underline{x}'$ and $\underline{y}'$ by

$$\underline{x}' = \begin{bmatrix} x_1' \\ \vdots \\ x_{p-1}' \end{bmatrix} = Q' \begin{bmatrix} b_2 \\ \vdots \\ b_p \end{bmatrix},$$

$$\underline{y}' = \underline{y} + M \begin{bmatrix} -d_1^{-1}\phi(a_1,\ldots,a_n) \\ 0 \\ \vdots \\ 0 \end{bmatrix} + M' \begin{bmatrix} -q_{21}c_1^{-1}\phi(a_1,\ldots,a_n) \\ -q_{31}c_1^{-1}\phi(a_1,\ldots,a_n) \\ \vdots \\ -q_{p1}c_1^{-1}\phi(a_1,\ldots,a_n) \end{bmatrix}.$$

Then program P' computes $M'\underline{x}' + \underline{y}'$, where the components of $\underline{x}'$ are linearly independent linear combinations of p-1 indeterminates $\{b_2,\ldots,b_p\}$.

The program $\mathcal{P}'$ computes $M'\underline{x}' + \underline{y}'$, where M' has $\delta-1$ columns linearly independent modulo F^r. Since $\mathcal{P}'$ has at least $\delta-1$ nonscalar multiplications, $\mathcal{P}$ has at least δ. □

Theorem 2.10 ([W5]). Let F be a field, and let the fixed monic polynomial $\pi(z)$ of degree n have prime factorization $\pi_1(z)^{\nu_1}\cdots\pi_k(z)^{\nu_k}$ in $F[z]$. The multiplicative complexity of $F[z]/(\pi(z))$ is at least $2n-k$.

Proof. Part 1. Let $\alpha(z) = a_0 + a_1 z + \ldots + a_{n-1}z^{n-1}$ and $\beta(z) = b_0 + b_1 z + \ldots + b_{n-1}z^{n-1}$, where $a_0, \ldots, a_{n-1}, b_0, \ldots, b_{n-1}$ are distinct indeterminates. Let P be the $n \times n$ companion matrix for $\pi(z) = p_0 + p_1 z + \ldots + p_{n-1}z^{n-1} + z^n$:

$$P = \begin{bmatrix} 0 & & & 0 & -p_0 \\ 1 & 0 & & & -p_1 \\ & 1 & \ddots & & \vdots \\ & & \ddots & 0 & -p_{n-2} \\ 0 & & & 1 & -p_{n-1} \end{bmatrix},$$

Let $\underline{x} = [b_0\ b_1\ \ldots\ b_{n-1}]^t$, $\underline{\alpha} = [a_0\ a_1\ \ldots\ a_{n-1}]^t$, and $M = a_0 I_n + a_1 P + \ldots + a_{n-1}P^{n-1} = [\underline{\alpha} \mid P\underline{\alpha} \mid \ldots \mid P^{n-1}\underline{\alpha}]$. For $i = 1,\ldots,k$ let $n_i = \deg \pi_i(z)^{\nu_i}$, and let $\alpha_{i0} + \alpha_{i1}z + \ldots + \alpha_{i,n_i-1}z^{n_i-1} = \alpha(z) \bmod \pi_i(z)^{\nu_i}$ and $\beta_{i0} + \beta_{i1}z + \ldots + \beta_{i,n_i-1}z^{n_i-1} = \beta(z) \bmod \pi_i(z)^{\nu_i}$; let $\underline{x}^i = [\beta_{i0}\ \beta_{i1}\ \cdots\ \beta_{i,n_i-1}]^t$ and $\underline{\alpha}^i = [\alpha_{i0}\ \alpha_{i1}\ \cdots\ \alpha_{i,n_i-1}]^t$. The components of each $\underline{x}^i$ and $\underline{\alpha}^i$ are linear combinations of $\{b_0, \ldots, b_{n-1}\}$ and $\{a_0, \ldots, a_{n-1}\}$. Let P_i be the $n_i \times n_i$ companion matrix for $\pi_i^{\nu_i}$ and $M_i = \alpha_{i0}I_{n_i} + \alpha_{i1}P_i + \cdots + \alpha_{i,n_i-1}P_i^{n_i-1} = [\underline{\alpha}^i \mid P_i\underline{\alpha}^i \mid \ldots \mid P_i^{n_i-1}\underline{\alpha}^i]$ for each $i = 1,\ldots,k$.

Computing the coefficients of $\alpha(z)\beta(z) \bmod \pi(z)$ is equivalent to multiplying $M\underline{x}$; computing $\alpha(z)\beta(z) \bmod \pi_i(z)^{\nu_i}$ is equivalent to multiplying $M_i\underline{x}^i$. Since $F[z]/(\pi(z)) \cong F[z]/(\pi_1(z)^{\nu_1}) \oplus \ldots \oplus F[z]/(\pi_k(z)^{\nu_k})$, changing the basis of $M\underline{x}$ yields

$$\tilde{M}\tilde{\underline{x}} = \begin{bmatrix} M_1 & 0 & & 0 \\ 0 & M_2 & & \\ & & \ddots & \\ 0 & & & M_k \end{bmatrix} \begin{bmatrix} \underline{x}^1 \\ \underline{x}^2 \\ \vdots \\ \underline{x}^k \end{bmatrix}.$$

The Chinese Remainder Theorem guarantees that computing $M\underline{x}$ and computing $\tilde{M}\tilde{\underline{x}}$ require the same number of nonscalar multiplications: an answer to one matrix product yields, with some additions and scalar multiplications, a solution to the other product. We shall prove that any program computing $\tilde{M}\tilde{\underline{x}}$ must have at least 2n-k nonscalar multiplications.

Part 2. We demonstrate that all components of vectors $\underline{\alpha}^i$ are linearly independent linear combinations of $\{a_0, \ldots, a_{n-1}\}$; this result will imply that all components of the $\underline{x}^i$ are linearly independent linear combinations of $\{b_0, \ldots, b_{n-1}\}$.

The vector space V of matrices M generated from polynomials $\alpha(z)$ in Part 1 is spanned by $\{I_n, P, \ldots, P^{n-1}\}$, which is a basis for this space; this subspace V of $F^{n\times n}$ has dimension n. Changing the basis, i.e., mapping matrices M into matrices $\tilde{M}$, is an invertible transformation from V to $\tilde{V}$, a vector space of matrices $\tilde{M}$; since V has dimension n over F, the space $\tilde{V}$ also has dimension n.

The entries of each matrix M_i are linear combinations of $\{\alpha_{i0}, \ldots, \alpha_{i,n_i-1}\}$. Let R be the set of pairs (i,j) such that $\{\alpha_{ij} \mid (i,j) \notin R\}$

is the set of all α_{ij} that can be obtained from linear combinations of $\{\alpha_{ij} \mid (i,j) \in R\}$. Suppose, contrary to what we want to prove, R has fewer than n pairs. Then we can obtain all matrices $\tilde{M}$ by linear combinations of

$$\left\{ \left. \begin{bmatrix} 0 & & & & & 0 \\ & \ddots & & & & \\ & & 0 & & & \\ & & & P_i^j & & \\ & & & & 0 & \\ & & & & & \ddots \\ 0 & & & & & & 0 \end{bmatrix} \right| (i,j) \in R \right\} .$$

But then $\tilde{V}$ is spanned by fewer than n matrices, and $\tilde{V}$ has dimension less than n. Having obtained a contradiction, we conclude that all the α_{ij} are linearly independent.

Part 3. We revise the proof of Theorem 4 in [W5] to show that computing $\tilde{M}\underline{\tilde{x}}$ requires at least 2n-k nonscalar multiplications.

For i = 1,...,k, define $\hat{W}_i$ to be the one dimensional subspace of F^{n_i} spanned by the row vector [1 0 0 ... 0]. One may verify that if $\underline{w} \in \hat{W}_i$, then for any one-variable polynomial ϕ of degree $\deg \phi < n_i$, $\underline{w}\phi(P_i) = 0$ only if $\phi = 0$. Let $\overline{W}_i$ be the subspace such that $F^{n_i} = \hat{W}_i \oplus \overline{W}_i$. This subspace $\overline{W}_i$ includes $\{\underline{w} \in F^{n_i} \mid$ there exists some nonzero polynomial ϕ of one variable with $\deg \phi < n_i$ such that $\underline{w}\phi(P_i) = 0\}$.

Let t be the minimum number of nonscalar multiplications required for computing $\alpha(z)\beta(z) \bmod \pi(z)$. By Theorem 2.1, we may assume $\tilde{M}\underline{\tilde{x}} = K\underline{\zeta}$, where K is a n×t matrix with entries in F, and the components of $\underline{\zeta}$ are the results of nonscalar multiplications involving the indeterminates $\{a_0, \ldots, a_{n-1}, b_0, \ldots, b_{n-1}\}$. Because the first column of $\tilde{M}$ contains only zeroes and the linearly independent expressions α_{10}, α_{11}, ..., α_{1,n_1-1}, for no row vector $\underline{v}$ in F^n does $\underline{v}\tilde{M}$ equal zero. Therefore, K has rank n; assume that the first n columns of K are linearly independent. Let Q be the n×n nonsingular matrix such that $QK = [I_n \mid K']$, where I_n is the n×n identity matrix and K' is a n×(t-n) matrix.

In order to apply Lemma 2.9, we shall find a row vector $\underline{c}$ in F^n with at most k nonzero components such that the n columns of the 1×n matrix $\underline{c}Q\tilde{M}$ are linearly independent. It will follow that computing $\underline{c}Q\tilde{M}\tilde{\underline{x}} = \underline{c}[I_n \mid K']\underline{\zeta}$ requires at least n nonscalar multiplications; because $\underline{c}$ has at most k nonzero components, $\underline{c}[I_n \mid K']\underline{\zeta}$ can be computed with only k + (t-n) nonscalar multiplications. Thus, we shall show that $k + t - n \geq n$.

To find this vector $\underline{c}$, we seek some linear combination $\underline{c}Q = \underline{d}$ of k rows of Q such that if $\underline{d}$ is partitioned $[\underline{d}_1 \mid \underline{d}_2 \mid \ldots \mid \underline{d}_k]$, where each each $\underline{d}_i$ has n_i components, then $\underline{d}_i \notin \overline{W}_i$. First, partition $Q = [Q^1 \mid \ldots \mid Q^k]$, where each Q^i comprises n_i columns of Q. For each i write $Q^i = \hat{Q}^i + \overline{Q}^i$ such that every row of $\hat{Q}^i$ is in $\hat{W}_i$, and every row of $\overline{Q}^i$ is in $\overline{W}_i$. Let $\hat{Q} = [\hat{Q}^1 \mid \ldots \mid \hat{Q}^k]$ and $\overline{Q} = [\overline{Q}^1 \mid \ldots \mid \overline{Q}^k]$. Because $Q = \hat{Q} + \overline{Q}$ is nonsingular, the map $\underline{v} \mapsto \underline{v}Q$ is surjective onto F^n. Thus, the range of $\hat{Q}$ is all of $\hat{W}_1 \oplus \ldots \oplus \hat{W}_k$; that is, for every $\underline{w}$ in $\hat{W}_1 \oplus \ldots \oplus \hat{W}_k$ there exists $\underline{v}$ in F^n such that $\underline{w} = \underline{v}\hat{Q}$. The matrix $\hat{Q}$ has k linearly independent rows $\{\hat{q}_{j_1}, \ldots, \hat{q}_{j_k}\}$ because $\hat{W}_1 \oplus \ldots \oplus \hat{W}_k$ has dimension k. Let $\underline{y} = [\underline{y}_1 \mid \ldots \mid \underline{y}_k]$ be any row vector in F^n such that $\underline{y}_i \in \hat{W}_i$ and $\underline{y}_i \neq 0$ for i = 1,...,k. Some linear combination of $\{\hat{q}_{j_1}, \ldots, \hat{q}_{j_k}\}$ equals $\underline{y}$. There exists a vector $\underline{c}$ in F^n having at most k nonzero components (viz., at positions $j_1, \ldots, j_k$) such that $\underline{c}Q = \underline{y}$. Let $\underline{d} = [\underline{d}_1 \mid \cdots \mid \underline{d}_k] = \underline{c}Q$. Since each $\underline{y}_i \neq 0$, each $\underline{d}_i \not\in \overline{W}_i$.

Finally, we demonstrate that the n columns of $\underline{c}Q\tilde{M} = \underline{d}\tilde{M}$ are linearly independent modulo $F^1 = F$. For any $u_{10}, u_{11}, \ldots, u_{1,n_1-1}, u_{20}, \ldots, u_{k,n_k-1}$ in F, if

$$\sum_{i=1}^{k} \underline{d}_i \sum_{j=0}^{n_i-1} u_{ij} P_i^j \underline{\alpha}^i$$

is in F, then it must be zero; this result implies that

$$\underline{d}_i \sum_{j=0}^{n_i-1} u_{ij} P_i^j = 0 \qquad \text{for } i = 1,\ldots,k$$

because all components of all $\underline{\alpha}^i$ are linearly independent. Since every $\underline{d}_i \not\in \overline{W}_i$, every $u_{ij} = 0$.

Ergo, by Lemma 2.9, computing $\underline{c}Q\tilde{M}\tilde{\underline{x}}$ requires at least n nonscalar multiplications. Hence, $k + t - n \geq n$, and $t \geq 2n - k$. □

In Winograd's proof [W5], the rows $\hat{q}_{j_1}, \ldots, \hat{q}_{j_k}$ are chosen according to a different criterion, essentially that the part of $\hat{q}_{j_\ell}$ in $\hat{Q}^\ell$ be nonzero; some of the j_ℓ may be the same. In this case the desired $\underline{c}$ vector might not exist. For instance, suppose $F = \mathbb{F}_2$, $k = 3$, $n = 6$, $n_1 = n_2 = n_3 = 2$,

$$Q = \left[\begin{array}{cc:cc:cc} 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{array}\right], \qquad \hat{Q} = \left[\begin{array}{cc:cc:cc} 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array}\right],$$

$j_1 = j_3 = 1$, and $j_2 = 2$; note that Q is nonsingular. Then $[1\ 1\ 0\ 0\ 0\ 0]\hat{Q} = [0\ 0 \mid 1\ 0 \mid 1\ 0]$, $[1\ 0\ 0\ 0\ 0\ 0]\hat{Q} = [1\ 0 \mid 0\ 0 \mid 1\ 0]$, and $[0\ 1\ 0\ 0\ 0\ 0]\hat{Q} = [1\ 0 \mid 0\ 0 \mid 0\ 0]$, none of which yields the vector $\underline{y} = [1\ 0 \mid 1\ 0 \mid 1\ 0]$.

3. REPRESENTATIONS AND ALGORITHMS

3A. Preliminaries

In this section we collect definitions and theorems from standard textbooks [C3] and [H1], assuming that the reader is familiar with algebra at the level of [J1]. As usual, all algebras are associative and finite-dimensional over a field.

Definition. A ring satisfies the minimum condition if it has no infinite descending chain of ideals.

Every finite-dimensional, associative algebra satisfies the minimum condition.

Definition. A ring is simple if it has no nontrivial two-sided ideals.

Theorem 3.1 (Wedderburn). Every simple ring with minimum condition is isomorphic to a full matrix ring over a division ring.

Definition. A ring is semisimple if it includes no nilpotent left ideals.

Theorem 3.2. Every semisimple ring with minimum condition is isomorphic to a direct sum of a finite number of simple subrings with minimum condition.

Theorem 3.3 (Maschke). The algebra $F[G]$ of a finite group G over a field F is semisimple if the characteristic of F does not divide the order of G: i.e., if char $F \nmid$ card G.

Definition. A *representation of an algebra* $\mathcal{A}$ over a field F is an algebra homomorphism from $\mathcal{A}$ into the endomorphisms (set of linear transformations) $\mathrm{End}_F(V)$ of a vector space V over F; the space V is the *representation space*. The *dimension* of a representation is the dimension of its representation space. If the representation space is $\mathcal{A}$, then the representation is *regular*.

Definition. A *matrix representation of dimension r* of an algebra $\mathcal{A}$ over a field F is an algebra homomorphism from $\mathcal{A}$ into a subalgebra of $F^{r\times r}$, the ring of $r\times r$ matrices over F. The *left regular matrix representation* ρ_L with respect to a basis $\{v_1, \ldots, v_n\}$ of $\mathcal{A}$ is an algebra homomorphism from $\mathcal{A}$ into $F^{n\times n}$, where $n = [\mathcal{A}:F]$; for any α in $\mathcal{A}$,

$$\alpha v_j = \rho_L(\alpha)_{1j}v_1 + \ldots + \rho_L(\alpha)_{nj}v_n, \quad j = 1,\ldots,n,$$

defines the i,j entry of $\rho_L(\alpha)_{ij}$ of $\rho_L(\alpha)$.

A matrix representation is a coordinatized version of a representation. Any representation T is defined by the values $T(v_i)$ for $i = 1,\ldots,n$ if $\{v_1, \ldots, v_n\}$ is a basis for the algebra. We shall use the word *representation* for both the homomorphism and its values.

Definition. A matrix representation of dimension r of a group G over a field F is a group homomorphism from G into the group of invertible matrices in $F^{r\times r}$.

Any matrix representation of a group G over a field F extends naturally to a matrix representation of the group algebra F[G] and, furthermore, to the algebra E[G] over any field E that includes F.

Example. With respect to the basis $\{1, \hat{\imath}, \hat{\jmath}, \hat{k}\}$ the left regular matrix representation of the real quaternions $\mathcal{Q}_{\mathbb{R}}$ is

$$\rho_L: \ a_0 + a_1\hat{\imath} + a_2\hat{\jmath} + a_3\hat{k} \mapsto \begin{bmatrix} a_0 & -a_1 & -a_2 & -a_3 \\ a_1 & a_0 & -a_3 & a_2 \\ a_2 & a_3 & a_0 & -a_1 \\ a_3 & -a_2 & a_1 & a_0 \end{bmatrix}.$$

Example. One can define the regular matrix representation of the polynomial algebra $F[z]/(\pi(z))$ in terms of the companion matrix P for $\pi(z) = p_0 + p_1 z + \dots + p_{n-1}z^{n-1} + z^n$:

$$P = \begin{bmatrix} 0 & & & 0 & -p_0 \\ 1 & 0 & & & -p_1 \\ & 1 & \ddots & & \vdots \\ & & \ddots & 0 & -p_{n-2} \\ 0 & & & 1 & -p_{n-1} \end{bmatrix}.$$

The left regular representation with respect to the basis $\{1, z, \dots, z^{n-1}\}$ is

$$\rho_L\colon\ a_0 + a_1 z + \ldots + a_{n-1} z^{n-1} \mapsto a_0 I_n + a_1 P + \ldots + a_{n-1} P^{n-1},$$

where I_n is the n×n identity matrix.

Example. Let $\{g_1, \ldots, g_n\}$ be an ordering of a finite group G. The i,j entry of the left regular matrix representation of the element $\alpha = \sum_{i=1}^{n} a_i g_i$ of F[G] is $\rho_L(\alpha)_{ij} = a_k$ if, and only if, $g_i = g_k g_j$.

Definition. Two matrix representations T and U of A of dimension r are equivalent if there exists a fixed invertible matrix S in $F^{r\times r}$ such that $T(\alpha)S = SU(\alpha)$ for all α in A.

Definition. A representation T of an algebra A with nonzero representation space V is reducible if there exists a nontrivial proper invariant subspace W of V satisfying

$$T(\alpha)w \in W \qquad\qquad \text{for all } \alpha \in A, \text{ all } w \in W;$$

otherwise, the representation is irreducible. The representation is completely reducible if for every such invariant subspace W of V there exists an invariant subspace $\bar{W}$ such that $V \cong W \oplus \bar{W}$.

Theorem 3.4. A finite-dimensional, associative algebra is semisimple if, and only if, every reducible representation is completely reducible.

For any algebra A and representation T of A with representation space V one can make V into a (left) A-module by defining scalar multiplication by

$$\alpha v = T(\alpha)v \qquad\qquad \text{for all } \alpha \text{ in } A\text{, all } v \text{ in } V.$$

Call this A-module irreducible if T is irreducible.

Theorem 3.5. If A is a semisimple algebra, then every irreducible A-module is isomorphic to some minimal left ideal of A (i.e., a left ideal that properly contains no nontrivial left ideals).

Theorem 3.6 (Schur's Lemma). If A is a finite-dimensional, associative, semisimple algebra, then the endomorphisms of any irreducible left A-module form a division algebra; if the base field F is algebraically closed and X is an irreducible left A-module, then $\mathrm{End}_A(X) \cong F$. The simple subalgebra of A including a minimal left ideal J of A is isomorphic to a matrix ring over this division algebra defined by J.

Definition. A field F is a splitting field for the finite group G if every irreducible representation of G over F extends to an irreducible representation of E[G] over E for every extension field E of F.

The splitting field of a group is related to the splitting field of a polynomial. A field F is a splitting field for the cyclic group Z_n of order n if, and only if, it is a splitting field for the polynomial $z^n - 1$ because $F[Z_n] \cong F[z]/(z^n - 1)$.

Corollary. If G is a finite group, then any algebraically closed field is a splitting field for G.

3B. The General Element and its Characteristic Polynomial

From the theorems in the last section one can deduce that a (finite-dimensional, associative) semisimple algebra decomposes into a direct sum of full matrix rings over certain division algebras. We define the general element ξ of an algebra and study the characteristic polynomial of the regular representation of ξ. The factorization of this polynomial yields information on the sizes of these matrix rings and division algebras.

Definition. For a semisimple algebra A over a field F the general element of A is $\xi = x_1v_1 + x_2v_2 + \ldots + x_nv_n$, where $\{v_1, \ldots, v_n\}$ is a basis for A and $x_1, \ldots, x_n$ are distinct indeterminates over F.

The left regular matrix representation of the general element ξ is

$$\rho_L(\xi) = x_1\rho_L(v_1) + \ldots + x_n\rho_L(v_n),$$

which has entries in $F[x_1,\ldots,x_n]$. Let $\chi(\lambda)$ denote the characteristic polynomial of $\rho_L(\xi)$: i.e., $\chi(\lambda) = \det\ (\lambda I_n - \rho_L(\xi))$, where I_n is the $n{\times}n$ identity matrix.

We first establish a lemma crucial to the proof of our main result. If $\psi(x_1,\ldots,x_n)$ is a polynomial in $F[x_1,\ldots,x_n]$, then write $\psi(u_1,\ldots,u_n)$ for the value in F resulting from the substitution of u_i for x_i, $i = 1,\ldots,n$, in ψ. Recall that the degree of a multivariate polynomial is the maximum of the degrees of its terms:

$$\deg \psi(x_1,\ldots,x_n) = \max \{\nu_1 + \ldots + \nu_n \mid x_1^{\nu_1}\ldots x_n^{\nu_n} \text{ appears in } \psi\}.$$

For instance, $\deg \chi(\lambda) = n$.

Lemma 3.7. Let F be a field of cardinality $f \geq n$ (f may be infinite). Let $x_1, \ldots, x_n$ be indeterminates and let $R = F[x_1,\ldots,x_n]$. Let $\phi_1(\lambda,x_1,\ldots,x_n)$ and $\phi_2(\lambda,x_1,\ldots,x_n)$ be monic polynomials in $R[\lambda]$ of degree less than n in $F[\lambda,x_1,\ldots, x_n]$. If for all $u_1, \ldots, u_n$ in F the irreducible factors of $\phi_1(\lambda,u_1,\ldots,u_n)$ divide $\phi_2(\lambda,u_1,\ldots,u_n)$ in $F[\lambda]$, then all irreducible factors of $\phi_1(\lambda,x_1,\ldots,x_n)$ divide $\phi_2(\lambda,x_1,\ldots,x_n)$ in $R[\lambda]$.

Proof. It suffices to demonstrate that if $\phi_1(\lambda,u_1,\ldots,u_n) \mid \phi_2(\lambda,u_1,\ldots,u_n)$ in $F[\lambda]$ for all elements $u_1, \ldots, u_n$ in F, then $\phi_1 \mid \phi_2$ in $R[\lambda]$.

Since ϕ_1 and ϕ_2 are monic in $R[\lambda]$, one can use the division algorithm for polynomials to write

$$\phi_2(\lambda,x_1,\ldots,x_n) = \phi_1(\lambda,x_1,\ldots,x_n)q(\lambda,x_1,\ldots,x_n) + r(\lambda,x_1,\ldots,x_n).$$

Let $r(\lambda,x_1,\ldots,x_n) = \psi_d(x_1,\ldots,x_n)\lambda^d + \ldots + \psi_0(x_1,\ldots,x_n)$ and $\delta = \deg r$; since $\deg \phi_2 < n$, $\delta < n$. By hypothesis, $\psi_j(u_1,\ldots,u_n) = 0$ for each j and every selection of $u_1, \ldots, u_n$ from F.

Case: F infinite. This property of each ψ_j implies that $\psi_j = 0$ for all j. Therefore, $r = 0$ and $\phi_1 \mid \phi_2$ in $R[\lambda]$.

Case: F finite. Suppose $\psi_j \neq 0$. If for some $u_2, \ldots, u_n$ in F we have $\psi_j(x_1,u_2,\ldots,u_n) \neq 0$, then the polynomial

$$\prod_{u \in F} (x_1 - u) = x_1^f - x_1$$

divides $\psi_j(x_1,u_2,\ldots,u_n)$ in $F[x_1]$ because $\psi_j(u,u_2,\ldots,u_n) = 0$ for every u in F; but since the highest power of x_1 in ψ_j is at most δ and $f \geq n > \delta$, we have a contradiction. Therefore, for all selections of $u_2, \ldots, u_n$ in F we have $\psi_j(x_1,u_2,\ldots,u_n) = 0$. Now consider ψ_j as a polynomial in x_1: $\psi_j = \sum_k \psi_{jk}(x_2,\ldots,x_n)x_1^k$, where each $\psi_{jk} \in F[x_2,\ldots,x_n]$. We know that $\psi_{jk}(u_2,\ldots,u_n) = 0$ for every selection of $u_2, \ldots, u_n$ in F, and deg $\psi_{jk} \leq \delta < f$ for every j,k. By induction every $\psi_{jk} = 0$, and hence $\psi_j = 0$. Since $\psi_j = 0$ for every j, r = 0, and $\phi_1 | \phi_2$. □

The hypothesis on the cardinality of the field F cannot be weakened. For example, let $\phi_1(\lambda,x_1,x_2,x_3,x_4) = \lambda + (x_1 + x_2 + x_3 + x_4)$, $\phi_2(\lambda,x_1,x_2,x_3,x_4) = \lambda^2(\lambda + (x_1 + x_2 + x_3 + x_4)) + x_1^2x_2 + x_2^2x_1 + x_3^2x_4 + x_4^2x_3$, and $F = \mathbb{F}_2$, the field of two elements. Then $\phi_1(\lambda,u_1,u_2,u_3,u_4) | \phi_2(\lambda,u_1,u_2,u_3,u_4)$ for every selection u_1, u_2, u_3, u_4 from F, but $\phi_1(\lambda,x_1,x_2,x_3,x_4) \nmid \phi_2(\lambda,x_1,x_2,x_3,x_4)$ in $F[\lambda,x_1,x_2,x_3,x_4]$. Thus, the ideas expressed by Lemma 3.7 seem inadequate for establishing our main result for arbitrary fields.

Because $F[\lambda,x_1,\ldots,x_n]$ is a unique factorization domain, $\chi(\lambda)$ has a unique prime factorization into powers of irreducible factors.

Theorem 3.8. Let $\chi(\lambda)$ be the characteristic polynomial of the regular representation of the general element ξ of a semisimple algebra A of dimension n over a field $\mathbf{F}$ of cardinality card $F \geq n$. There is a one-one correspondence between the irreducible factors of $\chi(\lambda)$ and the inequivalent, irreducible representations of A over F; for each irreducible representation T the characteristic polynomial of $T(\xi)$ is a power of an irreducible factor of $\chi(\lambda)$.

Proof. Part 1. All inequivalent, irreducible representations appear (with some multiplicity) in the regular representation (Theorem 3.5).

Part 2. If representation T is reducible, then clearly the characteristic polynomial of $T(\xi)$ is reducible.

Part 3. We must show that if T is a representation of dimension $r \leq n$ such that the characteristic polynomial $\phi(\lambda,x_1,\ldots,x_n)$ of $T(\xi)$ has nontrivial distinct irreducible factors in $F[\lambda,x_1,\ldots,x_n]$, then T is reducible.

Let $\phi(\lambda,x_1,\ldots,x_n) = \phi_1(\lambda,x_1,\ldots,x_n)\phi_2(\lambda,x_1,\ldots,x_n)$, where ϕ_1 and ϕ_2 have no irreducible factors in common; equivalently, the greatest common divisor of ϕ_1 and ϕ_2 is a scalar in F. Assume $\phi_1(\lambda)$ and $\phi_2(\lambda)$ are monic and neither is a scalar. Since the dimension of T is at most n, deg $\phi \leq n$, and so deg $\phi_1 < n$ and deg $\phi_2 < n$.

Let $\{v_1, \ldots, v_n\}$ be a basis for A. Define $W = \{w \in F^r |$ for all $u = u_1v_1 + \ldots u_nv_n$ in A, $\phi_1(T(u),u_1,\ldots,u_n)w = 0\}$. It suffices to show that W is a nontrivial proper subspace of F^r invariant under T. If $w \in W$ and

$u = u_1v_1 + \ldots + u_nv_n \in A$, then $T(u)w \in W$ because $\phi_1(T(u),u_1,\ldots,u_n)T(u)w = T(u)\phi_1(T(u),u_1,\ldots,u_n)w = 0$. Thus, W is invariant under every $T(u)$.

Claim: $W \neq F^r$. Suppose, to the contrary, $W = F^r$. Then $\phi_1(T(u),u_1,\ldots,u_n) = 0$ for every $u = u_1v_1 + \ldots + u_nv_n$ in A, and the minimum polynomial of $T(u)$ divides $\phi_1(\lambda,u_1,\ldots,u_n)$. Since the minimum polynomial and the characteristic polynomial have the same irreducible factors, all the irreducible factors of $\phi_2(\lambda,u_1,\ldots,u_n)$ are factors of $\phi_1(\lambda,u_1,\ldots,u_n)$ for every selection of $u_1, \ldots, u_n$ from F. By Lemma 3.7, the irreducible factors of $\phi_2(\lambda,x_1,\ldots,x_n)$ divide $\phi_1(\lambda,x_1,\ldots,x_n)$. Contradiction.

Claim: $W \neq 0$. Suppose $W = 0$. Then for no nonzero w in F^r does $\phi_1(T(\xi),x_1,\ldots,x_n)w = 0$. Therefore, $\phi_1(T(\xi),x_1,\ldots,x_n) \neq 0$. Since $\phi(T(\xi),x_1,\ldots,x_n) = 0$ and $\phi(\lambda,x_1,\ldots,x_n) = \phi_1(\lambda,x_1,\ldots,x_n) \cdot \phi_2(\lambda,x_1,\ldots,x_n)$, we must have $\phi_2(T(\xi),x_1,\ldots,x_n) = 0$. For every $u = u_1v_1 + \ldots + u_nv_n$ in A, $\phi_2(T(u),u_1,\ldots,u_n) = 0$, and the minimum polynomial of $T(u)$ divides $\phi_2(\lambda,u_1,\ldots,u_n)$. As in the last paragraph we derive a contradiction.

Ergo, W is a nontrivial proper subspace of F^r invariant under $T(u)$ for all $u \in A$, and T is reducible. □

Let $\{J_1, \ldots, J_m\}$ be a complete set of nonisomorphic minimal left ideals of the semisimple algebra A over field F whose cardinality is at least as large as the dimension of A. Each J_i is isomorphic to the left A-module engendered by an irreducible representation. For i = 1, ...,m define the division algebra $\mathcal{D}_i = \mathrm{End}_A(J_i)$ and let $\mu_i = [J_i : \mathcal{D}_i]$.

From [C3] we know that $\mathcal{A} \cong \mathcal{A}_1 \oplus \ldots \oplus \mathcal{A}_m$, where each $\mathcal{A}_i$ is a full $\mu_i \times \mu_i$ matrix ring over $\mathcal{D}_i$, and $n = [\mathcal{A}:F] = \sum_{i=1}^{m} \mu_i^2 [\mathcal{D}_i:F]$. Each simple algebra $\mathcal{A}_i$ is a direct sum of μ_i copies of J_i. Let $\chi(\lambda)$ be the characteristic polynomial of the regular representation of the general element of $\mathcal{A}$, with prime factorization $\chi(\lambda) = \chi_1(\lambda)^{\nu_1} \cdots \chi_k(\lambda)^{\nu_k}$. Let $\deg \chi_i = n_i$. Then $k = m$ and we can reorder the χ_i or J_i so that $n_i \nu_i = \mu_i^2 [\mathcal{D}_i:F]$ for every i.

For each $i = 1,\ldots,m$, because $\mathcal{A}_i$ is a direct sum of μ_i copies of J_i, the ith representation occurs μ_i times in the regular representation, and therefore, $\mu_i | \nu_i$. Thus, in some cases we can calculate each μ_i and $[\mathcal{D}_i:F]$ from n_i and ν_i. For instance, if $n_4 = 3$ and $\nu_4 = 2$, then $\mu_4 = 1$ and $[\mathcal{D}_4:F] = 6$.

Example. Since the real quaternions $\mathcal{Q}_{\mathbb{R}}$ form a division algebra, $\mathcal{Q}_{\mathbb{R}}$ has one simple component (itself); $\mathcal{Q}_{\mathbb{R}}$ is isomorphic to a 1×1 matrix ring over itself. **The characteristic polynomial of the regular representation** of the general element $x_0 + x_1\hat{\imath} + x_2\hat{\jmath} + x_3\hat{k}$ of $\mathcal{Q}_{\mathbb{R}}$ is $(\lambda^2 - 2x_0\lambda + (x_0^2 + x_1^2 + x_2^2 + x_3^2))^2$. In this case $n = 4$, $k = 1$, $n_1 = 2$, $\nu_1 = 2$, $\mu_1 = 1$, and $[\mathcal{D}_1:\mathbb{R}] = 4$.

Intuitively, the general element epitomizes all elements of an algebra; the characteristic polynomial of the regular representation of the general element summarizes the properties of the representations of the algebra.

3C. An Algorithm for Products in Semisimple Algebras

We present a class of algorithms for multiplying elements of a semisimple algebra A over a field F. As in the last section, let A decompose into a direct sum of m simple algebras A_i, $i = 1,\dots,m$; each A_i is a full $\mu_i \times \mu_i$ matrix ring over a unique division algebra $\mathcal{D}_i$. Let $[A:F] = n$ and let $\{v_1, \dots, v_n\}$ be a basis for A. Let $\{T_1, \dots, T_m\}$ be the inequivalent, irreducible matrix representations of A over F, with T_i corresponding to A_i for each i.

Like Winograd [W3], we can cast the problem of multiplication in A in terms of a matrix-vector product. If $\alpha = a_1v_1 + \dots + a_nv_n$, $\beta = b_1v_1 + \dots + b_nv_n$, and $\alpha\beta = c_1v_1 + \dots + c_nv_n$, then

$$\begin{bmatrix} c_1 \\ \vdots \\ c_n \end{bmatrix} = \rho_L(\alpha)\begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix}. \qquad (3.1)$$

Changing the basis of the algebra changes the matrix and vectors in (3.1) by only additions and scalar multiplications. Thus, a change of basis yields a computationally equivalent matrix-vector problem: a program for the new problem can be modified to produce a program for the original problem with the same number of nonscalar multiplications, and vice versa. The caracteristic polynomial of the regular representation of the general element of A remains invariant under change of basis.

With a particular change of basis we can block-diagonalize the regular matrix representation so that each matrix representation T_i occurs μ_i times along the diagonal. Let $n_i = [A_i:F] = \mu_i^2[\mathcal{D}_i:F]$, and let

$\{v'_{11}, \dots, v'_{1n_1}, v'_{21}, \dots, v'_{mn_m}\}$ be this new basis, where for each i $\{v'_{i1}, \dots, v'_{in_i}\}$ is a basis for A_i. Let $\rho'_L(\alpha)$ be the left regular matrix representation of α with respect to the new basis. Let $\beta = b'_{11}v'_{11} + \dots + b'_{mn_m}v'_{mn_m}$. Computing (3.1) is equivalent to computing

$$\rho'_L(\alpha)\begin{bmatrix} b'_{11} \\ \vdots \\ b'_{mn_m} \end{bmatrix} = \tag{3.2}$$

$$\begin{bmatrix} \overbrace{T_1(\alpha) \ddots T_1(\alpha)}^{\mu_1} & & & & 0 \\ & T_2(\alpha) & & & \\ & & \ddots & & \\ & & & \overbrace{T_m(\alpha) \ddots T_m(\alpha)}^{\mu_m} & \\ 0 & & & & \end{bmatrix} \begin{bmatrix} b'_{11} \\ \vdots \\ b'_{1n_1} \\ b'_{21} \\ \vdots \\ b'_{m1} \\ \vdots \\ b'_{mn_m} \end{bmatrix}.$$

Each T_i has dimension $\mu_i[\mathcal{D}_i:F] = n_i/\mu_i$. For each i we have

$$\begin{bmatrix} T_i(\alpha) & & 0 \\ & \ddots & \\ 0 & & T_i(\alpha) \end{bmatrix} \begin{bmatrix} b'_{i1} \\ \vdots \\ b'_{in_i} \end{bmatrix}; \tag{3.3}$$

this matrix-vector product involves μ_i smaller products of T_i by a n_i/μ_i-vector. The ith product, (3.3), can be calculated with $\mu_i(n_i/\mu_i)^2 = \mu_i{}^3[\mathcal{D}_i:F]^2$ nonscalar multiplications. Therefore, we can calculate (3.2) with $\sum_{i=1}^{m} \mu_i{}^3[\mathcal{D}_i:F]^2$ nonscalar multiplications.

Example. In the quotient ring $\mathbb{Q}[z]/(z^2 + z - 6) \cong \mathbb{Q}[z]/(z + 3) \oplus \mathbb{Q}[z]/(z - 2)$ the computation of $c_0 + c_1z = (a_0 + a_1z)(b_0 + b_1z) \bmod (z^2 + z - 6)$ can be expressed in two equivalent ways:

$$\begin{bmatrix} c_0 \\ c_1 \end{bmatrix} = \begin{bmatrix} a_0 & 6a_1 \\ a_1 & a_0-a_1 \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \end{bmatrix}, \quad \begin{bmatrix} c_0+2c_1 \\ c_0-3c_1 \end{bmatrix} = \begin{bmatrix} a_0+2a_1 & 0 \\ 0 & a_0-3a_1 \end{bmatrix} \begin{bmatrix} b_0+2b_1 \\ b_0-3b_1 \end{bmatrix}.$$

The first matrix is the left regular matrix representation of $a_0 + a_1z$ with respect to the basis $\{1, z\}$, the second with respect to the basis $\{\frac{3}{5} + \frac{z}{5}, \frac{2}{5} - \frac{z}{5}\}$.

We now present an algorithm that also uses $\sum_{i=1}^{m} \mu_i^3 [\mathcal{D}_i : F]^2$ nonscalar multiplications in the worst case, but fewer in several important cases. Let η_i be the natural epimorphism of A onto A_i. Because $A \cong A_1 \oplus \ldots \oplus A_m$, there is a one-one correspondence between elements of A and m-tuples of their projections:

$$\alpha \leftrightarrow (\eta_1(\alpha), \ldots, \eta_m(\alpha)) \qquad \text{for all } \alpha \in A.$$

Each $\eta_i(\alpha)$ is a $\mu_i \times \mu_i$ matrix with entries in $\mathcal{D}_i$.

Algorithm 3.9. To compute the coefficients $\{c_1, \ldots, c_n\}$ of the product $\gamma = \alpha\beta = c_1v_1 + \ldots + c_nv_n$ of $\alpha = a_1v_1 + \ldots + a_nv_n$ and $\beta = b_1v_1 + \ldots + b_nv_n$ in A from $\{a_1, \ldots, a_n, b_1, \ldots, b_n\}$.

Step 1. For each simple algebra A_i calculate the projections $A_i = \eta_i(\alpha)$ of α and $B_i = \eta_i(\beta)$ of β in A_i. Each A_i and B_i is a $\mu_i \times \mu_i$ matrix with entries in $\mathcal{D}_i$.

Step 2. Multiply corresponding matrices. Let $C_i = A_iB_i$ for $i = 1,\ldots,m$.

Step 3. Find the element γ of $\mathcal{A}$ such that the projection of γ in $\mathcal{A}_i$ is $C_i = \eta_i(\gamma)$ for $i = 1,\ldots,m$. □

In essence, we multiply the projections $\{A_1, \ldots, A_m\}$ and $\{B_1, \ldots, B_m\}$ of α and β in the simple algebras $\mathcal{A}_1, \ldots, \mathcal{A}_m$. The result γ is uniquely determined by $\{C_1, \ldots, C_m\}$.

The algorithm employs a generalization of the discrete Fourier transform, which maps convolutional multiplication of two vectors into pointwise multiplication of the transformed vectors. For general semisimple algebras we "transform" an element into a set of matrices (its projections), not all of which may be 1×1. The "transform" of the product of two elements equals the set of products of corresponding matrices.

To perform Step 2 we multiply $\mu_i \times \mu_i$ matrices with entries in $\mathcal{D}_i$. The matrix multiplication in $\mathcal{A}_i$ requires at most μ_i^3 products of elements of $\mathcal{D}_i$. Each of these products in $\mathcal{D}_i$ requires at most $[\mathcal{D}_i:F]^2$ nonscalar multiplications over F. Because Steps 1 and 3 comprise only additions and scalar multiplications, Algorithm 3.9 uses at most $\sum_{i=1}^m \mu_i^3[\mathcal{D}_i:F]^2$ nonscalar multiplications.

Alternatively, one could compute $\{c_1, \ldots, c_n\}$ by calculating $\{a_ib_j \mid 1 \le i \le n,\ 1 \le j \le n\}$ and taking linear combinations of these products. This naive method uses n^2 nonscalar multiplications. Since $n^2 = (\sum_{i=1}^m \mu_i^2[\mathcal{D}_i:F])^2 \ge \sum_{i=1}^m \mu_i^3[\mathcal{D}_i:F]^2$, Algorithm 3.9 is superior.

We consider three special cases.

1. The divison algebra $\mathcal{D}_i$ is commutative. For example, if the base field F is finite, then each divison algebra $\mathcal{D}_i$ is finite; thus, by a celebrated theorem of Wedderburn (e.g., [J1, p. 431]), each $\mathcal{D}_i$ must be commutative. One can use Winograd's algorithm [W2] to calculate the matrix product A_iB_i with only $2\mu_i^2 + (\mu_i^2 - 2\mu_i)\lfloor\frac{1}{2}(\mu_i + 1)\rfloor$ products of elements of $\mathcal{D}_i$ because $\mathcal{D}_i$ is commutative. Moreover, the commutativity of $\mathcal{D}_i$ implies that it is a finite-dimensional extension field of F. If $\mathcal{D}_i$ is separable over F (e.g., if F is perfect), then $\mathcal{D}_i$ has a primitive element: $\mathcal{D}_i \cong F[z]/(\pi_i(z))$ for some irreducible polynomial π_i of degree $[\mathcal{D}_i:F]$. If F has at least $2[\mathcal{D}_i:F]-1$ distinct elements, then Algorithm 2.5 may be used to multiply elements of $\mathcal{D}_i$ with only $2[\mathcal{D}_i:F]-1$ nonscalar multiplications, meeting the lower bound of Theorem 2.4.

2. All $[\mathcal{D}_i:F] = 1$. In this case the projections A_i and B_i in Algorithm 3.9 coincide with the irreducible, inequivalent matrix representations $\{T_1, \ldots, T_m\}$ of A over F: $A_i = T_i(\alpha)$ and $B_i = T_i(\beta)$ for $i = 1,\ldots,m$. Let $\rho_L(\alpha) \mapsto S\rho_L(\alpha)S^{-1}$ be the similarity transformation (change of basis) that block-diagonalizes the left regular matrix representation--cf. (3.2). Calculating all products A_iB_i simultaneously is equivalent to multiplying $(S\rho_L(\alpha)S^{-1})(S\rho_L(\beta)S^{-1})$.

According to Schur's Lemma (Theorem 3.6), all $[\mathcal{D}_i:F] = 1$ when F is algebraically closed. If A is a group algebra F[G] and F is a splitting field for G, then all $[\mathcal{D}_i:F] = 1$. In this latter case, to multiply $(\sum_{g\in G} a_g g)(\sum_{g\in G} b_g g) = (\sum_{g\in G} c_g g)$, we calculate

$$A_i = T_i(\alpha) = \sum_{g\in G} a_g T_i(g), \quad B_i = T_i(\beta) = \sum_{g\in G} b_g T_i(g),$$

for $i = 1,\ldots,m$ in Step 1 of Algorithm 3.9. Modifying the inverse transform of equation (1.1), we obtain

$$c_g = \sum_{i=1}^{m} \sum_{j=1}^{\mu_i} \sum_{k=1}^{\mu_i} \frac{\mu_i}{n}(A_iB_i)_{jk}T_i(g^{-1})_{kj},$$

where $(A_iB_i)_{jk}$ is the j,k entry of A_iB_i and $T_i(g^{-1})_{kj}$ is the k,j entry of $T_i(g^{-1})$.

3. The simple algebra $\mathcal{A}_i$ is commutative. Because $\mathcal{A}_i$ is a full ring of $\mu_i \times \mu_i$ matrices, $\mu_i = 1$; otherwise, if $\mu_i > 1$, we could find a pair of matrices that do not commute. Then $\mathcal{A}_i$ is isomorphic to the division algebra $\mathcal{D}_i$, and the discussion for a commutative $\mathcal{D}_i$ (above) applies. Fiduccia and Zalcstein [F2] studied commutative semisimple algebras over perfect fields, in which the simple component algebras $\mathcal{A}_i$ are separable extension fields (Theorem 2.7).

In general, we might need fewer than $\mu_i^3[\mathcal{D}_i:F]$ nonscalar multiplications to multiply the matrices A_iB_i in Step 2 of Algorithm 3.9. Multiplying a pair of $\mu\times\mu$ matrices over any ring requires at most $\mathcal{O}(\mu^{\log_2 7})$ nonscalar multiplications (over that ring) [A1]. Fischer [F3] has demonstrated that at most $3.912\mu^{\log_2 7}$ total arithmetic operations are required if $\mu > 13$. Let $\mathcal{M}(\mu,R)$ be the number of nonscalar multiplications required to multiply two $\mu\times\mu$ matrices with entries in a ring R. For every ring R, $\mathcal{M}(1,R) = 1$, $\mathcal{M}(2^k,R) \leq 7^k$ for any positive integer k [S1], $\mathcal{M}(2,R) \geq 7$ [W4], and $\mathcal{M}(3,R) \leq 23$ [L1]. (Actually, Winograd [W4] showed that over $\mathbb{Q}$ multiplying a pair of 2×2 matrices requires at least 7 nonscalar multiplications, but his proof holds for any commutative ring;

since any program for multiplying matrices with entries in any (possibly noncommutative) ring can be used if the ring is commutative, we conclude that $\mathcal{M}(2,R) = 7$ for every R.)

Fiduccia [private communication] speculates that the multiplicative complexity of any division algebra is linear in its dimension. Fast algorithms for multiplying elements of general noncommutative division algebras have not yet been discovered.

Conjecture. The multiplicative complexity of $\mathcal{A}$ over F is at least $\sum_{i=1}^{m} \mathcal{M}(\mu_i,\mathcal{D}_i)(2[\mathcal{D}_i:F] - 1)$.

Proving this conjecture seems to require two results: (1) the multiplicative complexity of $\mathcal{A}$ equals the sum of the complexities of the A_i; (2) the multiplicative complexity of each A_i is at least $\mathcal{M}(\mu_i,\mathcal{D}_i)(2[\mathcal{D}_i:F] - 1)$. Part (1) is related to the "Direct Sum Conjecture" of [F2] and [W5]: that the multiplicative complexity of

$$\begin{bmatrix} M_1(\underline{\zeta}^{(1)}) & 0 \\ 0 & M_2(\underline{\zeta}^{(2)}) \end{bmatrix} \begin{bmatrix} \underline{\tau}^{(1)} \\ \underline{\tau}^{(2)} \end{bmatrix},$$

where components of $\underline{\zeta}^{(1)}$, $\underline{\zeta}^{(2)}$, $\underline{\tau}^{(1)}$, and $\underline{\tau}^{(2)}$ are disjoint sets of indeterminates, equals the sum of the complexities of the matrix-vector products $M_1(\underline{\zeta}^{(1)})\underline{\tau}^{(1)}$ and $M_2(\underline{\zeta}^{(2)})\underline{\tau}^{(2)}$. Instead of disjoint sets of indeterminates, however, we have linearly independent linear combinations of indeterminates. Part (2) appears to be a generalization of Theorem 2.4.

In summary, Algorithm 3.9 enables us to compute the product of two elements of $\mathcal{A}$ with at most $\sum_{i=1}^{m} \mu_i^{3}[\mathcal{D}_i:F]^2$ nonscalar multiplications. The factor μ_i^{3} can be replaced by the number of nonscalar multiplications with which a pair of $\mu_i \times \mu_i$ matrices with entries in $\mathcal{D}_i$ can be multiplied. For the factor $[\mathcal{D}_i:F]^2$ one may substitute the number of nonscalar multiplications with which products in $\mathcal{D}_i$ can be calculated; if $\mathcal{D}_i$ is commutative and F is a perfect field with at least $2[\mathcal{D}_i:F]-1$ distinct elements, then the factor $[\mathcal{D}_i:F]^2$ can be replaced by $2[\mathcal{D}_i:F]-1$.

4. EXAMPLES

4A. Polynomial Algebras

Consider the quotient polynomial algebra $F[z]/(\pi(z))$ in which the irreducible factors of $\pi(z) = p_0 + p_1 z + \ldots + p_{n-1} z^{n-1} + z^n$ all have multiplicity 1: $\pi(z)$ has prime factorization $\pi_1(z)\cdots\pi_k(z)$, where π_1, ..., π_k are distinct irreducible polynomials in $F[z]$. This algebra is semisimple:

$$F[z]/(\pi(z)) \cong F[z]/(\pi_1(z)) \oplus \ldots \oplus F[z]/(\pi_k(z));$$

each simple component algebra $F[z]/(\pi_i(z))$ is a finite-dimensional extension field of F.

For semisimple algebras Algorithm 2.5 is a special case of Algorithm 3.9. The product $(\alpha(z) \bmod \pi_i(z))(\beta(z) \bmod \pi_i(z))$ is merely the product of the projections of α and β in the simple algebra $F[z]/(\pi_i(z))$. The algorithms of Moenck and Borodin [M1] [B2] can be used to perform for each π_i the interpolation and multiple evaluation in Steps 3 and 4 of Algorithm 2.5 with $O(n_i(\log n_i)^2)$ additions and scalar multiplications. Furthermore, one can employ these algorithms to calculate rapidly the residues $\alpha(z) \bmod \pi_i(z)$ and $\beta(z) \bmod \pi_i(z)$ as well as the reconstruction of $\alpha(z)\beta(z)$ from the $\alpha(z)\beta(z) \bmod \pi_i(z)$ $(i = 1,\ldots,k)$.

In Algorithm 2.5 modular representation and Chinese remaindering are used in two different ways. Steps 2 and 5 perform a transform $\hat{\alpha}_i(z) = \alpha(z) \bmod \pi_i(z)$ and its inverse; the polynomials $\pi_1(z)$, ..., $\pi_k(z)$ are uniquely determined by $\pi(z)$. This transform corresponds to the transform of Algorithm 3.9: each $\hat{\alpha}_i$ is the projection of α in the simple algebra $F[z]/(\pi_i(z))$. In contrast, Steps 3 and 4 of Algorithm 2.5 multiply the transformed polynomials $\hat{\alpha}_i$ and $\hat{\beta}_i$ in each $F[z]/(\pi_i(z))$. The elements u_1, ..., u_{2n_i-1} are chosen freely; for instance, if F contains a primitive $(2n_i-1)$th root of unity ω and we select $u_j = \omega^j$ for $j = 1,\ldots,2n_i-1$, then Steps 3 and 4 may be performed with a fast Fourier transform.

When $\pi(z) = z^n - 1$, the algebra $F[z]/(\pi(z))$ is isomorphic to $F[Z_n]$, the algebra of the cyclic group Z_n over F. According to Maschke's Theorem (Theorem 3.3), the algebra $F[Z_n]$ is semisimple if char $F \nmid n$. The irreducible factors of $z^n - 1$ are the well-known cyclotomic polynomials when $F = \mathbb{Q}$ [J1]; the number of irreducible factors of $z^n - 1$ in $\mathbb{Q}[z]$ is $\Delta(n)$, the number of positive integral divisors of n. If $F = \mathbb{R}$, then $z^n - 1$ has only linear and quadratic irreducible factors: a total of $\lceil\frac{n+1}{2}\rceil$. For many n, $\lceil\frac{n+1}{2}\rceil \geq \Delta(n)$. By Theorems 2.6 and 2.10, the multiplicative complexity of $\mathbb{Q}[z]/(z^5 - 1)$ is $2\cdot 5 - \Delta(5) = 8$, whereas the multiplicative complexity of $\mathbb{R}[z]/(z^5 - 1)$ is $2\cdot 5 - \lceil\frac{5+1}{2}\rceil = 7$. For any positive integer n the multiplicative complexity of $\mathbb{C}[z]/(z^n - 1)$ is n because $z^n - 1$ splits into linear factors over $\mathbb{C}$. Clearly, the multiplicative complexity of a group algebra depends on the base field. We expect a minimum complexity for a group algebra if the base field is a splitting field for the group.

Example. The left regular matrix representation of the general element $\xi = x_0 + x_1z + x_2z^2$ of $F[z]/(z^3 - 1) \cong F[Z_3]$ with respect to the basis $\{1, z, z^2\}$ is

$$\rho_L(\xi) = \begin{bmatrix} x_0 & x_2 & x_1 \\ x_1 & x_0 & x_2 \\ x_2 & x_1 & x_0 \end{bmatrix}.$$

Suppose $F = \mathbb{Q}$. We know that $\mathbb{Q}[z]/(z^3 - 1) \cong \mathbb{Q}[z]/(z - 1) \oplus \mathbb{Q}[z]/(1 + z + z^2)$, so $\mathbb{Q}[z]/(z^3 - 1)$ has two irreducible, inequivalent representations. Appropriately changing the basis of $\mathbb{Q}[z]/(z^3 - 1)$ changes $\rho_L(\xi)$ to

$$\begin{bmatrix} x_0+x_1+x_2 & 0 & 0 \\ 0 & x_0-x_2 & x_2-x_1 \\ 0 & x_1-x_2 & x_0-x_1 \end{bmatrix}.$$

Matrix representations of dimensions 1 and 2 appear on the diagonal. We can interpret $\mathbb{Q}[z]/(1 + z + z^2)$ to be a ring of 1×1 matrices over itself, a commutative division algebra of dimension 2 over $\mathbb{Q}$. The multiplicative complexity of $\mathbb{Q}[Z_3]$ is $1 + 1^3(2 \cdot 2 - 1) = 4$.

If $F = \mathbb{C}$ instead of $\mathbb{Q}$, then $F[z]/(z^3 - 1)$ has three simple component algebras: $\mathbb{C}[z]/(z^3 - 1) \cong \mathbb{C}[z]/(z - 1) \oplus \mathbb{C}[z]/(z - \omega) \oplus \mathbb{C}[z]/(z - \omega^2)$, where ω is a primitive cube root of unity in $\mathbb{C}$. We can find a basis for $\mathbb{C}[z]/(z^3 - 1)$ so that $\rho_L(\xi)$ is

$$\begin{bmatrix} x_0+x_1+x_2 & 0 & 0 \\ 0 & x_0+\omega x_1+\omega^2 x_2 & 0 \\ 0 & 0 & x_0+\omega^2 x_1+\omega x_2 \end{bmatrix}.$$

When the base field F is perfect one can calculate the characteristic polynomial of the regular representation of the general element of $F[z]/(\pi(z))$ indirectly.

Theorem 4.1 (e.g., [C3; p. 208]). Let N be a n×n matrix with entries in a field F and let $\sigma_1, \ldots, \sigma_n$ be its eigenvalues. For any polynomial $\phi(z)$ in F[z] the eigenvalues of $\phi(N)$ are $\phi(\sigma_1), \ldots, \phi(\sigma_n)$.

Suppose F is perfect. Let P be the companion matrix for $\pi(z) = p_0 + p_1 z + \ldots + p_{n-1}z^{n-1} + z^n = \pi_1(z)\cdots\pi_k(z)$, where $\pi_1, \ldots, \pi_k$ are distinct irreducible polynomials:

$$P = \begin{bmatrix} 0 & & & 0 & -p_0 \\ 1 & 0 & & & -p_1 \\ & 1 & \ddots & & \vdots \\ & & \ddots & 0 & -p_{n-2} \\ 0 & & & 1 & -p_{n-1} \end{bmatrix};$$

since the characteristic polynomial of P is π, the eigenvalues of P are the roots of π. Let field E split π over F: in E[z] we have $\pi_i(z) = (z - \sigma_{i1})\cdots(z - \sigma_{in_i})$ for $i = 1,\ldots,k$, where $n_i = \deg \pi_i$; all roots σ_{ij} are distinct because F is perfect. By Theorem 4.1, the characteristic polynomial of $\rho_L(\xi) = x_0 + x_1 P + \ldots + x_{n-1}P^{n-1}$ is $\chi(\lambda) = \det(\lambda I_n - \rho_L(\xi)) = \chi_1(\lambda)\cdots\chi_k(\lambda)$, where

$$\chi_i(\lambda) = \prod_{j=1}^{n_i} [\lambda - (x_0 + x_1\sigma_{ij} + x_2\sigma_{ij}^2 + \ldots + x_{n-1}\sigma_{ij}^{n_i-1})].$$

For each i, the polynomial $\chi_i(\lambda)$, when written in $F[\lambda, x_0, \ldots, x_{n-1}]$, involves only symmetric polynomials in $\{\sigma_{i1}, \ldots, \sigma_{in_i}\}$. Because the coefficients of $\pi_i(z)$, with suitable changes of sign, are the values of the elementary symmetric polynomials in $\{\sigma_{i1}, \ldots, \sigma_{in_i}\}$, one can compute each $\chi_i(\lambda)$ from the coefficients of $\pi_i(z)$.

Example. Suppose the polynomial $\pi(z) = p_0 + p_1 z + p_2 z^2 + p_3 z^3$ is irreducible in $F[z]$. Let σ_1, σ_2, and σ_3 be the distinct roots of $\pi(z)$ in some splitting field E: $\pi(z) = (z - \sigma_1)(z - \sigma_2)(z - \sigma_3)$ in $E[z]$. The elementary symmetric polynomials in $\{\sigma_1, \sigma_2, \sigma_3\}$ are $\varepsilon_0 = \sigma_1\sigma_2\sigma_3$, $\varepsilon_1 = \sigma_1\sigma_2 + \sigma_2\sigma_3 + \sigma_3\sigma_1$, and $\varepsilon_2 = \sigma_1 + \sigma_2 + \sigma_3$; clearly, $\varepsilon_0 = -p_0$, $\varepsilon_1 = p_1$, and $\varepsilon_2 = -p_2$. The characteristic polynomial of the regular representation of the general element $x_0 + x_1 z + x_2 z^2$ is

$$\begin{aligned}\chi(\lambda) &= [\lambda-(x_0+x_1\sigma_1+x_2\sigma_1^2)][\lambda-(x_0+x_1\sigma_2+x_2\sigma_2^2)][\lambda-(x_0+x_1\sigma_3+x_2\sigma_3^2)] \\ &= \lambda^3 - [3x_0 + \varepsilon_2 x_1 + (\varepsilon_2^2 - 2\varepsilon_1)x_2]\lambda^2 \\ &\quad + [3x_0^2 + \varepsilon_1 x_1^2 + (\varepsilon_1^2 - 2\varepsilon_0\varepsilon_2)x_2^2 + 2\varepsilon_2 x_0 x_1 \\ &\qquad + 2(\varepsilon_2^2 - 2\varepsilon_1)x_0 x_2 + (\varepsilon_2\varepsilon_1 - 3\varepsilon_0)x_1 x_2]\lambda \\ &\quad - [x_0^3 + \varepsilon_2 x_0^2 x_1 + (\varepsilon_2^2 - 2\varepsilon_1)x_0^2 x_2 + \varepsilon_0 x_1^3 \\ &\qquad + \varepsilon_1 x_1^2 x_0 + \varepsilon_2\varepsilon_0 x_1^2 x_2 + \varepsilon_0^2 x_2^3 + (\varepsilon_1^2 - 2\varepsilon_0\varepsilon_2)x_2^2 x_0 \\ &\qquad + \varepsilon_1\varepsilon_0 x_2^2 x_1 + (\varepsilon_2\varepsilon_1 - 3\varepsilon_0)x_0 x_1 x_2].\end{aligned}$$

Thus, when F is perfect, the characteristic polynomial $\chi(\lambda)$ is the product of distinct linear factors in $E[\lambda, x_0, \ldots, x_{n-1}]$ and hence has no repeated factors in $F[\lambda, x_0, \ldots, x_{n-1}]$. Each factor χ_i, irreducible in $F[\lambda, x_0, \ldots, x_{n-1}]$, corresponds to a factor π_i of π, and hence also to the simple subalgebra $F[z]/(\pi_i(z))$, as guaranteed by Theorem 3.8.

4B. Abelian Group Algebras

We examine the semisimple algebra $F[G]$ of a finite abelian group G over a field F. Let $n = \text{card } G$. By Maschke's Theorem, $F[G]$ is semisimple if char $F \nmid n$.

First, we present a new proof of Chalkley's calculation [Cl] of the characteristic polynomial of the regular representation of the general element of $F[G]$ when F has a primitive nth root of unity. (Of course, if F has a primitive nth root of unity, then char $F \nmid n$ necessarily.) The proof hinges on the decomposition of G into a direct product of cyclic groups and on Ingraham's observation [Il] that the determinant is transitive.

Lemma 4.2 ([Il]). Let M be a $mr\times mr$ matrix with entries in field F. Suppose M can be partitioned into a $m\times m$ matrix of $r\times r$ matrices N_{ij} such that these $r\times r$ matrices commute. Let R be a commutative subring of $F^{r\times r}$ containing all these N_{ij}. Then

$$\det M = \det(\det_R M),$$

where $\det_R M$ is the determinant of M as a $m\times m$ matrix with entries in R. Moreover, if I_{mr} is the $mr\times mr$ identity matrix, then

$$\det(\lambda I_{mr} - M) = \det(\det_{R[\lambda]}(\lambda I_{mr} - M)).$$

Proof. [Il] or [Jl, pp. 407-408]. □

Lemma 4.3. Let $x_0, x_1, \ldots, x_{n-1}$ be indeterminates and let ω be a primitive nth root of unity. The characteristic polynomial of

$$M = \begin{bmatrix} x_0 & x_{n-1} & x_{n-2} & \cdots & x_1 \\ x_1 & x_0 & x_{n-1} & \cdots & x_2 \\ x_2 & x_1 & x_0 & \cdots & x_3 \\ \vdots & \vdots & \vdots & & \vdots \\ x_{n-1} & x_{n-2} & x_{n-3} & \cdots & x_0 \end{bmatrix}$$

is $\sum_{i=1}^{n} [\lambda - (x_0 + x_1\omega^i + x_2\omega^{2i} + \ldots + x_{n-1}\omega^{(n-1)i})]$.

Proof. Let P be the companion matrix for $z^n - 1$. The matrix M is the left regular matrix representation of $x_0 + x_1 z + \ldots + x_{n-1}z^{n-1}$ in $F[z]/(z^n - 1)$ with respect to the basis $\{1, z, \ldots, z^{n-1}\}$; i.e., $M = x_0 + x_1 P + \ldots + x_{n-1}P^{n-1}$. The lemma follows from Theorem 4.1. □

Theorem 4.4 ([C1]). Let F be a field with a primitive nth root of unity ω. Let G be a finite abelian group of order n isomorphic to a direct product of r cyclic groups $\langle h_1\rangle, \ldots, \langle h_r\rangle$ of orders $n_1, \ldots, n_r$. The characteristic polynomial of the regular representation of the general element $\xi = \sum_{i_1=0}^{n_1-1} \cdots \sum_{i_r=0}^{n_r-1} x(i_1,\ldots,i_r)h_1^{i_1}\cdots h_r^{i_r}$ of $F[G]$ is

$$\chi(\lambda) = \prod_{j_1=0}^{n_1-1} \cdots \prod_{j_r=0}^{n_r-1} [\lambda - \sum_{i_1=0}^{n_1-1} \cdots \sum_{i_r=0}^{n_r-1} x(i_1,\ldots,i_r)\omega_1^{i_1 j_1}\cdots\omega_r^{i_r j_r}],$$

where each $\omega_k = \omega^{n/n_k}$.

Proof. We proceed by induction on r.

Case: $r = 1$. When $G \cong Z_n$, the recherché result is a restatement of Lemma 4.3.

Case: $r > 1$. Assume the result for direct products of r-1 cyclic groups. Suppose $n_r = 3$; the proof for arbitrary values of n_r is analogous. Then $G \cong G_0 \times Z_3$, where G_0 comprises r-1 cyclic groups. With respect to a basis ordering the elements of the group properly, the left regular matrix representation of an element $\sum_{i_1=0}^{n_1-1} \cdots \sum_{i_r=0}^{n_r-1} u(i_i,\ldots,i_r)h_1^{i_1}\cdots h_r^{i_r}$ of $F[G]$ has the form

$$M = \begin{bmatrix} N_0 & N_2 & N_1 \\ N_1 & N_0 & N_2 \\ N_2 & N_1 & N_0 \end{bmatrix},$$

where N_0, N_1, and N_2 are left regular $\frac{n}{n_r} \times \frac{n}{n_r}$ matrix representations of certain elements of $F[G_0]$. The matrices N_0, N_1, and N_2 commute because G_0 is abelian. Since ω_r is a primitive cube root of unity, Lemma 4.2 implies that the characteristic polynomial of M is

$$\det(\lambda I_n - M) = \det[(\lambda I_{n/n_r} - (N_0+N_1+N_2))\cdot$$
$$(\lambda I_{n/n_r} - (N_0+\omega_r N_1+\omega_r^2 N_2))(\lambda I_{n/n_r} - (N_0+\omega_r^2 N_1+\omega_r^4 N_2))].$$

By the inductive hypothesis, the characteristic polynomials of $N_0+N_1+N_2$, $N_0+\omega_r N_1+\omega_r^2 N_2$, and $N_0+\omega_r^2 N_1+\omega_r^4 N_2$ are

$$\det(\lambda I - (N_0+N_1+N_2)) =$$
$$\Pi_{j_1}\cdots\Pi_{j_{r-1}}[\lambda - \textstyle\sum_{i_1}\cdots\sum_{i_{r-1}}(u(i_1,\ldots,i_{r-1},0)+$$
$$u(i_1,\ldots,i_{r-1},1)+u(i_1,\ldots,i_{r-1},2))\omega_1^{i_1 j_1}\cdots\omega_{r-1}^{i_{r-1}j_{r-1}}],$$

$$\det(\lambda I - (N_0+\omega_r N_1+\omega_r^2 N_2)) =$$
$$\Pi_{j_1}\cdots\Pi_{j_{r-1}}[\lambda - \textstyle\sum_{i_1}\cdots\sum_{i_{r-1}}(u(i_1,\ldots,i_{r-1},0)+$$
$$\omega_r u(i_1,\ldots,i_{r-1},1)+\omega_r^2 u(i_1,\ldots,i_{r-1},2))\omega_1^{i_1j_1}\cdots\omega_{r-1}^{i_{r-1}j_{r-1}}],$$

$$\det(\lambda I - (N_0+\omega_r^2 N_1+\omega_r^4 N_2)) =$$
$$\Pi_{j_1}\cdots\Pi_{j_{r-1}}[\lambda - \textstyle\sum_{i_1}\cdots\sum_{i_{r-1}}(u(i_1,\ldots,i_{r-1},0)+$$
$$\omega_r^2 u(i_1,\ldots,i_{r-1},1)+\omega_r^4 u(i_1,\ldots,i_{r-1},2))\omega_1^{i_1j_1}\cdots\omega_{r-1}^{i_{r-1}j_{r-1}}].$$

Because $\det(\lambda I_n - M) = \det(\lambda I - (N_0+N_1+N_2))\cdot\det(\lambda I - (N_0+\omega_r N_1+\omega_r^2 N_2))\cdot$ $\det(\lambda I - (N_0+\omega_r^2 N_1+\omega_r^4 N_2)) =$

$$\prod_{j_1=0}^{n_1-1}\cdots\prod_{j_r=0}^{n_r-1}[\lambda - \sum_{i_1=0}^{n_1-1}\cdots\sum_{i_r=0}^{n_r-1} u(i_1,\ldots,i_r)\omega_1^{i_1j_1}\cdots\omega_r^{i_rj_r}]$$

for every selection of elements $u(i_1,\ldots,i_r)$ of F, (4.1) holds incontrovertibly for indeterminates $\{x(i_1,\ldots,i_r)\}$. □

Note that although the abelian group G decomposes into a direct product of cyclic groups of prime power order, the proof of Theorem 4.4 did not require this property of the orders of the cyclic groups.

Since $\chi(\lambda)$ has n distinct linear factors over $F[\lambda,x(0,\ldots,0),\ldots,$ $x(n_1-1,\ldots,n_r-1)]$, Theorem 3.8 asserts that there are n irreducible, inequivalent 1×1 matrix representations of F[G], each of which occurs once in the regular representation. These representations are

$$T_{j_1\cdots j_r}\colon \quad \sum_{i_1=0}^{n_1-1}\cdots\sum_{i_r=0}^{n_r-1} u(i_1,\ldots,i_r)h_1^{i_1}\cdots h_r^{i_r} \mapsto$$

$$\sum_{i_1=0}^{n_1-1}\cdots\sum_{i_r=0}^{n_r-1} u(i_1,\ldots,i_r)\omega_1^{i_1j_1}\cdots\omega_r^{i_rj_r}$$

for $0 \le j_1 \le n_1-1, \ldots, 0 \le j_r \le n_r-1$. Following Chalkley [Cl], one can construct a matrix Ω whose entries are various powers of ω such that $\Omega\rho_L(\xi)\Omega^{-1}$ is diagonal, where ρ_L is the left regular representation with respect to the basis comprising the elements of G. For each $\alpha = \sum_{i_1}\cdots\sum_{i_r} a(i_1,\ldots,i_r)h_1^{i_1}\cdots h_r^{i_r}$ in F[G] the diagonal entries of $\Omega\rho_L(\alpha)\Omega^{-1}$ are the representations $T_{j_1\cdots j_r}(\alpha)$. An n-vector $\hat{\alpha}$ with components $T_{j_1\cdots j_r}(\alpha)$ $(0 \le j_1 \le n_1-1, \ldots, 0 \le j_r \le n_r-1)$ is the multidimensional discrete Fourier transform of an n-vector with components $a(i_1,\ldots,i_r)$ $(0 \le i_1 \le n_1-1, \ldots, 0 \le i_r \le n_r)$. To compute $\gamma = \alpha\beta$ in F[G], one multiplies the multidimensional transforms $\hat{\alpha}$ and $\hat{\beta}$ componentwise to form $\hat{\gamma}$. This method is equivalent to calculating $(\Omega\rho_L(\alpha)\Omega^{-1})(\Omega\rho_L(\alpha)\Omega^{-1}) = \Omega\rho_L(\alpha)\Omega^{-1}$, which involves simultaneously the products of the n representations: $T_{j_1\cdots j_r}(\alpha)T_{j_1\cdots j_r}(\beta) = T_{j_1\cdots j_r}(\gamma)$. This use of the multidimensional discrete Fourier transform is a special case of the procedure outlined in Section 3C.

When F does not contain a primitive nth root of unity, however, a multidimensional approach may be inferior to Algorithm 3.9. For example, consider $\mathbb{Q}[Z_3 \times Z_5] \cong \mathbb{Q}[s,t]/((s^3 - 1)(t^5 - 1))$, where $Z_3 \times Z_5 = \{s^i t^j \mid 0 \le i \le 2, 0 \le j \le 4\}$. To compute $\alpha\beta =$

$(\sum_{i,j} a_{ij}s^it^j)(\sum_{i,j} b_{ij}s^it^j)$, one could calculate $\alpha\beta$ as polynomials in s: one could multiply

$$[\sum_i (\sum_j a_{ij}t^j)s^i][\sum_i (\sum_j b_{ij}t^j)s^i] \bmod (s - 1)$$

with one product of polynomials in t (mod $(t^5 - 1)$) and

$$[\sum_i (\sum_j a_{ij}t^j)s^i][\sum_i (\sum_j b_{ij}t^j)s^i] \bmod (s^2 + s + 1)$$

with three products of polynomials in t (mod $(t^5 - 1)$). By Theorem 2.6, each product of polynomials in t can be calculated with $2\cdot5 - 2 = 8$ nonscalar multiplications because $t^5 - 1$ has two irreducible factors in $\mathbb{Q}[t]$. This multidimensional method requires $4\cdot8 = 32$ nonscalar multiplications. But because $\mathbb{Q}[Z_3 \times Z_5] \cong \mathbb{Q}[Z_{15}] \cong \mathbb{Q}[z]/(z^{15} - 1)$ and $z^{15} - 1$ has 4 distinct irreducible factors over $\mathbb{Q}$, we can multiply elements of $\mathbb{Q}[Z_3 \times Z_5]$ with only $2\cdot15 - 4 = 26$ nonscalar multiplications.

Instead of using a multidimensional approach for $\mathbb{Q}[G]$, one could decompose $\mathbb{Q}[G]$ into a direct sum of simple algebras, as in Theorem 2.7 and Algorithm 3.9. For instance,

$$\mathbb{Q}[Z_3 \times Z_3] \cong \frac{\mathbb{Q}[s,t]}{((s - 1), (t - 1))} \oplus \frac{\mathbb{Q}[s,t]}{((s^2 + s + 1), (t - 1))} \oplus \frac{\mathbb{Q}[s,t]}{((s - 1), (t^2 + t + 1))} \oplus \frac{\mathbb{Q}[s,t]}{((s^2 + s + 1), (t^2 + t + 1))}.$$

Because $\mathbb{Q}[G]$ is commutative, each direct summand is a finite-dimensional extension field of $\mathbb{Q}$, which is perfect. Let $\Delta(m)$ be the number of positive integral divisors of an integer m. In $\mathbb{Q}[z]$ the polynomial $z^n - 1$ has $\Delta(n)$ irreducible factors. If G is the direct product of cyclic groups of orders $n_1, \ldots, n_r$, then the number of fields (simple subalgebras) into which $\mathbb{Q}[G]$ decomposes is $\Delta(n_1)\cdots\Delta(n_r)$. By Theorem 2.7, the multiplicative complexity of $\mathbb{Q}[G]$ is at most $2n - \Delta(n_1)\cdots\Delta(n_r)$, where $n = n_1\cdots n_r$ is the cardinality of G. The multidimensional approach described in the last paragraph requires $\Pi_{i=1}^{r}(2n_i - \Delta(n_i))$ nonscalar multiplications. An easy inductive argument shows that $\Pi_{i=1}^{r}(2n_i - \Delta(n_i)) \geq 2n - \Delta(n_1)\cdots\Delta(n_r)$: if $\Pi_{i=1}^{r-1}(2n_i - \Delta(n_i)) \geq 2n/n_r - \Delta(n_1)\cdots\Delta(n_{r-1})$, then

$$\begin{aligned}\prod_{i=1}^{r}(2n_i - \Delta(n_i)) &\geq (2n/n_r - \Delta(n_1)\cdots\Delta(n_{r-1}))(2n_r - \Delta(n_r)) \\ &\geq 2n - \Delta(n_1)\cdots\Delta(n_r) + \\ &\qquad 2[1 - \Delta(n_r)/n_r][n - \Delta(n_1)\cdots\Delta(n_{r-1})n_r];\end{aligned}$$

and because $1 \geq \Delta(n_r)/n_r$ and $n \geq \Delta(n_1)\cdots\Delta(n_{r-1})n_r$, the conclusion follows; the multidimensional approach is inferior to Algorithm 3.9.

4C. Dihedral Group Algebras

The dihedral group D_r of order 2r is generated by {s, t} with relations $s^r = t^2 = stst = e$, the identity of D_r. We wish to compute products $\alpha\beta = (a_0e + a_1s + \ldots + a_{r-1}s^{r-1} + a_rt + a_{r+1}st + \ldots + a_{2r-1}s^{r-1}t)\cdot(b_0e + b_1s + \ldots + b_{r-1}s^{r-1} + b_rt + b_{r+1}st + \ldots + b_{2r-1}s^{r-1}t) = \gamma = (c_0e + c_1s + \ldots + c_{r-1}s^{r-1} + c_rt + c_{r+1}st + \ldots + c_{2r-1}s^{r-1}t)$ in $F[D_r]$. One may verify that

$$\textstyle\sum_{i=0}^{r-1} c_i z^i = (\sum_{i=0}^{r-1} a_i z^i)(b_0+b_1 z+\dots b_{r-2}z^{r-2}+b_{r-1}z^{r-1}) +$$
$$\textstyle(\sum_{i=0}^{r-1} a_{r+i} z^i)(b_r+b_{2r-1}z+\dots+b_{r+2}z^{r-2}+b_{r+1}z^{r-1}) \bmod (z^r-1),$$

$$\textstyle\sum_{i=0}^{r-1} c_{r+i} z^i = (\sum_{i=0}^{r-1} a_i z^i)(b_r+b_{r+1}z+\dots+b_{2r-2}z^{r-2}+b_{2r-1}z^{r-1}) +$$
$$\textstyle(\sum_{i=0}^{r-1} a_{r+i} z^i)(b_0+b_{r-1}z+\dots+b_2 z^{r-2}+b_1 z^{r-1}) \bmod (z^r-1).$$

Exploiting the inequivalent, irreducible matrix representations of $\mathbb{C}[D_r]$ ([C3, p. 339]), Willsky discovered an efficient procedure for computing products in $\mathbb{C}[D_r]$. The algebra $\mathbb{C}[D_r]$ has two 1×1 matrix representations if r is odd, four if r is even; and $\lfloor\frac{1}{2}(r-1)\rfloor$ 2×2 representations:

$$\begin{array}{ll} T_1(s) = 1, & T_1(t) = 1; \\ T_2(s) = 1, & T_2(t) = -1; \\ T_3(s) = -1, & T_3(t) = 1; \\ T_4(s) = -1, & T_4(t) = -1; \end{array}$$

(the last two lines: $\Big\}$ if r is even;)

and for $k = 1,2,\dots,\lfloor\frac{1}{2}(r-1)\rfloor$,

$$U_k(s) = \begin{bmatrix} \omega^k & 0 \\ 0 & \omega^{-k} \end{bmatrix}, \qquad U_k(t) = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix},$$

where ω is a primitive rth root of unity. As in Section 3C, to compute $\alpha\beta$ one multiplies corresponding representations of α and β--a total of $2 + 7\cdot\frac{r-1}{2}$ nonscalar multiplications if r is odd, $4 + 7\cdot\frac{r-2}{2}$ nonscalar multiplications if r is even (if one multiplies a pair of 2×2 matrices with 7 nonscalar multiplications). Willsky noted that the $U_k(\alpha)$ and $U_k(\beta)$ can be calculated with fast Fourier transforms:

$$U_k(\alpha) = \begin{bmatrix} \sum_{i=0}^{r-1} a_i\omega^{ik} & \sum_{i=0}^{r-1} a_{r+i}\omega^{ik} \\ \sum_{i=0}^{r-1} a_{r+i}\omega^{i(r-k)} & \sum_{i=0}^{r-1} a_i\omega^{i(r-k)} \end{bmatrix}$$

for $k = 1,\ldots,\lfloor\frac{r-1}{2}\rfloor$.

The field $\mathbb{R}$ of real numbers is a splitting field for D_r. The representations U_k above are equivalent to

$$\tilde{U}_k(s) = \begin{bmatrix} Re(\omega^k) & -[Im(\omega^k)]^2 \\ 1 & Re(\omega^k) \end{bmatrix}, \quad \tilde{U}_k(t) = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \qquad (4.2)$$

for $k = 1,\ldots,\lfloor\frac{r-1}{2}\rfloor$ because, letting

$$S_k = \begin{bmatrix} 1 & \sqrt{-1}Im(\omega^k) \\ 1 & -\sqrt{-1}Im(\omega^k) \end{bmatrix},$$

we have $U_k(s)S_k = S_k\tilde{U}_k(s)$ and $U_k(t)S_k = S_k\tilde{U}_k(t)$ for $k = 1,\ldots,\lfloor\frac{r-1}{2}\rfloor$. Thus, the multiplicative complexity of $\mathbb{R}[D_r]$ equals that of $\mathbb{C}[D_r]$ for every dihedral group D_r.

From (4.2) it is apparent that the field $\mathbb{Q}$ of rational numbers is a splitting field for D_3, D_4, and D_6: the division algebras that constitute the simple subalgebras of $\mathbb{Q}[D_3]$, $\mathbb{Q}[D_4]$, and $\mathbb{Q}[D_6]$ coincide with $\mathbb{Q}$. The projections of an element of $\mathbb{Q}[D_3]$ (or $\mathbb{Q}[D_4]$ or $\mathbb{Q}[D_6]$) into these simple subalgebras are its matrix representations over $\mathbb{Q}$. Consider the computation of the product of $\alpha = a_0e + a_1s + a_2s^2 + a_3t + a_4st + a_5s^2t$ and $\beta = b_0e + b_1s + b_2s^2 + b_3t + b_4st + b_5s^2t$, yielding $\gamma = c_0e + c_1s + c_2s^2 + c_3t + c_4st + c_5s^2t$ in $\mathbb{Q}[D_3]$. We multiply corresponding matrix representations of α and β to obtain the representations of γ:

$$T_0(\gamma) = T_0(\alpha)T_0(\beta) = (a_0+a_1+a_2+a_3+a_4+a_5)(b_0+b_1+b_2+b_3+b_4+b_5),$$
$$T_1(\gamma) = T_1(\alpha)T_1(\beta) = (a_0+a_1+a_2-a_3-a_4-a_5)(b_0+b_1+b_2-b_3-b_4-b_5),$$
$$\overline{U}(\gamma) = \overline{U}(\alpha)\overline{U}(\beta) =$$
$$\begin{bmatrix} a_0-a_1+a_3-a_5 & a_1-a_2-a_4+a_5 \\ -a_1+a_2+a_3-a_4 & a_0-a_2-a_3+a_5 \end{bmatrix}\begin{bmatrix} b_0-b_1+b_3-b_5 & b_1-b_2-b_4+b_5 \\ -b_1+b_2+b_3-b_4 & b_0-b_2-b_3+b_5 \end{bmatrix};$$

the representation $\overline{U}$ is equivalent to the representation $\tilde{U}_1$ for $\mathbb{Q}[D_3]$ in (4.2).

The field $\mathbb{Q}$ is not a splitting field for every D_r, however. For instance, no 2×2 matrix with entries in $\mathbb{Q}$ is equivalent to the representation $\tilde{U}_1(s)$ of s in D_5 because the characteristic polynomial of $\tilde{U}_1(s)$ in (4.2) is $\lambda^2 + \frac{1}{2}(\sqrt{5} - 1)\lambda + 1$.

Nevertheless, we construct an efficient algorithm for computing products in $\mathbb{Q}[D_5]$. The characteristic polynomial of the regular representation of the general element $\xi = x_0e + x_1s + \dots + x_4s^4 + x_5t + x_6st + \dots + x_9s^4t$ of $\mathbb{Q}[D_5]$ is $\chi(\lambda) = \chi_1(\lambda)\chi_2(\lambda)\chi_3(\lambda)^2$, where

$$\chi_1(\lambda) = \lambda - (x_0+x_1+x_2+x_3+x_4+x_5+x_6+x_7+x_8+x_9),$$
$$\chi_2(\lambda) = \lambda - (x_0+x_1+x_2+x_3+x_4-x_5-x_6-x_7-x_8-x_9),$$
$$\chi_3(\lambda) = \Theta_0^2 - (\Theta_1 + \Theta_2)\Theta_0 - \Theta_1^2 - \Theta_2^2 + 3\Theta_1\Theta_2,$$
$$\Theta_0 = (x_0-\lambda)^2+x_1^2+x_2^2+x_3^2+x_4^2-x_5^2-x_6^2-x_7^2-x_8^2-x_9^2,$$
$$\Theta_1 = (x_0-\lambda)x_1 + x_1x_2+x_2x_3+x_3x_4+x_4(x_0-\lambda)$$
$$-x_5x_6-x_6x_7-x_7x_8-x_8x_9-x_9x_5,$$
$$\Theta_2 = (x_0-\lambda)x_2+x_2x_4+x_4x_1+x_1x_3+x_3(x_0-\lambda)$$
$$-x_5x_7-x_7x_9-x_9x_6-x_6x_8-x_8x_5.$$

Because $\chi_3(\lambda)$ factors into $(\Theta_0 + \frac{1}{2}(\sqrt{5} - 1)\Theta_1 - \frac{1}{2}(\sqrt{5} + 1)\Theta_2)\cdot$

$(\Theta_0 - \frac{1}{2}(\sqrt{5} + 1)\Theta_1 + \frac{1}{2}(\sqrt{5} - 1)\Theta_2)$ in $\mathbb{R}[\lambda, x_0, \ldots, x_9]$, which is a unique factorization domain, $\chi_3(\lambda)$ is irreducible in $\mathbb{Q}[\lambda, x_0, \ldots, x_9]$. Thus, by Theorem 3.8, $\mathbb{Q}[D_5]$ has three irreducible, inequivalent matrix representations. Let $\mathbb{Q}[D_5] \cong A_1 \oplus A_2 \oplus A_3$, where A_i corresponds to χ_i for i = 1,2,3. Then $A_1 \cong A_2 \cong \mathbb{Q}$; the simple algebras A_1 and A_2 are rings of 1×1 matrices over $\mathbb{Q}$. Let A_3 be a full $\mu_3 \times \mu_3$ matrix ring over some division algebra $\mathcal{D}_3$. According to the discussion at the end of Section 3B, $[A_3 : \mathbb{Q}] = \mu_3^2 [\mathcal{D}_3 : \mathbb{Q}] = \deg(\chi_3(\lambda)^2) = 8$; moreover, since χ_3 has multiplicity 2, $\mu_3 | 2$. From (4.2) we deduce that $\mu_3 = 2$ and $\mathcal{D}_3 = \mathbb{Q}(\sqrt{5}) \cong \mathbb{Q}[z]/(z^2-5)$. The projection of any element $\alpha = a_0 e + a_1 s + \ldots + a_4 s^4 + a_5 t + a_6 st + \ldots + a_9 s^4 t$ into A_3 is defined by

$$s \mapsto \begin{bmatrix} \frac{1}{4}(\sqrt{5} - 1) & -\frac{1}{8}(\sqrt{5} + 5) \\ 1 & \frac{1}{4}(\sqrt{5} - 1) \end{bmatrix}, \quad t \mapsto \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

To compute the product $\alpha\beta$ in $\mathbb{Q}[D_5]$ with Algorithm 3.9 we use the projections of α and β in the simple subalgebras A_1, A_2, and A_3. Multiplying projections of α and β in A_1 and A_2 requires 1 nonscalar multiplication each. Multiplying the projections of α and β in A_3 requires at most 7•3 nonscalar multiplications; the factor 7 arises from 2×2 matrix multiplication, the factor 3 from multiplying elements of $\mathcal{D}_3 \cong \mathbb{Q}[z]/(z^2 - 5)$. The multiplicative complexity of $\mathbb{Q}[D_5]$ is at most $1 + 1 + 7 \cdot 3 = 23$.

4D. Generalized Quaternion Group Algebras

Related to the dihedral group D_{2r}, the generalized quaternion group H_r of order 4r is generated by $\{s, t\}$ with relations $s^{2r} = t^2 s^r = stst^{-1} = e$, the identity of H_r. If $r \geq 2$, then H_r is noncommutative.

The algebra $\mathbb{C}[H_r]$ has four irreducible, inequivalent representations of dimension 1 and r-1 of dimension 2 [C3, p. 339]:

$$\begin{array}{ll} T_1(s) = 1, & T_1(t) = 1; \\ T_2(s) = 1, & T_2(t) = -1; \\ T_3(s) = -1, & T_3(t) = 1; \\ T_4(s) = -1, & T_4(t) = -1; \end{array} \qquad (4.3)$$

and for $k = 1,\ldots,r-1$,

$$U_k(s) = \begin{bmatrix} \omega^k & 0 \\ 0 & \omega^{-k} \end{bmatrix}, \qquad U_k(t) = \begin{bmatrix} 0 & \omega^{-kr} \\ 1 & 0 \end{bmatrix}, \qquad (4.4)$$

where ω is a primitive $2r$ th root of unity. According to Section 3C, one can exploit these representations to compute products in $\mathbb{C}[H_r]$ with $4 + 7(r-1)$ nonscalar multiplications: if α and β are in $\mathbb{C}[H_r]$, then each product $T_k(\alpha)T_k(\beta)$ requires 1 nonscalar multiplication and each 2×2 matrix product $U_k(\alpha)U_k(\beta)$ for $k = 1,\ldots,r-1$ requires 7.

Computing products in $\mathbb{R}[H_r]$ and $\mathbb{Q}[H_r]$ apparently requires more nonscalar multiplications than computations in $\mathbb{C}[H_r]$. The algebra $\mathbb{R}[H_2]$ decomposes into $\mathbb{R} \oplus \mathbb{R} \oplus \mathbb{R} \oplus \mathbb{R} \oplus \mathcal{Q}_{\mathbb{R}}$, a direct sum of simple algebras, each of which is a 1×1 matrix ring over a division algebra. The projections into the first four summands are the one-dimensional representations

(4.3) of $\mathbb{R}[H_2]$. A projection of $\alpha = a_0 + a_1s + a_2s^2 + a_3s^3 + a_4t + a_5st + a_6s^2t + a_7s^3t$ into the last summand, the real quaternions $\mathcal{Q}_{\mathbb{R}}$, is

$$\alpha \mapsto (a_0 - a_2) + (a_1 - a_3)\hat{\imath} + (a_4 - a_6)\hat{\jmath} + (a_5 - a_7)\hat{k}. \qquad (4.5)$$

Dobkin [D1] proved that one can multiply quaternions over any field whose characteristic is not 2 with 8 nonscalar multiplications. Products in $\mathbb{R}[H_2]$ can be computed with 1 + 1 + 1 + 1 + 8 = 12 nonscalar multiplications. This analysis also applies to $\mathbb{Q}[H_2] \cong \mathbb{Q} \oplus \mathbb{Q} \oplus \mathbb{Q} \oplus \mathbb{Q} \oplus \mathcal{Q}_{\mathbb{Q}}$: projections into the first four summands are defined by (4.3), and a projection into the last summand, the rational quaternions $\mathcal{Q}_{\mathbb{Q}}$, is given by (4.5). One can calculate products in $\mathbb{C}[H_2]$ with 12 nonscalar multiplications.

The optimality of these procedures for $\mathbb{R}[H_2]$ and $\mathbb{Q}[H_2]$ seems to hinge on the "Direct Sum Conjecture" (see Section 3C) and the assessment of the multiplicative complexity of the quaternions; both problems have been studied extensively by other researchers. Theorem 2.4 provides the best lower bound known for the multiplicative complexity of $\mathcal{Q}_{\mathbb{R}}$ in our model of computation: 7 nonscalar multiplications.

The algebra $\mathbb{R}[H_3]$ is more complicated than $\mathbb{R}[H_2]$. The characteristic polynomial of the regular representation of the general element $\xi = x_0e + x_1s + \ldots + x_5s^5 + x_6t + x_7st + \ldots + x_{11}s^5t$ of $\mathbb{R}[H_3]$ has prime power factorization $\chi_1(\lambda)\chi_2(\lambda)\chi_3(\lambda)\chi_4(\lambda)(\chi_5(\lambda))^2(\chi_6(\lambda))^2$ in $\mathbb{R}[\lambda, x_0, \ldots, x_{11}]$, where

$$\chi_1(\lambda) = \lambda - (x_0+x_1+x_2+x_3+x_4+x_5+x_6+x_7+x_8+x_9+x_{10}+x_{11}),$$
$$\chi_2(\lambda) = \lambda - (x_0+x_1+x_2+x_3+x_4+x_5-x_6-x_7-x_8-x_9-x_{10}-x_{11}),$$

$$\chi_3(\lambda) = \lambda - (x_0-x_1+x_2-x_3+x_4-x_5+x_6-x_7+x_8-x_9+x_{10}-x_{11}),$$

$$\chi_4(\lambda) = \lambda - (x_0-x_1+x_2-x_3+x_4-x_5-x_6+x_7-x_8+x_9-x_{10}+x_{11}),$$

$$\begin{aligned}\chi_5(\lambda) = {} & (x_0+x_3-x_1-x_4+x_6+x_9-x_8-x_{11}-\lambda)(x_0+x_3-x_2-x_5-x_6-x_9+x_8+x_{11}-\lambda) \\ & - (-x_1-x_4+x_2+x_5+x_6+x_9-x_7-x_{10})(x_1+x_4-x_2-x_5-x_7-x_{10}+x_8+x_{11}),\end{aligned}$$

$$\begin{aligned}\chi_6(\lambda) = {} & \Theta_0(x_0-\lambda,x_1,x_2,x_3,x_4,x_5)^2 + \frac{3}{2}\Theta_1(x_1,x_2,x_4,x_5)^2 \\ & + \Theta_0(x_6,x_7,x_8,x_9,x_{10},x_{11})^2 + \frac{3}{2}\Theta_1(x_7,x_8,x_{10},x_{11})^2\end{aligned}$$

$$\left.\begin{aligned}\Theta_0(z_0,z_1,z_2,z_3,z_4,z_5) &= z_0 + \frac{1}{2}z_1 - \frac{1}{2}z_2 - z_3 - \frac{1}{2}z_4 + \frac{1}{2}z_5, \\ \Theta_1(z_1,z_2,z_4,z_5) &= z_1 + z_2 - z_4 - z_5.\end{aligned}\right\} \tag{4.6}$$

Let $\mathbb{R}[H_3] \cong A_1 \oplus \ldots \oplus A_6$, where χ_i is the characteristic polynomial of ξ corresponding to the simple subalgebra A_i for $i = 1,\ldots,6$. Then $A_1 \cong A_2 \cong A_3 \cong A_4 \cong \mathbb{R}$; the projections of an element $\alpha = a_0e + a_1s + \ldots + a_5s^5 + a_6t + a_7st + \ldots + a_{11}s^5t$ of $\mathbb{R}[H_3]$ into these four subalgebras are the one-dimensional representations (4.3). A projection of α into A_5 is a 2×2 matrix representation over $\mathbb{R}$:

$$\alpha \to \begin{bmatrix} a_0+a_3-a_1-a_4+a_6+a_9-a_8-a_{11} & a_1+a_4-a_2-a_5-a_7-a_{10}+a_8+a_{11} \\ -a_1-a_4+a_2+a_5+a_6+a_9-a_7-a_{10} & a_0+a_3-a_2-a_5-a_6-a_9+a_8+a_{11} \end{bmatrix}. \tag{4.7}$$

This representation is equivalent to (4.4) in which $k = 2$ and ω is a primitive sixth root of unity. Since A_5 is the ring of 2×2 matrices over $\mathbb{R}$, products in A_5 require exactly 7 nonscalar multiplications. Finally, $A_6 \cong \mathcal{Q}_{\mathbb{R}}$, the real quaternions. A projection of α into A_6 is

$$\alpha \mapsto \Theta_0(a_0,a_1,a_2,a_3,a_4,a_5) + \Theta_1(a_1,a_2,a_4,a_5)\,\frac{\sqrt{3}}{2}\,\hat{i} + \Theta_0(a_6,a_7,a_8,a_9,a_{10},a_{11})\hat{j} + \Theta_1(a_7,a_8,a_{10},a_{11})\,\frac{\sqrt{3}}{2}\,\hat{k}, \qquad (4.8)$$

where Θ_0 and Θ_1 are defined in (4.6). One can compute a product of real quaternions (in A_6) with 8 nonscalar multiplications. Therefore, using Algorithm 3.9, we can compute products of elements of $\mathbb{R}[H_3]$ with $1 + 1 + 1 + 1 + 7 + 8 = 19$ nonscalar multiplications.

The structure of $\mathbb{Q}[H_3]$ resembles that of $R[H_3]$: $\mathbb{Q}[H_3] \cong A_1' \oplus A_2' \oplus A_3' \oplus A_4' \oplus A_5' \oplus A_6'$, where $A_1' \cong A_2' \cong A_3' \cong A_4' \cong \mathbb{Q}$, and A_5' is a ring of 2×2 matrices over $\mathbb{Q}$; projections of $\alpha = a_0e + a_1s + \ldots + a_5s^5 + a_6t + a_7st + \ldots + a_{11}s^5t$ into A_1', A_2', A_3', and A_4' are defined by (4.3), and a projection into A_5' is given by (4.7). The simple subalgebra A_6' is a division algebra that resembles the rational quaternions; a basis for A_6' is $\{1, \frac{\sqrt{3}}{2}\hat{i}, \hat{j}, \frac{\sqrt{3}}{2}\hat{k}\}$, and (4.8) is a projection of α into A_6'. Fiduccia's decomposition procedure [FO] yields an algorithm for computing $(u_0 + u_1\frac{\sqrt{3}}{2}\hat{i} + u_2\hat{j} + u_3\frac{\sqrt{3}}{2}\hat{k})(w_0 + w_1\frac{\sqrt{3}}{2}\hat{i} + w_2\hat{j} + w_3\frac{\sqrt{3}}{2}\hat{k}) = y_0 + y_1\frac{\sqrt{3}}{2}\hat{i} + y_2\hat{j} + y_3\frac{\sqrt{3}}{2}\hat{k}$ in A_6' with 10 nonscalar multiplications: let

$$\begin{aligned}
m_1 &= u_1(w_0 - \tfrac{3}{2}w_1), & m_6 &= u_3(\tfrac{3}{2}w_1 - w_2),\\
m_2 &= u_1(w_2 - \tfrac{3}{2}w_3), & m_7 &= (u_0 - u_1 - u_2 - u_3)w_0,\\
m_3 &= u_2(w_0 - w_2), & m_8 &= (u_0 + \tfrac{3}{2}u_1 + u_2 - \tfrac{3}{2}u_3)w_1,\\
m_4 &= u_2(-w_1 + w_3), & m_9 &= (u_0 - u_1 + u_2 + u_3)w_2,\\
m_5 &= u_3(w_0 - \tfrac{3}{2}w_3), & m_{10} &= (u_0 + \tfrac{3}{2}u_1 - u_2 + \tfrac{3}{2}u_3)w_3;
\end{aligned}$$

then $y_0 = m_1 + m_3 + m_5 + m_7$, $y_1 = m_1 + m_4 + m_6 + m_8$, $y_2 = m_2 + m_3 +$

$m_6 + m_9$, and $y_3 = m_2 + m_4 + m_5 + m_{10}$. The multiplicative complexity of $\mathbb{Q}[H_3]$ is at most $1 + 1 + 1 + 1 + 7 + 10 = 21$.

It is not known whether these algorithms for generalized quaternion group algebras are optimal. Even if Algorithm 3.9 were optimal for semisimple algebras in general, it might be nonoptimal when applied to a specific algebra such as $\mathbb{R}[H_3]$.

5. FURTHER PROBLEMS

In addition to the important conjecture formulated at the end of Section 3C we propose three subjects for further investigation.

1. Non-semisimple algebras. Examples include quaternions over fields of characteristic 2 and abelian group algebras $F[G]$ in which $\mathrm{char}\, F \mid \mathrm{card}\, G$.

Consider $F[Z_2]$, where Z_2 is the cyclic group of order 2. If $\mathrm{char}\, F \neq 2$, then $1 \neq -1$, and computing products in

$$F[Z_2] \cong F[z]/(z^2 - 1) \cong F[z]/(z - 1) \oplus F[z]/(z + 1)$$

requires exactly 2 nonscalar multiplications. If $\mathrm{char}\, F = 2$, then $1 = -1$, and by Theorems 2.6 and 2.10, computing products in

$$F[z]/(z^2 - 1) \cong F[z]/(z + 1)^2$$

requires exactly 3 nonscalar multiplications.

The technique of Section 3C cannot be used for non-semisimple algebras: unlike semisimple algebras, a non-semisimple algebra cannot be expressed as a direct sum of simple algebras. Equivalently, some reducible representations of a non-semisimple algebra might not be completely reducible. Changing the basis of the algebra might not completely block-diagonalize the regular matrix representation.

For example, if $\mathrm{char}\, F \neq 2$, one can compute products in $F[Z_2 \times Z_2]$ with 4 nonscalar multiplications:

$$\begin{aligned}F[Z_2 \times Z_2] &\cong F[y,z]/((y^2 - 1),(z^2 - 1)) \cong \\ &F[y,z]/((y - 1),(z - 1)) \oplus F[y,z]/((y - 1),(z + 1)) \oplus \\ &F[y,z]/((y + 1),(z - 1)) \oplus F[y,z]/((y + 1),(z + 1)).\end{aligned}$$

The group algebra $\mathbb{F}_4[Z_2 \times Z_2]$ seems to require more nonscalar multiplications. Let τ be a primitive element of $\mathbb{F}_4$ and let s and t generate $Z_2 \times Z_2$: $\mathbb{F}_4[Z_2 \times Z_2] = \{a_0e + a_1s + a_2t + a_3st \mid s^2 = t^2 = e,\ st = ts,$ every $a_i \in \{0, \tau^0, \tau^1, \tau^2\}\}$. Write $(a_0e + a_1s + a_2t + a_3st)(b_0e + b_1s + b_2t + b_3st) = (c_0e + c_1s + c_2t + c_3st)$ as a matrix-vector product in which the matrix is the left regular matrix representation of $a_0e + a_1s + a_2t + a_3st$ with respect to the basis $\{e, s, t, st\}$:

$$\begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} a_0 & a_1 & a_2 & a_3 \\ a_1 & a_0 & a_3 & a_2 \\ a_2 & a_3 & a_0 & a_1 \\ a_3 & a_2 & a_1 & a_0 \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ b_3 \end{bmatrix},$$

$$\begin{bmatrix} c_0+c_1+c_2+c_3 \\ c_1+c_0 \\ c_2+c_0 \\ c_3+c_0 \end{bmatrix} = \begin{bmatrix} a_0+a_1+a_2+a_3 & 0 & 0 & 0 \\ 0 & a_0+a_1 & a_3+a_2 & a_2+a_3 \\ 0 & a_3+a_1 & a_0+a_2 & a_1+a_3 \\ 0 & a_2+a_1 & a_1+a_2 & a_0+a_3 \end{bmatrix} \begin{bmatrix} b_0+b_1+b_2+b_3 \\ b_1+b_0 \\ b_2+b_0 \\ b_3+b_0 \end{bmatrix}.$$

The representation

$$\begin{bmatrix} a_0+a_1 & a_3+a_2 & a_2+a_3 \\ a_3+a_1 & a_0+a_2 & a_1+a_3 \\ a_2+a_3 & a_1+a_2 & a_0+a_3 \end{bmatrix}$$

is reducible, but not completely reducible. Since the subspace spanned by $[\tau^0 \ \ \tau^1 \ \ \tau^2]^t$ is fixed by this matrix, a change of basis yields

$$\begin{bmatrix} a_0+a_1+(a_2+a_1)\tau^1 & a_3+a_2+(a_1+a_2)\tau^1 & 0 \\ a_3+a_1+(a_2+a_1)\tau^2 & a_0+a_2+(a_1+a_2)\tau^2 & 0 \\ (a_2+a_1)\tau^1 & (a_1+a_2)\tau^1 & a_0+a_1+a_2+a_3 \end{bmatrix}.$$

How quickly can products in non-semisimple algebras be computed? In Section 2B we discussed quotient polynomial algebras that may be non-semisimple. If $\pi(z) \in F[z]$ and $\pi(z)$ has prime factorization $\pi_1(z)^{\nu 1} \ldots \pi_k(z)^{\nu k}$, then

$$F[z]/(\pi(z)) \cong F[z]/(\pi_1(z)^{\nu 1}) \oplus \ldots \oplus F[z]/(\pi_k(z)^{\nu k}).$$

This decomposition leads to Algorithm 2.5. Can the methods of Secion 2B be extended to other non-semisimple algebras?

2. Small fields. Algorithm 2.5 and Theorem 3.8 both require sufficiently large fields. Can these noisome cardinality constraints be removed? A fast method for multiplying elements of $\mathbb{F}_2[z]/(\pi(z))$ in which π is irreducible would be of interest to coding theorists [B1] [L2].

3. Calculating the transform. We have emphasized the minimization of nonscalar multiplications at the expense of additions and scalar multiplications. If we instead seek the shortest program, we should minimize the total number of arithmetic operations for calculating the transform (projections) of Algorithm 3.9 and its inverse. The algorithms of Moenck and Borodin [M1] [B2], which generalize the fast Fourier transform, apply to quotient polynomial algebras. We can characterize the Moenck-Borodin technique abstractly: for instance, to calculate the projections of an element α in the simple algebras A_1, A_2, A_3, and A_4 of

$A \cong A_1 \oplus A_2 \oplus A_3 \oplus A_4$, first calculate the projections α_{12} in $A_1 \oplus A_2$ and α_{34} in $A_3 \oplus A_4$; then calculate the projections of α_{12} in A_1 and A_2 and the projections of α_{34} in A_3 and A_4. Can this abstract characterization be developed sufficiently to permit formulation of fast procedures for calculating the transform (projections) and inverse described in Section 3C? Nicholson [N1] discussed theoretical underpinnings of the transform for abelian group rings. Can we obtain lower bounds on the complexity of calculating the transform?

6. SUMMARY AND CONCLUSIONS

In this thesis we examined computations of products in finite-dimensional, associative algebras, especially semisimple algebras.

We established the correctness of Theorem 2.10: the multiplicative complexity of the quotient polynomial algebra $F[z]/(\pi(z))$ is at lest $2n-k$, where $n = \deg \pi$ and k is the number of (nontrivial) distinct irreducible factors of π.

Having defined the general element ξ of a semisimple algebra A and the characteristic polynomial $\chi(\lambda)$ of the regular representation of ξ, we proved that the irreducible factors of χ correspond to the irreducible, inequivalent representations of A: when the base field of A is sufficiently large, the polynomial χ has m distinct irreducible factors if, and only if, A is the direct sum of m simple algebras; for each of the irreducible representations T_i the characteristic polynomial of $T_i(\xi)$ is the power of some irreducible factor of χ. The characteristic polynomial of the regular representation of the general element is invariant under change of basis, which yields a computationally equivalent problem.

We devised a procedure for multiplying elements of a semisimple algebra A with a generalization of the discrete Fourier transform. Let $A \cong A_1 \oplus \ldots \oplus A_m$, where each simple algebra A_i is a full $\mu_i \times \mu_i$ matrix ring over a division algebra $\mathcal{D}_i$. In essence, one multiplies the projections of the multiplicands in the simple subalgebras $A_1, \ldots, A_m$. This method uses at most $\sum_{i=1}^{m} \mu_i^3 [\mathcal{D}_i : F]^2$ nonscalar multiplications, though this number is smaller in many cases.

To elucidate these ideas, we demonstrated that known algorithms for computing products in quotient polynomial algebras and certain abelian group algebras are instances of our general technique. We presented a new explicit calculation of the characteristic polynomial of the regular representation of a general element of an abelian group algebra in which the base field contains a primitive root of unity. Finally, we constructed algorithms for products in algebras of dihedral groups and generalized quaternion groups.

We have focused on minimizing the number of nonscalar multiplications for computing $\{c_1, \ldots, c_n\}$ from $\{a_1, \ldots, a_n, b_1, \ldots, b_n\}$, where $\sum_{i=1}^{n} c_i v_i = (\sum_{i=1}^{n} a_i v_i)(\sum_{i=1}^{n} b_i v_i)$ in an algebra with basis $\{v_1, \ldots, v_n\}$; consequently, we have drawn many of our ideas and techniques from linear algebra. Theorem 2.1 guarantees that since $\{c_1, \ldots, c_n\}$ are bilinear forms, one may assume that each nonscalar multiplications step is the product of linear combinations of $\{a_1, \ldots, a_n, b_1, \ldots, b_n\}$; the $\{c_1, \ldots, c_n\}$ are linear combinations of the results of these steps. The proofs of Lemma 2.9 and Theorem 2.10 appeal to linear independence of vectors. The transform defined in Section 3C is essentially a change of basis: calculating projections in simple subalgebras involves linear combinations of the indeterminates. Morover, methods based on linear algebra have established only linear lower bounds on nonscalar multiplicative complexity (e.g., [F2], [W3]). Have we exhausted the insights available from applications of linear algebra? To obtain novel results--efficient algorithms, stronger (e.g., nonlinear) lower bounds--it seems that we must use more sophisticated mathematical methods such as concepts from algebraic geometry [S2].

NOTATION

$\mathbb{C}$	field of complex numbers
$\mathbb{F}_r$	finite field with r elements
$\mathbb{Q}$	field of rational numbers
$\mathbb{R}$	field of real numbers
$\mathcal{Q}_F$	algebra of quaternions over a field F
D_r	dihedral group of order 2r
H_r	generalized quaternion group of order 4r
I_n	n×n identity matrix
Z_n	cyclic group of order n
$\Delta(m)$	number of positive integral divisors of an integer m
$\mathcal{E}(\theta_k)$	evaluation of a statement in a straight-line program
$\mathcal{O}(f(n))$	a function bounded by cf(n) for some constant c
$\rho_L(\alpha)$	left regular matrix representation of α
card Y	cardinality of an object Y
char F	characteristic of a field F
deg ψ	degree of a polynomial ψ
det A	determinant of a matrix A
$\mathrm{End}_R(X)$	ring of endomorphisms of an R-module X
$\mathcal{I}m(u)$	imaginary part of a complex number u
$\mathcal{R}e(u)$	real part of a complex number u
$\phi_1 \equiv \phi_2 \pmod{\psi}$	the polynomials ϕ_1 and ϕ_2 are congruent modulo ψ
$\phi_1 = \phi_2 \bmod \psi$	the polynomial ϕ_1 is the residue (remainder) when ϕ_2 is divided by ψ

$[A:F]$	dimension of an algebra A over a field F
$F[G]$	algebra of a group G over a field F
$F[z]/(\pi(z))$	ring of polynomials in $F[z]$ modulo $\pi(z)$
$F[x_1,\ldots,x_n]$	ring of polynomials in indeterminates $x_1, \ldots, x_n$
$F^{n\times n}$	ring of $n\times n$ matrices with entries in F (homomorphic to $\mathrm{End}_F(F^n)$)
$(\pi(z))$	ideal generated by $\pi(z)$
$<h>$	cyclic group generated by h
$\lfloor u \rfloor$	largest integer not greater than u
$\lceil u \rceil$	smallest integer not less than u
$m\mid n \quad (m\nmid n)$	m does (does not) divide n
A^t	transpose of a matrix A
$\cong$	isomorphism of objects
$\oplus$	direct sum
$\square$	end of proof or algorithm

REFERENCES

[A1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman, The Design and Analysis of Computer Algorithms. Reading, Mass.: Addison-Wesley, 1974.

[B1] E. R. Berlekamp, Algebraic Coding Theory. New York: McGraw-Hill, 1968.

[B2] A. Borodin and R. Moenck, "Fast Modular Transforms." J. Computer and Systems Sci. 8 (1974) 366-386.

[B3] A. Borodin and I. Munro, The Computational Complexity of Algebraic and Numeric Problems. New York: American Elsevier, 1975.

[C1] R. Chalkley, "Matrices Derived from Finite Abelian Groups." Mathematics Magazine 49 (1976) 121-129.

[C2] J. W. Cooley and J. W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series." Mathematics of Computation 19 (1965) 297-301.

[C3] C. W. Curtis and I. Reiner, Representation Theory of Finite Groups and Associative Algebras. New York: Wiley Interscience, 1962.

[D1] D. Dobkin, "On the Arithmetic Complexity of a Class of Arithmetic Computations." Research Report No. 23, Dept. of Computer Science, Yale University, 1973.

[F0] C. M. Fiduccia, "Fast Matrix Multiplication." Proc. Third Annual ACM Symposium on Theory of Computing, Shaker Heights, Ohio, 1971, pp. 45-49.

[F1] C. M. Fiduccia, "Polynomial Evaluation via the Division Algorithm: the Fast Fourier Transform Revisited." Proc. Fourth Annual ACM Symposium on Theory of Computing, Denver, Colo., 1972, pp. 88-93.

[F2] C. M. Fiduccia and Y. Zalcstein, "Algebras Having Linear Multiplicative Complexities." Technical Report No. 46, Department of Computer Science, State University of New York at Stony Brook, August 1975.

[F3] P. C. Fischer, "Further Schemes for Combining Matrix Algorithms." Automata, Languages and Programming (Second Colloquium on Automata, Languages, and Programming, Saarbrücken, Germany), ed. J. Loeckx. Berlin: Springer-Verlag, 1974, pp. 428-436.

[H1] I. N. Herstein, Noncommutative Rings. Menasha, Wis.: Mathematical Association of America, 1968.

[H2] J. E. Hopcroft and L. R. Kerr, "On Minimizing the Number of Multiplications Necessary for Matrix Multiplication." SIAM J. Applied Mathematics 20 (1971) 30-36.

[I1] M. H. Ingraham, "A Note on Determinants." Bull. AMS 43 (1937) 579-580.

[J1] N. Jacobson, Basic Algebra I. San Francisco: W. H. Freeman, 1974.

[K1] D. E. Knuth, The Art of Computer Programming vol. 2: Seminumerical Algorithms. Reading, Mass.: Addison-Wesley, 1971.

[L1] J. D. Laderman, "A Noncommutative Algorithm for Multiplying 3×3 Matrices Using 23 Multiplications." Bull. AMS 82 (1976) 126-128.

[L2] J. H. van Lint, Coding Theory. Berlin: Springer-Verlag, 1971.

[M1] R. Moenck and A. Borodin, "Fast Modular Transformations via Division." Proc. Thirteenth Annual IEEE Symposium on Switching and Automata Theory, 1972, pp. 90-96.

[M2] J. Morgenstern, "Note on a Lower Bound of the Linear Complexity of the Fast Fourier Transform." J. ACM 20 (1973) 305-306.

[N1] P. J. Nicholson, "Algebraic Theory of Finite Fourier Transforms." J. Computer and Systems Sci. 5 (1971) 524-547.

[P1] R. L. Probert, "Commutativity, Non-commutativity, and Bilinearity." Information Processing Letters 5 (1976) 46-49.

[R1] C. M. Rader, "Discrete Fourier Transform when the Number of Data Samples is Prime." Proc. IEEE 56 (1968) 1107-1108.

[S1] V. Strassen, "Gaussian Elimination is not Optimal." Numerische Mathematik 13 (1969) 354-356.

[S2] V. Strassen, "Die Berechnungskomplexität von elementarysymmetrischen Funktionen und von Interpolationskoeffizienten." Numerische Mathematik 20 (1973) 238-251.

[W1] A. S. Willsky, "Filtering for Random Finite Group Homomorphic Sequential Systems." Mathematical Systems Theory (Proc. of the International Symposium, Undine, Italy, 1975), ed. G. Marchesini and S. K. Mitter. Berlin: Springer-Verlag, 1976, pp. 312-321.

[W2] S. Winograd, "A New Algorithm for Inner Product." IEEE Trans. Computers C-17 (1968) 693-694.

[W3] S. Winograd, "On the Number of Multiplications Required to Compute Certain Functions." Comm. Pure and Applied Mathematics 23 (1970) 165-179.

[W4] S. Winograd, "On Multiplication of 2×2 Matrices." Linear Algebra and its Applications 4 (1971) 381-388.

[W5] S. Winograd, "Some Bilinear Forms whose Multiplicative Complexity Depends on the Field of Constants." IBM Research Report RC5669, October 1975.

[W6] S. Winograd, "On Computing the Discrete Fourier Transform." Proc. National Academy of Sciences, USA 73 (1976) 1005-1006.