

NEARLY OPTIMAL ALGORITHMS FOR CANONICAL MATRIX FORMS

MARK GIESBRECHT *

Abstract. A Las Vegas type probabilistic algorithm is presented for finding the Frobenius canonical form of an $n \times n$ matrix T over any field \mathbf{K} . The algorithm requires $O^*(\text{MM}(n)) = \text{MM}(n) \cdot (\log n)^{O(1)}$ operations in \mathbf{K} , where $O(\text{MM}(n))$ operations in \mathbf{K} are sufficient to multiply two $n \times n$ matrices over \mathbf{K} . This nearly matches the lower bound of $\Omega(\text{MM}(n))$ operations in \mathbf{K} for this problem, and improves on the $O(n^4)$ operations in \mathbf{K} required by the previously best known algorithms. We also demonstrate a fast parallel implementation of our algorithm for the Frobenius form, which is processor-efficient on a PRAM. As an application we give an algorithm to evaluate a polynomial $g \in \mathbf{K}[x]$ at T which requires only $O^*(\text{MM}(n))$ operations in \mathbf{K} when $\deg g \leq n^2$. Other applications include sequential and parallel algorithms for computing the minimal and characteristic polynomials of a matrix, the rational Jordan form of a matrix, for testing whether two matrices are similar, and for matrix powering, which are substantially faster than those previously known.

Key words. Frobenius form, Jordan form, evaluating polynomials at matrices, matrix powering, matrix multiplication, processor-efficient parallel algorithms.

AMS subject classifications. 15-04, 15A21, 65Y05, 65Y20

Computing a canonical or normal form of an $n \times n$ matrix T over any field \mathbf{K} is a classical mathematical problem with many practical applications. A fundamental theorem of linear algebra states that any $T \in \mathbf{K}^{n \times n}$ is similar to a unique matrix $S \in \mathbf{K}^{n \times n}$ of the block diagonal form

$$S = \text{diag}(C_{f_1}, C_{f_2}, \dots, C_{f_k}) = \begin{pmatrix} \boxed{C_{f_1}} & & & \mathbf{0} \\ & & & \\ & & \boxed{C_{f_2}} & \\ & & & \ddots \\ \mathbf{0} & & & & \boxed{C_{f_k}} \end{pmatrix} \in \mathbf{K}^{n \times n}, \quad (0.1)$$

where each C_{f_i} is the companion matrix of some monic $f_i \in \mathbf{K}[x]$ for $1 \leq i \leq k$, and $f_i \mid f_{i-1}$ for $2 \leq i \leq k$ (we write $T \sim S$ to denote similarity). Recall that the companion matrix C_g of a monic $g = \sum_{0 \leq j \leq r} b_j x^j \in \mathbf{K}[x]$ has the form

$$C_g = \begin{pmatrix} 0 & 0 & & & -b_0 \\ 1 & 0 & & & -b_1 \\ & 1 & & \mathbf{0} & \vdots \\ & & \ddots & & \vdots \\ 0 & & & \ddots & 1 & -b_{r-1} \end{pmatrix} \in \mathbf{K}^{r \times r}.$$

* Department of Computer Science, University of Manitoba, Winnipeg, Manitoba, Canada, R3T 2N2. Research was supported in part by Natural Sciences and Engineering Research Council of Canada research grant OGP0155376.

A matrix S with these properties, called the *Frobenius form* of T , always exists and is unique and *rational*, i.e., it remains the same even if the entries of T are allowed to lie in an algebraic closure of \mathbf{K} . The polynomials $f_1, \dots, f_k \in \mathbf{K}[x]$ are the *invariant factors* of T , and the first of these, f_1 , is the *minimal polynomial* of T , that is f_1 is the polynomial of smallest degree in $\mathbf{K}[x] \setminus \{0\}$ such that $f_1(T) = 0$. The product $f_1 \cdots f_k$ is the characteristic polynomial of T . Since T is similar to S , by definition there exists an invertible $U \in \mathbf{K}^{n \times n}$ such that $U^{-1}TU = S$. Our approach to finding the Frobenius form is to find such a matrix U , from which we find S . Two excellent references for the background information are Gantmacher (1990), Chapter 7, and Hoffman & Kunze (1971), Chapter 7.

In Sections 1 through 4 we link the cost of computing the Frobenius form to the cost of multiplying two matrices. Specifically, if $\text{MM}(n)$ operations in \mathbf{K} are sufficient to multiply two $n \times n$ matrices over a field \mathbf{K} , we show that $O(\text{MM}(n)) = \text{MM}(n) \cdot (\log n)^{O(1)}$ operations in \mathbf{K} suffice to compute the Frobenius form of an $n \times n$ matrix over \mathbf{K} . This algorithm is of the Las Vegas type — it is allowed to choose elements randomly and uniformly from a finite subset of \mathbf{K} at unit cost, and with probability at least $1/4$ it returns the correct answer, otherwise it reports failure. An incorrect answer is never returned. Our algorithm only works as stated when $\#\mathbf{K} \geq n^2$. When \mathbf{K} has $q < n^2$ elements, we embed \mathbf{K} in a field \mathbf{F} of degree $O(\log n)$ over \mathbf{K} which does possess n^2 elements. The Frobenius form of T lies in $\mathbf{K}^{n \times n}$ since it is rational, but $U \in \mathbf{F}^{n \times n}$ may have entries in $\mathbf{F} \setminus \mathbf{K}$. In this case the running time of our algorithms is multiplied by a small power of $\log_q n$. Details are presented in Section 4. In Section 5 we show how to implement our algorithm for the Frobenius form in a processor-efficient manner on a PRAM. As a by-product we obtain the first processor-efficient parallel algorithm for the characteristic polynomial of a matrix.

The best previously known algorithms for finding the Frobenius form of a matrix in $\mathbf{K}^{n \times n}$, by Ozello (1987) and Lüneburg (1987), require $O(n^4)$ operations in \mathbf{K} (see also Kannan & Bachem 1979, and Kannan 1985). Kaltofen *et al.* (1987, 1990) present probabilistic algorithms for finding the Frobenius form of a matrix in the parallel complexity class RNC^2 . They achieve their result through the use of a more general algorithm to compute the Smith normal form of a polynomial matrix. From the Smith normal form of the matrix $\lambda I - T \in \mathbf{K}[\lambda]^{n \times n}$ (where λ is an indeterminate) the Frobenius form S of T can be derived easily. A (deterministic) NC^2 algorithm for computing the Frobenius form is demonstrated by Villard (1994).

Our algorithm for computing the Frobenius form is nearly optimal in that there is a lower bound for the problem of $\Omega(\text{MM}(n))$ operations in \mathbf{K} . If we can compute the Frobenius form, then we can find the characteristic polynomial $f \in \mathbf{K}[x]$ of T with $O(n^2)$ additional operations in \mathbf{K} . The determinant of T is $f(0)$, and it is shown in Baur & Strassen (1982) that computing the determinant requires $\Omega(\text{MM}(n))$ operations in \mathbf{K} , whence computing the Frobenius form of T requires $\Omega(\text{MM}(n))$ operations in \mathbf{K} (one must be careful of the model of computation here as Baur & Strassen's result is for the arithmetic circuit model – see Giesbrecht 1993, Section 1.1 for details).

As applications of our algorithm for computing the Frobenius form, we obtain fast algorithms for a number of interesting problems. One of the most striking results is for the problem of evaluating a polynomial at a matrix. In Section 6 we show that a polynomial $g \in \mathbf{K}[x]$, of degree r , can be evaluated at any matrix $T \in \mathbf{K}^{n \times n}$ with $O(\text{MM}(n) + r)$ operations in \mathbf{K} . We also demonstrate a lower bound for evaluating a fixed non-linear polynomial at a matrix of $\Omega(\text{MM}(n))$ operations in \mathbf{K} , so

in fact our algorithm is nearly optimal. The algorithm we present here improves upon the previously fastest algorithm of Paterson & Stockmeyer (1973), which requires $O(\text{MM}(n)\sqrt{r})$ operations in \mathbf{K} . More generally, we show that an arithmetic circuit or straight-line program can be evaluated at a matrix in nearly optimal time sequentially, and processor-efficiently in parallel. In brief, the Frobenius form provides a mechanism to transform the problem of evaluating a polynomial in $\mathbf{K}[x]$ at a matrix in $\mathbf{K}^{n \times n}$ from the multiplicative semi-group of $\mathbf{K}^{n \times n}$ to the multiplicative semi-group of a modular polynomial ring, where computation can be performed much more quickly.

As noted above, the companion matrix of the minimal polynomial forms the first block in the Frobenius form. Also, any two similar matrices have the same Frobenius form. Thus, our algorithm for computing the Frobenius form yields Las Vegas algorithms to determine the similarity of any two matrices in $\mathbf{K}^{n \times n}$, and to find the minimal polynomial of a matrix, which require $O^{\sim}(\text{MM}(n))$ operations in \mathbf{K} . Our algorithm for computing the minimal polynomial is also nearly optimal: it is shown by Wiedemann (see also Kaltofen 1992) that if we can compute the minimal polynomial of an $n \times n$ matrix over a field \mathbf{K} with $t(n)$ operations in \mathbf{K} , then there is a Las Vegas algorithm to compute the determinant of any matrix in $\mathbf{K}^{n \times n}$ for which $O(t(n))$ operations in \mathbf{K} are sufficient (see Giesbrecht 1993 for details). The best previously known algorithm for computing the minimal polynomial of an $n \times n$ matrix is the Monte Carlo algorithm of Wiedemann (1986), and requires an expected $O(n^3)$ operations (a Monte Carlo type algorithm has the ability to select random elements uniformly from a finite subset of \mathbf{K} and with constant probability returns the correct answer, but with some controllably small probability may return an incorrect answer). To determine similarity, the best previously known sequential algorithms are the Frobenius form algorithms of Lüneburg (1987) and Ozello (1987), which require $O(n^4)$ field operations. An algorithm by Zalcstein & Garzon (1987), based on a very different method, runs in the parallel complexity class NC^2 .

The Jordan normal form is probably the most commonly encountered of all canonical matrix forms, and is also one of the more difficult to compute. In particular, it requires that we determine both the geometric structure of the matrix (as captured in the Frobenius form) and the factorization of the minimal polynomial of the matrix into linear factors. In Section 7 we introduce the rational Jordan form as a slight generalization of the usual Jordan form. The rational Jordan form always exists and coincides with the usual Jordan form whenever that form exists. We show that computing the rational Jordan form of an $n \times n$ matrix over any field \mathbf{K} can be accomplished by a Las Vegas type algorithm with an expected number of $O^{\sim}(\text{MM}(n))$ operations in \mathbf{K} , given the complete factorization (into irreducible factors in $\mathbf{K}[x]$) of the minimal polynomial of that matrix. A simple lower bound of $\Omega(\text{MM}(n))$ operations in \mathbf{K} is shown for this problem. The requirement of a complete factorization of the minimal polynomial is also necessary, given that the complete factorization of any polynomial can be read off the rational Jordan form of the companion matrix of that polynomial.

Computational Model and Complexity Assumptions. The algorithms presented in this paper are generally given for both sequential and parallel models of computation. For the sequential algorithms we employ the arithmetic RAM, described more formally in von zur Gathen (1993). Informally, this is just a standard (Boolean) Random Access Machine (or RAM, see Aho *et al.* 1974, Section 1.2), with an additional memory for holding elements of some field \mathbf{K} . The instructions of the usual

RAM are supplemented with instructions for basic field operations in \mathbf{K} ($+$, \times , $-$, \div), as well as a test for the zero element in \mathbf{K} , and appropriate input and output instructions for elements in \mathbf{K} . We report the cost of our algorithms as the number of operations in \mathbf{K} required as a function of the number of elements of \mathbf{K} in the input.

For parallel algorithms we employ the arithmetic Parallel Random Access Machine (arithmetic PRAM) model over a field \mathbf{K} (see Karp & Ramachandran 1990). This is a collection of arithmetic RAM's (over \mathbf{K}) communicating by means of a set of shared memory cells. All accesses to global (and local) memory require unit time. The time $t(n)$ required by an arithmetic PRAM algorithm is the maximum of the number of field operations required by each of the component arithmetic RAM's in the PRAM on input size n . Also of interest is $p(n)$, the largest number of processors involved in the computation on any input of size n , as is the work $w(n) = t(n) \cdot p(n)$. It is useful to compare the work required by an algorithm running on a PRAM with the time required by an optimal sequential algorithm. Assume that any sequential algorithm for some problem requires time $\Omega(t(n))$ on input size n . A PRAM algorithm for this problem is called *processor-efficient* if it requires time $(\log n)^{O(1)}$ when run on $t(n) \cdot (\log n)^{O(1)}$ processors. That is, it is a fast parallel algorithm for which the work is within a poly-logarithmic factor of the optimal.

It is occasionally convenient, especially in summarizing results, to ignore logarithmic factors using the “soft O ” notation: for any $g, h: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$, $g = O^\sim(h)$ if and only if there exists a constant $k \geq 0$ such that $g = O(h(\log h)^k)$.

We isolate the cost of algorithms for polynomial arithmetic and linear algebra as named functions (M and MM , respectively) of their input sizes when their costs appear in the cost of some algorithm using these operations. Over a ring F , we assume $O(M_F(n))$ operations in F are sufficient to multiply two polynomials in $F[x]$ of degree at most n . The fast integer multiplication algorithms of Schönhage & Strassen (1971) can be recast as polynomial multiplication algorithms (see Schönhage 1977 and Nussbaumer 1980), and allow us to choose $M_F(n) = n \log n \log \log n$ when F is a field. More generally, for any ring F , the algorithm of Cantor & Kaltofen (1991) allows us to choose $M_F(n) = n \log n \log \log n$. If \mathbf{K} is a field, and $f, g \in \mathbf{K}[x]$ have degree at most n , we can compute the division with remainder of f by g , that is, we can find $Q, R \in \mathbf{K}[x]$ such that $f = Qg + R$ with $R = 0$ or $\deg R < \deg g$ in $O(M_F(n))$ operations in \mathbf{K} . We can compute $\gcd(f, g)$ with $O(M_F(n) \log n)$ operations in \mathbf{K} (see Aho *et al.* 1974, Section 8.9). Generally we will simply write $M(n)$ for $M_F(n)$ when F is clear from context.

For any field \mathbf{K} , let $\mathbf{K}^{n \times n}$ be the ring of $n \times n$ matrices with entries in \mathbf{K} . We assume that $O(MM(n))$ operations in \mathbf{K} are sufficient to multiply two matrices in $\mathbf{K}^{n \times n}$. Currently, the asymptotically best algorithm for matrix multiplication is by Coppersmith & Winograd (1990), with $MM(n) = n^{2.376}$. For convenience we assume throughout this paper that $MM(n) = \Omega(n^{2+\epsilon})$ for some $\epsilon > 0$. $O(MM(n))$ operations in \mathbf{K} are also sufficient to compute the rank and determinant of an $n \times n$ matrix over a field \mathbf{K} , as well as to invert a non-singular matrix in $\mathbf{K}^{n \times n}$. Also, we can solve a system of n linear equations in n unknowns over \mathbf{K} (which may be singular) with $O(MM(n))$ operations in \mathbf{K} (Bunch & Hopcroft 1974). Corresponding parallel algorithms for the polynomial and matrix operations above are discussed in Section 5.

1. Finding a Modular Cyclic Decomposition. The Frobenius normal form of the matrix $T \in \mathbf{K}^{n \times n}$ corresponds to a decomposition of the vector space $\mathbf{K}^{n \times 1}$, called the *cyclic decomposition* of $\mathbf{K}^{n \times 1}$ with respect to T . It is convenient to simply consider the linear map T as a \mathbf{K} -endomorphism of a finite dimensional vector space

V over \mathbf{K} (i.e., $T \in \text{End}_{\mathbf{K}} V$), ignoring for now its representation as a matrix (in the case of $T \in \mathbf{K}^{n \times n}$ we have $V = \mathbf{K}^{n \times 1}$). V is a $\mathbf{K}[x]$ -module as follows: any $g = \sum_{0 \leq i \leq r} b_i x^i \in \mathbf{K}[x]$ acts on V as the endomorphism $\sum_{0 \leq i \leq r} b_i T^i \in \text{End}_{\mathbf{K}} V$. We write $f \circ v = f(T)v \in V$. An important role is played in this theory by polynomials which annihilate (i.e., map to $0 \in V$) vectors and subspaces of V . For any $v \in V$, define the *annihilator* $\text{Ann}(T; v) \subseteq \mathbf{K}[x]$ of v as the ideal in $\mathbf{K}[x]$ of all $f \in \mathbf{K}[x]$ such that $f \circ v = 0$. Since $\mathbf{K}[x]$ is a principal ideal domain, $\text{Ann}(T; v)$ is generated by a unique monic polynomial in $\mathbf{K}[x]$, the *minimal polynomial* $\min(T; v) \in \mathbf{K}[x]$ of the vector $v \in V$. The subspace $\text{Orb}(T; v) \subseteq V$ spanned by $v, Tv, T^2v, \dots \in V$ is the *cyclic subspace* of V generated by v . $\text{Orb}(T; v)$ is T -invariant, that is, $Tw \in \text{Orb}(T; v)$ for all $w \in \text{Orb}(T; v)$, and $\dim \text{Orb}(T; v) = \deg(\min(T; v))$. For any subspace $W \subseteq V$, the set of polynomials $\text{Ann}(T; W) \subseteq \mathbf{K}[x]$ which annihilate *every* vector in W is also an ideal in $\mathbf{K}[x]$, called the *annihilator* of W . As an ideal, it too is generated by a unique monic polynomial, the *minimal polynomial* $\min(T; W) \in \mathbf{K}[x]$ of T on W .

The following very important theorem is shown in Hoffman & Kunze (1971), Section 7.2, Theorem 3, and Gantmacher (1990), Section 7.5.

FACT 1.1 (CYCLIC DECOMPOSITION THEOREM). *Let V be a finite dimensional vector space over a field \mathbf{K} and $T \in \text{End}_{\mathbf{K}} V$. There exist vectors $v_1, v_2, \dots, v_k \in V$ such that*

$$V = V_1 \oplus V_2 \oplus \dots \oplus V_k,$$

where $V_i = \text{Orb}(T; v_i)$, and

$$\min(T; v_i) = \min(T; V_i \oplus \dots \oplus V_k) \in \mathbf{K}[x]$$

for $1 \leq i \leq k$.

Given a cyclic decomposition of V under T as $V = V_1 \oplus \dots \oplus V_k$, where $V_i = \text{Orb}(T; v_i)$ for some $v_i \in V$ for $1 \leq i \leq k$, $f_i = \min(T; v_i) \in \mathbf{K}[x]$ is the *i th invariant factor* of T , for $1 \leq i \leq k$. Recall that the companion matrices of the invariant factors are found, in order, along the diagonal of the Frobenius form of T .

We once again consider T as a matrix in $\mathbf{K}^{n \times n}$, and the relationship between the cyclic decomposition of $\mathbf{K}^{n \times 1}$ under T and T 's Frobenius form. Suppose $\deg f_i = d_i$ for $1 \leq i \leq k$. The matrix

$$U = \left[\begin{array}{c|c|c|c|c|c|c|c|c|c} v_1 & Tv_1 & T^2v_1 & \dots & T^{d_1-1}v_1 & \dots & v_k & Tv_k & \dots & T^{d_k-1}v_k \end{array} \right] \in \mathbf{K}^{n \times n} \quad (1.2)$$

has the property that $S = U^{-1}TU$ is in Frobenius normal form (see, for example, Hoffman & Kunze 1971, Section 7.5).

Our algorithm for finding the Frobenius normal form S of a matrix $T \in \mathbf{K}^{n \times n}$ proceeds by constructing the matrix U as in (1.2), and then computing $S = U^{-1}TU$. To do this we must somehow find the invariant factors $f_1, \dots, f_k \in \mathbf{K}[x]$ and an associated set of vectors $v_1, \dots, v_k \in V$ generating the cyclic components. Unfortunately, such vectors are very rare, and the approach of choosing them randomly will fail miserably. Ozello (1987), and Lüneburg (1987) construct them deterministically, but using more time than we allow ourselves. Our approach will be to first choose randomly a list of vectors $w_1, \dots, w_k \in V$. With high probability these generate a “modular cyclic decomposition” of V (which we define later), from which the invariant factors and the Frobenius form S of T can be determined. The vectors w_1, \dots, w_k are then “purified” to obtain vectors v_1, \dots, v_k generating the components of the cyclic decomposition.

Our construction follows approximately the geometric proof of the Cyclic Decomposition Theorem by Gantmacher (1990), Chapter 7 (see also Hoffman & Kunze 1971, Section 7.2)

A weaker form of the cyclic decomposition theorem can be stated in terms of modular T -invariant vector spaces. If W is a T -invariant subspace of V then the space V/W , or V modulo W , is a \mathbf{K} -vector space of dimension $\dim V - \dim W$. For $w_1, w_2 \in V$ we denote by $w_1 + W$ and $w_2 + W$ the images of w_1 and w_2 in V/W respectively, and say $w_1 \equiv w_2 \pmod{W}$ when $w_1 - w_2 \in W$. Since W is T -invariant, V/W has a $\mathbf{K}[x]$ -module structure induced by the action of $\mathbf{K}[x]$ on V . Let $T_W: V/W \rightarrow V/W$ be T reduced modulo W , carrying $v + W$ to $Tv + W$ for any $v \in V$. Then V/W is a $\mathbf{K}[x]$ -module, where $f \in \mathbf{K}[x]$ acts as $f(T_W): V/W \rightarrow V/W$, and we write $f \circ (v + W) = f(T_W)(v + W) = f(T)v + W$. Thus, $\min(T_W; v + W) \in \mathbf{K}[x]$ is the minimal polynomial of $v + W \in V/W$ in $\mathbf{K}[x]$, that is, the monic $f \in \mathbf{K}[x] \setminus \{0\}$ of minimal degree such that $f \circ (v + W) \equiv 0 \pmod{W}$.

An important role is played in the decomposition theory of a vector space by *maximal* vectors $v \in V$, those satisfying $\min(T; v) = \min(T; V)$. A proof of the following lemma can be found in Gantmacher (1990), Section 7.4:

FACT 1.2. *Let V be any vector space over \mathbf{K} and $T: V \rightarrow V$ a linear map with invariant factors $f_1, \dots, f_k \in \mathbf{K}[x]$, where $f_i | f_{i-1}$ for $2 \leq i \leq k$. There exists a maximal $v_1 \in V$ such that $\min(T; v_1) = \min(T; V) = f_1$, and for all such v_1 there exists a T -invariant subspace $V' \subseteq V$ such that $V = \text{Orb}(T; v_1) \oplus V'$. Furthermore, the invariant factors of T on V' are f_2, \dots, f_k .*

The utility of this lemma comes in the realization that, for v_1 and V' as above, V' is isomorphic as a $\mathbf{K}[x]$ -module to $V/\text{Orb}(T; v_1)$. This follows since every element $w \in V$ can be written uniquely as $w = \bar{w} + w'$ for some $\bar{w} \in \text{Orb}(T; v_1)$ and $w' \in V'$, whence $w \equiv w' \pmod{\text{Orb}(T; v_1)}$. Applying Fact 1.2 recursively to V' yields a decomposition of V as follows (see Gantmacher 1990, Section 7.4).

FACT 1.3 (MODULAR CYCLIC DECOMPOSITION). *For $i \geq 1$ let $w_i \in V$ and define $V_i = \text{Orb}(T; w_1) + \dots + \text{Orb}(T; w_i)$ and $T_i: V/V_i \rightarrow V/V_i$ as T reduced modulo V_i (with $V_0 = \{0\}$ and $T_0 = T$). Assume that k is the smallest integer such that $V_k = V$, and that for $1 \leq i \leq k$ we have $\min(T_{i-1}; w_i + V_{i-1}) = \min(T_{i-1}; V/V_{i-1})$. Then $\min(T_{i-1}; w_i + V_{i-1}) = f_i$, the i th invariant factor of T on V , for $1 \leq i \leq k$.*

If w_1, \dots, w_k are as in the above theorem, then we say that they generate a *modular cyclic basis*

$$w_1, Tw_1, \dots, T^{d_1-1}w_1, w_2, Tw_2, \dots, T^{d_2-1}w_2, \dots, w_k, Tw_k, \dots, T^{d_k-1}w_k$$

for V . Of course the summation $V = \text{Orb}(T; w_1) + \dots + \text{Orb}(T; w_k)$ in Fact 1.3 is not direct, and once we have found w_1, \dots, w_k , we must somehow “purify” them to get a direct sum. This is accomplished in Section 3.

Our approach to finding $w_1, \dots, w_k \in V$ is to choose them randomly. This will only work when $\#\mathbf{K} \geq n^2$. When $\#\mathbf{K} < n^2$, we choose a small extension field \mathbf{F} of \mathbf{K} such that $\#\mathbf{F} \geq n^2$. This will be discussed in Section 4. For the remainder of this section, assume that $\#\mathbf{K} \geq n^2$.

LEMMA 1.4. *Let W be any vector space of dimension at most n over \mathbf{K} , spanned by $u_1, \dots, u_n \in W$, and let $T: W \rightarrow W$ be a \mathbf{K} -linear map. Let L be a subset of \mathbf{K} containing at least n^2 elements. Then*

$$\text{Prob}_{(a_1, \dots, a_n) \in L^n} \left\{ \min(T; \sum_{1 \leq i \leq n} a_i u_i) = \min(T; W) \right\} \geq 1 - \frac{1}{n}.$$

Proof. Let $y_1, \dots, y_l \in W$ be such that

$$W = \text{Orb}(T; y_1) \oplus \text{Orb}(T; y_2) \oplus \dots \oplus \text{Orb}(T; y_l)$$

is the cyclic decomposition of W with respect to T , where T has l invariant factors. For $1 \leq i \leq l$, let $g_i = \min(T; y_i) \in \mathbf{K}[x]$ be the i th invariant factor of T , so $g_i \mid g_{i-1}$ for $2 \leq i \leq l$.

We choose random elements $a_1, \dots, a_n \in L$ and assign $v = \sum_{1 \leq i \leq n} a_i u_i \in W$. In what follows we only consider the component of v in $\text{Orb}(T; y_1)$. Let $h_i \in \mathbf{K}[x]$ be such that the component of u_i in $\text{Orb}(T; y_1)$ is $h_i \circ y_1$, with $\deg h_i < \deg g_1$, for $1 \leq i \leq n$. The component of v in $\text{Orb}(T; y_1)$ is then $v_1 = \sum_{1 \leq i \leq n} a_i h_i \circ y_1$. Let $T_1: \text{Orb}(T; y_1) \rightarrow \text{Orb}(T; y_1)$ be the restriction of T to $\text{Orb}(T; y_1)$. Certainly $\min(T; v) = g_1$ when $\min(T_1; v_1) = g_1$; $\min(T; v)$ must annihilate *all* of v 's components. When is $\min(T_1; v_1) = g_1$? Let $h = \sum_{1 \leq i \leq n} a_i h_i$, so that $v_1 = h \circ y_1$. Clearly $g_1 \circ v_1 = (g_1 h) \circ y_1 = 0$ so $g = \min(T; v_1)$ divides g_1 . Also, $(gh) \circ y_1 = 0$, which is true only when $gh \equiv 0 \pmod{g_1}$. If $\gcd(h, g_1) = 1$ then $g \equiv 0 \pmod{g_1}$, whence $g = g_1$. Conversely, if $\gcd(h, g_1) \neq 1$ then g is a proper divisor of g_1 . In summary, $\min(T; v_1) = g_1$ if and only if $\gcd(h, g_1) = 1$.

To determine the probability that, for randomly chosen $a_1, \dots, a_n \in L$, we have $\gcd(\sum_{1 \leq i \leq n} a_i h_i, g_1) = 1$, consider the polynomial $\rho = \sum_{1 \leq i \leq n} x_i h_i \in \mathbf{K}[x_1, \dots, x_n][x]$, in the indeterminates x_1, \dots, x_n, x . To prove the theorem it is sufficient to show that

$$\text{Prob}_{(a_1, \dots, a_n) \in L^n} \left\{ \gcd(\rho(a_1, \dots, a_n)(x), g_1(x)) = 1 \right\} \geq 1 - \frac{1}{n}.$$

This is accomplished with the aid of resultants (see van der Waerden 1970, Section 5.8). The resultant of ρ and g_1 , considered as polynomials in x with coefficients in the integral domain $\mathbf{K}[x_1, \dots, x_n]$ is a polynomial $R \in \mathbf{K}[x_1, \dots, x_n]$ of degree at most n . Moreover, $\gcd(\rho(a_1, \dots, a_n)(x), g_1(x)) = 1$ if and only if $R(a_1, \dots, a_n) \neq 0$.

The polynomial R is non-zero as follows. The component of y_1 in $\text{Orb}(T; y_1)$ is y_1 itself, and hence has minimal polynomial g_1 . Assume $y_1 = \sum_{1 \leq j \leq n} b_j u_j$ for some $b_1, \dots, b_n \in \mathbf{K}$. Then $\sum_{1 \leq j \leq n} b_j h_j \equiv 1 \pmod{g_1}$, and since each h_j has degree less than g_1 by definition, $\sum_{1 \leq j \leq n} b_j h_j = 1$. Thus, $R(b_1, \dots, b_n) \neq 0$.

We now apply Corollary 1 of Schwartz (1980) to obtain

$$\text{Prob}_{(a_1, \dots, a_n) \in L^n} \left\{ R(a_1, \dots, a_n) \neq 0 \right\} \geq 1 - \frac{n}{\#L} \geq 1 - \frac{1}{n}.$$

□

A modular cyclic decomposition is now constructed by simply choosing the generating vectors “randomly” from V . Actually, since we do not know the number k of invariant factors beforehand, we choose $w_1, \dots, w_n \in V$ and consider the probability that the first k of these generate V in the desired manner.

THEOREM 1.5. *Let L be a subset of \mathbf{K} containing at least n^2 elements. For $1 \leq i \leq n$, randomly choose $w_i \in L^{n \times 1} \subseteq V$ and define $V_i = \text{Orb}(T; w_1) + \dots + \text{Orb}(T; w_i)$ and $T_i: V/V_i \rightarrow V/V_i$ as T reduced modulo V_i (with $V_0 = \{0\}$ and $T_0 = T$). With probability at least $1/4$, for all $1 \leq i \leq k$ (where k is the number of invariant factors of T) the i th invariant factor f_i of T satisfies $f_i = \min(T_{i-1}; w_i + V_{i-1}) = \min(T_{i-1}; V/V_{i-1})$.*

Proof. First we note, for $1 \leq i \leq n$, that V_i is a T -invariant subspace, so V/V_i is well defined as a $\mathbf{K}[T]$ -module. We show that for each $1 \leq i \leq n$ that the probability that

$\min(T_{i-1}; w_i + V_{i-1}) = \min(T_{i-1}; V/V_{i-1})$ is at least $1 - 1/n$. To see this, consider the standard basis $e_1, \dots, e_n \in \mathbb{K}^{n \times 1}$ for V , i.e., e_i is the vector with a one in the i th row and zeros in all other rows. We have chosen $w_i = \sum_{1 \leq i \leq n} a_i e_i$ for some randomly selected $a_1, \dots, a_n \in L$. The vectors $e_1 + V_{i-1}, \dots, e_n + V_{i-1}$ span V/V_{i-1} and

$$w_i + V_{i-1} = \sum_{1 \leq i \leq n} a_i e_i + V_{i-1} = \sum_{1 \leq i \leq n} a_i (e_i + V_{i-1}) \in V/V_{i-1}.$$

Thus, by Lemma 1.4,

$$\text{Prob}\left\{\min(T_{i-1}; w_i + V_{i-1}) = \min(T_{i-1}; V/V_{i-1})\right\} \geq 1 - 1/n.$$

The number of components k in the cyclic decomposition is certainly at most n , whence

$$\prod_{1 \leq i \leq k} \text{Prob}_{w_i \in L^{n \times 1}} \left\{\min(T_{i-1}; w_i + V_{i-1}) = \min(T_{i-1}; V/V_{i-1})\right\} \geq \left(1 - \frac{1}{n}\right)^n \geq \frac{1}{4}.$$

If all the w_i are chosen correctly then we also know, from Fact 1.3, that $f_i = \min(T_{i-1}; w_i + V_{i-1})$ for $1 \leq i \leq k$. \square

2. Computing the Invariant Factors of T on V . Assume now that we are given vectors $w_1, \dots, w_n \in V$ for V such that w_1, \dots, w_k generate a modular cyclic basis for V , and let V_i, T_i be as in Fact 1.3, for $0 \leq i \leq k$. How do we find k and the invariant factors f_1, \dots, f_k ? We adapt an algorithm of Keller-Gehrig (1985) to accomplish this with $O(\text{MM}(n) \log n)$ operations in \mathbb{K} .

FACT 2.1 (KELLER-GEHRIG 1985). *Let $T \in \mathbb{K}^{n \times n}$ and $u_1, \dots, u_n \in V$. Matrices $H_i \in \mathbb{K}^{n \times d_i}$, for $1 \leq i \leq n$ with the following properties can be computed with $O(\text{MM}(n) \log n)$ operations in \mathbb{K} :*

- (i) $\sum_{1 \leq i \leq n} d_i = n$, where $d_i \in \mathbb{N}$ may equal 0 for notational convenience;
- (ii) for $1 \leq i \leq n$, $H_i = [u_i | Tu_i] \cdots [T^{d_i-1} u_i]$ if $d_i > 0$;
- (iii) the columns of H_1, \dots, H_i form a basis of the $\mathbb{K}[x]$ -module generated by u_1, \dots, u_i , for $1 \leq i \leq n$ (that is, for $\text{Orb}(T; u_1) + \cdots + \text{Orb}(T; u_i)$).

Keller-Gehrig's (1985) algorithm is essentially an asymptotically fast version of the algorithm of Danilevsky (1937) for computing the characteristic polynomial of a matrix (see Faddeev & Faddeeva 1963). Applying Fact 2.1 to the vectors w_1, \dots, w_n yields k , bases for V_1, \dots, V_k , and the degrees of the invariant factors f_1, \dots, f_k . That is, the columns of H_1, \dots, H_i form a basis for $V_i = \text{Orb}(T; w_1) + \cdots + \text{Orb}(T; w_i)$ and $\deg f_i = d_i$ for $1 \leq i \leq k$. We know $d_i = 0$ for $k < i \leq n$ since w_1, \dots, w_k are assumed to generate V as a $\mathbb{K}[x]$ -module.

Let $H = [H_1 | H_2] \cdots [H_k] \in \mathbb{K}^{n \times n}$, which we call a *modular cyclic transition matrix* for T . Combining Fact 2.1 with Theorem 1.5 gives a probabilistic algorithm to find a modular cyclic transition matrix, which is correct with probability at least $1/4$. We summarize this algorithm below.

Algorithm: FindModCyc

Input: $T \in \mathbb{K}^{n \times n}$, where $\#\mathbb{K} \geq n^2$ and L is a subset of \mathbb{K} with $\#L \geq n^2$;

Output: a modular cyclic transition matrix $H = [H_1 | \cdots | H_k] \in \mathbb{K}^{n \times n}$ for T ,
 where $H_i = [w_i | \cdots | T^{d_i-1} w_i] \in \mathbb{K}^{n \times d_i}$ for some $w_i \in V$ ($1 \leq i \leq k$);

- (1) Choose w_1, \dots, w_n randomly from $L^{n \times 1}$.
- (2) Find $H = [H_1 | \cdots | H_k]$ such that $H_i = [w_i | \cdots | T^{d_i-1} w_i] \in \mathbb{K}^{n \times d_i}$, where
 - $\sum_{1 \leq i \leq k} d_i = n$
 - for $1 \leq i \leq k$, the columns of H_1, \dots, H_i form a basis of the $\mathbb{K}[x]$ -module generated by w_1, \dots, w_i ,
 using Keller-Gehrig's algorithm.

End.

THEOREM 2.2. *Let $\#\mathbb{K} \geq n^2$. Given $T \in \mathbb{K}^{n \times n}$, the algorithm **FindModCyc** returns $w_1, \dots, w_k \in V$ and matrices $H_i \in \mathbb{K}^{n \times d_i}$ for $1 \leq i \leq k$ defined as follows:*

- (i) $\sum_{1 \leq i \leq k} d_i = n$ with $d_i > 0$ for $1 \leq i \leq k$,
 - (ii) $H_i = [w_i | \cdots | T^{d_i-1} w_i]$ for $1 \leq i \leq k$,
 - (iii) The columns of H_1, \dots, H_i form a basis for $V_i = \text{Orb}(T; w_1) + \cdots + \text{Orb}(T; w_i)$,
 - (iv) $f_i = \min(T \bmod V_{i-1}; w_i + V_{i-1}) = \min(T \bmod V_{i-1}; V/V_{i-1})$ for $1 \leq i \leq k$,
- i.e., H is a modular cyclic transition matrix for T . The algorithm is probabilistic, and returns the correct answer with probability at least $1/4$ (it may produce an incorrect answer). It requires $O(\text{MM}(n) \log n)$ operations in \mathbb{K} .*

When H is a modular cyclic transition matrix it determines a change of basis on V under which T has the form

$$G = H^{-1}TH = \begin{pmatrix} \boxed{C_{f_1}} & \boxed{B_2} & \cdots & \\ & \boxed{C_{f_2}} & & \\ & & \ddots & \\ \mathbf{0} & & & \boxed{C_{f_k}} \end{pmatrix} \in \mathbb{K}^{n \times n}, \quad (2.3)$$

where $C_{f_i} \in \mathbb{K}^{d_i \times d_i}$ is the companion matrix of f_i , the i th invariant factor of T for $1 \leq i \leq k$ (see Keller-Gehrig 1985, Section 5). Each matrix $B_i \in \mathbb{K}^{t_i \times d_i}$, where $t_i = \sum_{1 \leq j < i} d_j$, is zero except for its last column; suppose this last column is

$$\vec{b}_i = (b_{i,1,0}, \dots, b_{i,1,d_1-1}, b_{i,2,0}, \dots, b_{i,2,d_2-1}, \dots, b_{i,i-1,0}, \dots, b_{i,i-1,d_{i-1}-1})^t \in \mathbb{K}^{t_i \times 1},$$

where $b_{ijl} \in \mathbb{K}$ for $2 \leq i \leq k$, $1 \leq j < i$, and $0 \leq l < d_j$. Under the change of basis induced by H , the vector $e_{t_i+j+1} \in \mathbb{K}^{n \times 1}$ (the column vector of all zeros except for a one in the $(t_i + j + 1)$ 'st row) in the basis given by the columns of H is the image of $T^j w_i$ in the standard basis for V .

The column \vec{b}_i in B_i and the last column $(-a_{i,0}, \dots, -a_{i,d_i-1})^t \in \mathbb{K}^{d_i \times 1}$ of C_{f_i} give the dependency of $T^{d_i} w_i$ on $T^l w_j$ for $1 \leq j \leq i$ and $0 \leq l < d_i$ by

$$Ge_{t_i+d_i} = H^{-1}T^{d_i} w_i = H^{-1} \sum_{0 \leq l < d_i} (-a_{il}) T^l w_i + H^{-1} \sum_{1 \leq j < i} \sum_{0 \leq l < d_j} c_{ijl} T^l w_j,$$

whence

$$T^{d_i} w_i + \sum_{0 \leq l < d_i} a_l T^l w_i = \sum_{1 \leq j < i} \sum_{0 \leq l < d_j} c_{ijl} T^l w_j,$$

or more simply

$$f_i \circ w_i = \sum_{1 \leq j < i} g_{ij} \circ w_j,$$

where $f_i = \sum_{0 \leq l < d_i} a_l x^l$, and $g_{ij} = \sum_{0 \leq l < d_j} c_{ijl} x^l$ with $a_l, c_{ijl} \in \mathbf{K}$ for $1 \leq j < i$ and $1 \leq i \leq k$. Thus $f_i \circ w_i \in V_{i-1}$, and since it has minimal degree, is the i th invariant factor of T on V . Note that by construction, $\sum_{1 \leq j < i} \deg g_{ij} \leq n$ for all $1 \leq i \leq k$, a fact which will be required in the next section.

THEOREM 2.3. *Let $w_1, \dots, w_n \in V$ be a basis for V such that w_1, \dots, w_k generate a modular cyclic basis for V , with $f_i = \min(T_{i-1}; w_i + V_{i-1})$ the i th invariant factor of T , for $1 \leq i \leq k$ and V_i, T_i as in Fact 1.3. Suppose also that we have computed a modular cyclic transition matrix for T . Then we can determine $k \in \mathbb{N}$, $f_1, \dots, f_k \in \mathbf{K}[x]$, and $g_{ij} \in \mathbf{K}[x]$ for $1 \leq j < i \leq k$ such that $f_i \circ w_i = \sum_{1 \leq j < i} g_{ij} \circ w_j$, with $O(\text{MM}(n))$ operations in \mathbf{K} . Furthermore, $\sum_{1 \leq j < i} \deg g_{ij} \leq n$ for all $1 \leq i \leq k$.*

Combining this theorem with Theorem 2.2 gives a Monte Carlo type probabilistic algorithm for computing the Frobenius form of T . Simply use the algorithm **FindModCyc** to find a modular cyclic transition matrix $H \in \mathbf{K}^{n \times n}$ for T . With probability at least $1/4$ the companion blocks of the the invariant factors $f_1, \dots, f_k \in \mathbf{K}[x]$ of T are along the diagonal of $H^{-1}TH$, from which we can construct the Frobenius form S of T . Of course, with positive probability we will get an erroneous Frobenius form, and since we do not get an invertible matrix $U \in \mathbf{K}^{n \times n}$ such that $S = U^{-1}TU$ from this procedure, we have no way of verifying that S is correct. This problem is addressed in the next section.

3. Purifying the Modular Decomposition. Once again, let $w_1, \dots, w_k \in V$ generate a modular cyclic basis for V , where V_i and T_i are as in Fact 1.3 for $0 \leq i \leq n$, and $f_i = \min(T_{i-1}; w_i + V_{i-1})$ is the i th invariant factor of T on V for $1 \leq i \leq k$. Furthermore, assume that $f_i \circ w_i = \sum_{1 \leq j < i} g_{ij} \circ w_j$ for some $g_{ij} \in \mathbf{K}[x]$ for all $1 \leq j < i \leq k$. This information can all be computed with the algorithm of the previous section (see Theorem 2.3), along with the matrix H defined there whose columns give a modular decomposition basis for V . In this section vectors $v_1, \dots, v_k \in V$ are constructed such that

$$V = \text{Orb}(T; v_1) \oplus \text{Orb}(T; v_2) \oplus \dots \oplus \text{Orb}(T; v_k),$$

and $f_i = \min(T; v_i)$, for $1 \leq i \leq k$. The key fact required is from Hoffman & Kunze (1971), Section 7.2, Theorem 3, Step 2:

FACT 3.1. *If $w_i \in V$, and $f_i, g_{ij} \in \mathbf{K}[x]$ are as above, then $f_i \mid g_{ij}$ for $1 \leq j < i \leq k$.*

Applying Fact 3.1, assume that $g_{ij} = f_i h_{ij}$ for some $h_{ij} \in \mathbf{K}[x]$, for all $1 \leq j < i \leq k$. We claim that the vectors $v_i = w_i - \sum_{1 \leq j < i} h_{ij} \circ w_j \in V$, for $1 \leq i \leq k$ are the ones desired.

THEOREM 3.2. *Let $w_1, \dots, w_k \in V$ and $V_i = \text{Orb}(T; w_1) + \dots + \text{Orb}(T; w_i)$ such that $V_k = V$ and $V_{k-1} \neq V$. Also, let $f_i = \min(T_{i-1}; w_i + V_{i-1})$ such that $f_i \circ w_i = \sum_{1 \leq j < i} g_{ij} \circ w_j$ and $g_{ij} = f_i h_{ij}$ for $g_{ij}, h_{ij} \in \mathbf{K}[x]$ with $1 \leq j < i \leq k$. If $v_i = w_i - \sum_{0 \leq j < i} h_{ij} \circ w_j$, then for $1 \leq i \leq k$,*

- (i) $f_i = \min(T; v_i)$,
- (ii) $V_i = \text{Orb}(T; v_1) \oplus \text{Orb}(T; v_2) \oplus \cdots \oplus \text{Orb}(T; v_i)$.

Proof. To show (i), note that for any i with $1 \leq i \leq k$ we have

$$\begin{aligned} f_i \circ v_i &= f_i \circ w_i - f_i \circ \left(\sum_{1 \leq j < i} h_{ij} \circ w_j \right) = f_i \circ w_i - \sum_{1 \leq j < i} f_i h_{ij} \circ w_j \\ &= f_i \circ w_i - \sum_{1 \leq j < i} g_{ij} \circ w_j = 0. \end{aligned}$$

Furthermore, since f_i is the polynomial of least degree such that $f_i \circ w_i \in V_{i-1}$, it follows that $f_i = \min(T; v_i)$.

To show (ii), we must show that $V_{i-1} \cap \text{Orb}(T; v_i) = \{0\}$, or equivalently that if $g \circ v_i \in V_{i-1}$ then $g \circ v_i = 0$ for any $g \in \mathbb{K}[x]$. Suppose $g \circ v_i \in V_{i-1}$ for some polynomial $g \in \mathbb{K}[x]$. This is true only if $g \circ w_i \in V_{i-1}$, since $v_i \equiv w_i \pmod{V_{i-1}}$. But $f_i = \min(T_{i-1}; w_i + V_{i-1})$, so $f_i \mid g$, and $g \circ v_i = 0$. Then $V_{i-1} \cap \text{Orb}(T; v_i) = \{0\}$ and $V_i = V_{i-1} \oplus \text{Orb}(T; v_i)$. \square

Note that this theorem does not assume that w_1, \dots, w_k generate a modular cyclic basis for V , only that they generate V as a $\mathbb{K}[x]$ -module.

To compute v_1, \dots, v_k , first compute $h_{ij} = g_{ij}/f_i \in \mathbb{K}[x]$ for all $1 \leq j < i \leq k$. For each $1 \leq i \leq k$ we note that $\sum_{1 \leq j < i} \deg g_{ij} \leq n$, so we can compute all g_{ij} for $1 \leq j < i$ with $O(M(n))$ operations in \mathbb{K} . The number k of invariant factors is at most n , so all the g_{ij} 's can be found with $O(nM(n))$ operations in \mathbb{K} .

Computing the h_{ij} 's gives a representation of the v_i 's in the modular decomposition basis for V given by the columns of H , not in the original (standard) basis for V . To find v_i in the standard basis, recall

$$v_i = w_i - \sum_{1 \leq j < i} h_{ij} \circ w_j = w_i - \sum_{1 \leq j < i} \sum_{0 \leq l < d_j} c_{ijl} T^l w_j.$$

We compute v_i by the matrix-vector product $v_i = H \bar{v}_i \in V$, where

$$\bar{v}_i = (-c_{i,1,1}, \dots, -c_{i,1,d_2-1}, \dots, -c_{i,i-1,1}, \dots, -c_{i,i-1,d_{i-1}}, 1, 0, \dots, 0)^t \in \mathbb{K}^{n \times 1}.$$

All of v_1, \dots, v_k can now be computed by a single matrix product

$$H \cdot [\bar{v}_1 \mid \cdots \mid \bar{v}_k] = [v_1 \mid \cdots \mid v_k],$$

with $O(MM(n))$ operations in \mathbb{K} .

THEOREM 3.3. *Given $w_i \in V$ and $f_i, g_{ij} \in \mathbb{K}[x]$ for $1 \leq i \leq k$ and $1 \leq j < i$ as above, we can compute $v_1, \dots, v_k \in V$ such that*

$$V = \text{Orb}(T; v_1) \oplus \cdots \oplus \text{Orb}(T; v_k)$$

and $f_i = \min(T; v_i)$ with $O(MM(n) + nM(n))$ operations in \mathbb{K} .

4. Computing the Frobenius Form. The complete algorithm for computing the Frobenius normal form of any $T \in \mathbb{K}^{n \times n}$, where $\#\mathbb{K} \geq n^2$, can now be stated. A modification which works over smaller fields is presented in the sequel.

Algorithm: FrobeniusForm

Input: $T \in \mathbb{K}^{n \times n}$, where $\#\mathbb{K} \geq n^2$;

Output: – The Frobenius form $S \in \mathbb{K}^{n \times n}$ of T ,
 – an invertible $U \in \mathbb{K}^{n \times n}$ such that $S = U^{-1}TU$;

- (1) Using **FindModCyc**, compute a modular cyclic transition matrix $H = [H_1 | \dots | H_k] \in \mathbb{K}^{n \times n}$, where
 - $H_i = [w_i | Tw_i | \dots | T^{d_i-1}w_i] \in \mathbb{K}^{n \times d_i}$ for $1 \leq i \leq k$,
 - the columns of H_1, \dots, H_i span the $\mathbb{K}[x]$ -module generated by the vectors w_1, \dots, w_i for $1 \leq i \leq k$;
 An erroneous computation of H will be detected in step (2) or (4).
- (2) If H is not invertible then quit, returning “failure”;
- (3) Compute $f_i, g_{ij} \in \mathbb{K}[x]$ from $G = H^{-1}TH$ such that $f_1 \circ w_1 = 0$ and $f_i \circ w_i = \sum_{1 \leq j < i} g_{ij} \circ w_j$ for $2 \leq i \leq k$ and $1 \leq j < i$;
- (4) If $f_i \nmid f_{i-1}$ for some $2 \leq i \leq k$, or $f_i \nmid g_{ij}$ for some $1 \leq j < i$ and $1 \leq i \leq k$, then return “failure” and quit;
- (5) Compute $h_{ij} = g_{ij}/f_i$ for $1 \leq j < i$ and $1 \leq i \leq k$;
- (6) Compute $v_i = w_i - \sum_{1 \leq j < i} h_{ij} \circ w_j$ for $1 \leq i \leq k$;
- (7) Compute the matrix $U \in \mathbb{K}^{n \times n}$ as in (1.2);
- (8) Output the Frobenius form $S = \text{diag}(C_{f_1}, \dots, C_{f_k})$ and the transition matrix U (where C_{f_i} is the companion matrix of f_i for $1 \leq i \leq k$).

End.

THEOREM 4.1. *Let $T \in \mathbb{K}^{n \times n}$ over any field \mathbb{K} with at least n^2 elements. With probability at least $1/4$, **FrobeniusForm** returns the Frobenius form S of T and an invertible $U \in \mathbb{K}^{n \times n}$ such that $S = U^{-1}TU$. Otherwise the algorithm reports “failure”. In either case $O(\text{MM}(n) \log n + nM(n))$ operations in \mathbb{K} are sufficient.*

Proof. First we prove the correctness of the output. We employ the algorithm **FindModCyc**, analyzed in Theorem 2.2, to compute a modular cyclic transition matrix H , which will be correct with probability at least $1/4$. Assume for now that H is computed correctly. By Fact 3.1 we know $f_i \mid g_{ij}$ for $1 \leq j < i$ and $1 \leq i \leq k$, and certainly $f_i \mid f_{i-1}$ for $2 \leq i \leq k$. Theorem 3.2 guarantees that v_1, \dots, v_k generate the cyclic composition factors of V and Gantmacher (1990), Section 7.5.2, shows that $S = U^{-1}TU$ is in Frobenius normal form.

If H is incorrectly computed, the algorithm reports failure in step (2) or (4). Certainly if H is singular then it is incorrect, and this is detected in step (2). Otherwise, assume that the tests in step (4) are passed, but w_1, \dots, w_k do not generate a modular cyclic basis. By Theorem 3.2, $V = \text{Orb}(T; v_1) \oplus \dots \oplus \text{Orb}(T; v_k)$. Let $l \geq 1$ be the smallest integer such that $f_l \neq \bar{f}_l$, where \bar{f}_l is the the actual l th invariant factor of T on V . In the vector space

$$\bar{V}_l = \text{Orb}(T; v_l) \oplus \dots \oplus \text{Orb}(T; v_k) \cong V/V_{l-1}$$

there exists a vector $w \in \bar{V}_l$ with $\min(T_{l-1}; w + V_{l-1}) = \bar{f}_l$ (where $T_{l-1}: V/V_{l-1} \rightarrow V/V_{l-1}$ is T reduced modulo V_{l-1}). Since $f_j \mid f_l$ for $l \leq j \leq k$, we know $f_l = \min(T_{l-1}; \bar{V}_l) = \bar{f}_l$, a contradiction.

By Theorem 2.2, step (1) requires $O(\text{MM}(n) \log n)$ operations in \mathbf{K} . Using Theorem 2.3, steps (2)-(4) require $O(\text{MM}(n))$ operations in \mathbf{K} (H is singular exactly when its last column is zero). By Theorem 3.3 steps (5)-(6) require $O(\text{MM}(n))$ operations in \mathbf{K} to complete. The matrix U in step (7) can be found using Fact 2.1, with $O(\text{MM}(n) \log n)$ operations in \mathbf{K} . \square

The algorithm **FrobeniusForm** requires that the field \mathbf{K} contain at least n^2 elements. If this is not the case, and $q = \#\mathbf{K} < n^2$, the chances of choosing w_1, \dots, w_k correctly may be very low. To remedy this we construct a field extension \mathbf{F} of \mathbf{K} containing n^2 elements. We then run the algorithm **FrobeniusForm** on $T \in \mathbf{F}^{n \times n}$. By Theorem 4.1 this correctly returns the Frobenius form $S \in \mathbf{F}^{n \times n}$ and an invertible $U \in \mathbf{F}^{n \times n}$ such that $S = U^{-1}TU$, with probability at least $1/4$. However, the Frobenius form of T is a unique rational invariant of T , regardless of the field in which the elements of T are embedded. Thus $S \in \mathbf{K}^{n \times n}$ and this is precisely what we are looking for. The only drawback to this technique, aside from the slightly increased cost of operations in \mathbf{F} , is that U is not necessarily a matrix over \mathbf{K} .

To construct the field $\mathbf{F} \supseteq \mathbf{K}$, we use the deterministic algorithm of Shoup (1993) to find a polynomial $\psi \in \mathbf{K}[x]$ of degree $\lceil 2 \log_q n \rceil$. Shoup's algorithm requires $O^\sim(n)$ operations in \mathbf{K} . Let $\mathbf{F} = \mathbf{K}[x]/(\psi)$, a finite field with at least n^2 elements, where each element is represented by a polynomial of degree less than $\lceil 2 \log_q n \rceil$. Elements of \mathbf{K} are simply represented as constant polynomials, giving a trivial embedding of \mathbf{K} into \mathbf{F} . Additions and subtractions in \mathbf{F} require $O(\log_q n)$ operations in \mathbf{K} , while multiplications in \mathbf{F} require $O(M(\log_q n))$ operations in \mathbf{K} and divisions in \mathbf{F} require $O(M(\log_q n) \log \log_q n)$ operations in \mathbf{K} .

THEOREM 4.2. *Let $T \in \mathbf{K}^{n \times n}$, where $q = \#\mathbf{K} < n^2$. The modification of **FrobeniusForm** described above computes the Frobenius form $S \in \mathbf{K}^{n \times n}$ of T , and an invertible matrix $U \in \mathbf{F}^{n \times n}$ such that $S = U^{-1}TU$. The field $\mathbf{F} = \mathbf{K}[x]/(\psi)$ is an extension of \mathbf{K} , given by an irreducible monic $\psi \in \mathbf{K}[x]$ of degree $\lceil 2 \log_q n \rceil$. The algorithm succeeds with probability at least $1/4$, and otherwise reports "failure". It requires $O((\text{MM}(n) \log n + nM(n)) \cdot M(\log_q n) \log \log_q n)$ or $O^\sim(\text{MM}(n))$ operations in \mathbf{K} .*

5. A Processor-Efficient Parallel Algorithm. In this section we exhibit a processor-efficient parallel algorithm for finding the Frobenius form of a matrix. Recall that computing the Frobenius form of a $T \in \mathbf{K}^{n \times n}$ is at least as hard as matrix multiplication, since from the Frobenius form we can quickly compute the determinant of T , which is known to be as hard as multiplying matrices (see Baur & Strassen, 1982). We demonstrate that our algorithm for computing the Frobenius form of any $T \in \mathbf{K}^{n \times n}$, can be implemented in a processor-efficient manner, and can be executed in (Las Vegas) time $(\log n)^{O(1)}$ using $O(\text{MM}(n))$ processors.

Actually, the only part of the algorithm **FrobeniusForm** which requires modification, other than specifying an implementation technique for each step from the "toolkit" of known parallel algorithms, is the subroutine **FindModCyc** used in step (1). The routine of Keller-Gehrig's (1985), used in step (2) of **FindModCyc** in Section 2, is the heart of his asymptotically fast version of Danilevsky's (1937)'s algorithm for finding the characteristic polynomial of a matrix. It is not immediately clear that Keller-Gehrig's algorithm can be implemented in a processor-efficient manner. In this section we exhibit fast parallel versions of Keller-Gehrig's algorithm for computing the characteristic polynomial of a matrix and of our routine **FindModCyc**, both of which are processor-efficient. We then give the implementation details of a processor-efficient version of **FrobeniusForm**.

We begin by collecting a number of useful parallel algorithms which we require as subroutines.

FACT 5.1. *Suppose \mathbf{K} is a field with characteristic p .*

- (i) *Suppose that there is a bilinear algorithm which, given $A, B \in \mathbf{K}^{n \times n}$ computes AB with $O(\text{MM}(n))$ operations in \mathbf{K} sequentially (for a definition of bilinear algorithm see Borodin & Munro 1975), then there exists a processor-efficient parallel algorithm to compute AB on $O(\text{MM}(n)n^\epsilon)$ processors in time $O(\log n)$, for any $\epsilon > 0$.*
- (ii) *Suppose $\#\mathbf{K} \geq n$ and $A \in \mathbf{K}^{n \times n}$ is invertible. If $p \geq n$ or $p = 0$ then we can compute A^{-1} with $O(\text{MM}(n))$ processors in time $O(\log^2 n)$. If $2 \leq p \leq n$ there is a Las Vegas type probabilistic algorithm to compute A^{-1} with time $O(\log^3 n / \log p)$ on $O(\text{MM}(n))$ processors.*
- (iii) *If $\#\mathbf{K} \geq n$ then given vectors $v_1, \dots, v_m \in \mathbf{K}^{n \times 1}$, where $m \leq 2n$, we can find the lexicographically first subset of v_1, \dots, v_m which form a basis for the vector space spanned by v_1, \dots, v_m . If $p \geq n$ or $p = 0$ this can be accomplished with $O(\log^4 n)$ time on $O(\text{MM}(n) / \log n)$ processors. If $2 \leq p < n$ this can be done with a Las Vegas algorithm requiring time $O(\log^5 n / \log p)$ on $O(\text{MM}(n) / \log n)$ processors.*
- (iv) *Given two polynomials $f, g \in \mathbf{K}[x]$ of degree at most n , we can compute fg in time $O(\log n)$ with $O(n \log \log n)$ processors.*
- (v) *Given two polynomials $f, g \in \mathbf{K}[x]$ of degree at most n , we can compute the unique $Q, R \in \mathbf{K}[x]$ such that $\deg R < \deg g$ and $f = Qg + R$ in time $O(\log n)$ on $O(n \log n)$ processors.*

Proof. Part (i) is shown by Pan & Reif (1985) Theorem A.1. A processor-efficient algorithm for matrix inversion is given by Kaltofen & Pan (1991) and Kaltofen & Pan (1992). Eberly (1991) shows (iii). The fast polynomial multiplication algorithms of Cantor & Kaltofen (1991) also have parallel time and processor bounds as stated in Part (iv). Part (v) is shown by Bini & Pan (1992). \square

We make the assumption, based on Fact 5.1(i), that there exists an algorithm which requires time $O(\log n)$ on $O(\text{MM}(n))$ processors to multiply two $n \times n$ matrices over a field \mathbf{K} .

Assume that we are given vectors $u_1, \dots, u_n \in V$. As in Fact 2.1, we want to find matrices $H_i \in \mathbf{K}^{n \times d_i}$ for $1 \leq i \leq n$ with the following properties:

- (i) $\sum_{1 \leq i \leq n} d_i = n$, where $d_i \in \mathbb{N}$ may equal 0 for notational convenience;
- (ii) for $1 \leq i \leq n$, $H_i = [u_i \mid Tu_i \mid \dots \mid T^{d_i-1}u_i]$ if $d_i > 0$;
- (iii) the columns of H_1, \dots, H_i form a basis of the $\mathbf{K}[x]$ -module generated by u_1, \dots, u_i , for $1 \leq i \leq n$.

We compute $H = [H_1 \mid \dots \mid H_n]$ by computing matrices $H^{(0)}, H^{(1)}, \dots, H^{(r)} \in \mathbf{K}^{n \times n}$ in sequence, where $r = \lceil \log_2 n \rceil$, $H^{(0)} = T$ and $H^{(r)} = H$. At step 0 we have $H^{(0)} = [H_1^{(0)} \mid \dots \mid H_n^{(0)}]$, where $H_i^{(0)}$ is the i th column of T for $1 \leq i \leq n$. The idea is that at the end of stage j we will have computed $H^{(j)} = [H_1^{(j)} \mid \dots \mid H_n^{(j)}]$, where

$$H_i^{(j)} = [u_i \mid Tu_i \mid \dots \mid T^{d_i^{(j)}-1}u_i] \in \mathbf{K}^{n \times d_i^{(j)}}, \quad \text{for } 1 \leq i \leq n.$$

Here $d_i^{(j)}$ is the smallest integer less than 2^j such that $T^{d_i^{(j)}}u_i$ lies in the vector space spanned by $u_i, Tu_i, \dots, T^{d_i^{(j)}-1}u_i$ along with all the columns of $H_1^{(j)}, \dots, H_{i-1}^{(j)}$. If

$u_i, Tu_i, \dots, T^{d_i^{(j)}} u_i$ and the columns of $H_1^{(j)}, \dots, H_{i-1}^{(j)}$ are linearly independent then $d_i^{(j)} = 2^j$. This can be expressed more concisely in the language of modular vector spaces. Let $V_{i-1}^{(j)}$ be the vector space spanned by the columns of $H_1^{(j)}, \dots, H_{i-1}^{(j)}$. Then $d_i^{(j)}$ is the smaller of 2^j and $\dim \text{Orb}(T \bmod V_{i-1}^{(j)}; u_i \bmod V_{i-1}^{(j)})$. Recall that we allow $d_i^{(j)}$ to be zero, meaning that u_i is linearly dependent on the columns of $H_1^{(j)}, \dots, H_{i-1}^{(j)}$. Since $2^r \geq n$, it is clear that $H^{(r)}$ will have the desired properties.

To find $H^{(j)}$ after computing $H^{(j-1)}$ (for $1 \leq j \leq r$), we first compute matrices $J_1^{(j)}, \dots, J_n^{(j)}$ as follows:

$$J_i^{(j)} = \begin{cases} H_i^{(j-1)} & \text{if } H_i^{(j-1)} \text{ has fewer than } 2^j \text{ columns,} \\ [H_i^{(j-1)} \mid T^{2^j} H_i^{(j-1)}] & \text{otherwise.} \end{cases}$$

This doubles the number of iterates of u_i under T if this is required. Note that if $T^l u_i$ is linearly dependent upon V_i together with $u_i, Tu_i, \dots, T^{l-1} u_i$ then so is $T^{l+1} u_i$. Thus, to find $d_1^{(j)}, \dots, d_n^{(j)}$ we identify and mark the lexicographically first set of n linearly independent columns of

$$J^{(j)} = [J_1^{(j)} \mid \dots \mid J_n^{(j)}].$$

The last column marked in $J_i^{(j)}$ is indexed by $d_i^{(j)}$. If none of the identified columns are in $J_i^{(j)}$ for some i then $d_i^{(j)} = 0$. Now let $H_i^{(j)} \in \mathbb{K}^{n \times d_i}$ consist of those columns of $J_i^{(j)}$ marked in this process.

This can be implemented in a processor-efficient manner using Fact 5.1. At stage j we compute $T^{2^j} = T^{2^{j-1}} \cdot T^{2^{j-1}}$, then we compute $J^{(j)}$ as described above from the product $T^{2^j} H^{(j-1)}$. The lexicographically first subset of the columns of $J^{(j)}$ are identified with the algorithm of Eberly (1991) (Fact 5.1(iii) in this paper), and $H^{(j)} \in \mathbb{K}^{n \times n}$ is constructed as above. This is repeated $r = \lceil \log_2 n \rceil$ times.

THEOREM 5.2. *If $p \geq n$ or $p = 0$ we can compute $H \in \mathbb{K}^{n \times n}$ and $d_1, \dots, d_n \in \mathbb{N}$ as in Fact 2.1 in time $O(\log^4 n)$ on $O(\text{MM}(n))$ processors. If $2 \leq p \leq n$ this algorithm requires time $O(\log^5 n / \log p)$ on $O(\text{MM}(n))$ processors.*

Combining Theorem 5.2 with Theorem 1.5 allows us to implement the algorithm **FindModCyc** in Section 2 in a processor-efficient manner.

THEOREM 5.3. *Let $\#\mathbb{K} \geq n^2$. Given $T \in \mathbb{K}^{n \times n}$, the algorithm **FindModCyc** implemented in a processor-efficient manner as described above, returns $w_1, \dots, w_k \in V$ and matrices $H_i \in \mathbb{K}^{n \times d_i}$ for $1 \leq i \leq k$ defined as follows:*

- (i) $\sum_{1 \leq i \leq k} d_i = n$ with $d_i > 0$ for $1 \leq i \leq k$,
- (ii) $H_i = [w_i \mid \dots \mid T^{d_i-1} w_i]$ for $1 \leq i \leq k$,
- (iii) The columns of H_1, \dots, H_i form a basis for $V_i = \text{Orb}(T; w_1) + \dots + \text{Orb}(T; w_i)$,
- (iv) $f_i = \min(T \bmod V_{i-1}; w_i + V_{i-1}) = \min(T \bmod V_{i-1}; V/V_{i-1})$ for $1 \leq i \leq k$.

The algorithm is probabilistic, and returns the correct answer with probability at least 1/4 (it may produce an incorrect answer). If $p \geq n$ or $p = 0$ it requires time $O(\log^4 n)$ on $O(\text{MM}(n))$ processors. If $2 \leq p \leq n$ it requires time $O(\log^5 n / \log p)$ on $O(\text{MM}(n))$ processors.

We now examine the parallel cost of the algorithm **FrobeniusForm**.

THEOREM 5.4. *The Las Vegas algorithm **FrobeniusForm** can be implemented in parallel in a processor-efficient manner. Let $T \in \mathbb{K}^{n \times n}$, where \mathbb{K} is a field with $\#\mathbb{K} \geq n^2$.*

- (i) If $p \geq n$ or $p = 0$ then **FrobeniusForm** requires time $O(\log^4 n)$ on $O(\text{MM}(n))$ processors.
- (ii) If $2 \leq p < n$ then **FrobeniusForm** requires time $O(\log^5 n / \log p)$ on $O(\text{MM}(n))$ processors.

Proof. The statement of the parallel algorithm does not vary from the description in Section 4. Only the manner in which each step is executed changes, and we address these changes now. Step (1) is accomplished using the processor-efficient implementation of **FindModCyc** described in Theorem 5.3. The matrix H is singular exactly when its last column is zero, and this can be tested with constant time and work to perform step (2). As in the sequential algorithm, the coefficients of f_i and g_{ij} for $2 \leq i \leq k$ and $1 \leq j < i$ are read from $G = H^{-1}TH$ as in Theorem 2.3, and G can be computed within the desired parallel time and processor bounds by Fact 5.1. Steps (4)–(5) can be done using processor-efficient polynomial division algorithms, as specified in Fact 5.1 parts (iv) and (v). Step (6) can be performed by a single matrix product as in Theorem 3.3, and so can be done in a processor-efficient manner. Applying Theorem 5.3 gives us U in step (7). \square

When $\#\mathbf{K} < n^2$ similar theorems hold, with running times and processor requirements multiplied by a small power of $\log_q n$.

THEOREM 5.5. *Let \mathbf{K} be a field with $q < n^2$ and $T \in \mathbf{K}^{n \times n}$. We can implement the algorithm **FrobeniusForm**, as modified in Theorem 4.2, in a processor-efficient manner.*

- (i) If $p \geq n$ and $p = 0$ it requires time $O(\log^4 n \log \log_q n)$ on $O(\text{MM}(n) \cdot \text{M}(\log_q n))$ processors.
- (ii) If $2 \leq p < n$ it requires time $O(\log^5 n \log \log_q n / \log p)$ on $O(\text{MM}(n) \cdot \text{M}(\log_q n))$ processors.

Proof. The proof is the same as that of Theorem 5.4 except that we work in an extension \mathbf{F} of \mathbf{K} of algebraic degree $\lceil \log_q n \rceil$ over \mathbf{K} , which does contain $O(n^2)$ elements (see Theorem 4.2). Each operations in \mathbf{F} requires $O(\log \log_q n)$ operations in \mathbf{K} on $O(\text{M}(\log_q n))$ processors. \square

Theorem 5.3 also has the following important corollary, essentially a processor-efficient version of the algorithms of Danilevsky (1937) and Keller-Gehrig (1985) for computing the characteristic polynomial of a matrix.

COROLLARY 5.6. *There is a processor-efficient Las Vegas algorithm for computing the characteristic polynomial $f \in \mathbf{K}[x]$ of any matrix $T \in \mathbf{K}^{n \times n}$ over any field \mathbf{K} . If $p \geq n$ or $p = 0$ we can compute f in time $O(\log^4 n)$ on $O(\text{MM}(n))$ processors. If $2 \leq p < n$ this algorithm requires time $O(\log^5 n / \log p)$ on $O(\text{MM}(n))$ processors.*

Proof. We first find a matrix $H \in \mathbf{K}^{n \times n}$ as in Theorem 5.2 such that $G = H^{-1}TH$ is as in (2.3). The inverse of H and product $f_1 \cdots f_k$ can be computed within the desired time and processor bounds using the methods described in Fact 5.1. \square

Note that the characteristic polynomial of $T \in \mathbf{K}^{n \times n}$ can also be computed as the product of the invariant factors of T , which can be read from the Frobenius form of T . The above corollary is a slight improvement on this for small fields. It relies on the fact that we do not need $H^{-1}TH$ to be in Frobenius form; the product of the polynomials whose companion matrices are on the diagonal of $H^{-1}TH$ always equals the characteristic polynomial of T .

6. Fast Evaluation of Matrix Functions. We now consider applications of our algorithm for finding the Frobenius form to evaluating polynomials at matrices, determining matrix similarity and finding the minimal polynomial of matrices. New algorithms requiring almost optimal time are presented. We also give lower bounds for the problems of evaluating a polynomial at matrix, and for finding the minimal polynomial of a matrix, which matches our upper bounds within poly-logarithmic factors.

Evaluating Polynomials at Matrices. Computing the Frobenius form $S \in \mathbb{K}^{n \times n}$ of a matrix $T \in \mathbb{K}^{n \times n}$ allows quick evaluation of $g(T) \in \mathbb{K}^{n \times n}$ for any polynomial $g \in \mathbb{K}[x]$. The problem of evaluating polynomials at matrices is certainly not new. Paterson & Stockmeyer (1973) give an algorithm to evaluate a polynomial $g \in \mathbb{K}[x]$ at any point in a ring extension \mathfrak{R} of \mathbb{K} , which requires $O(\sqrt{r})$ non-scalar multiplications in \mathfrak{R} , where $r = \deg g$. They apply their algorithm to evaluate $g(T)$ at any $T \in \mathbb{K}^{n \times n}$ with $O(\text{MM}(n)\sqrt{r})$ operations in \mathbb{K} (see also Brent & Kung 1978).

Theorem 6.2 gives a substantial improvement over this, especially when r is large, allowing evaluation of $g(T)$ with $O(\text{MM}(n) + r)$ operations in \mathbb{K} . The well known observation required is that if $S = U^{-1}TU$ for some invertible $U \in \mathbb{K}^{n \times n}$, and $S = \text{diag}(C_{f_1}, \dots, C_{f_k})$ as in (0.1), then

$$g(T) = U \cdot g(S) \cdot U^{-1} = U \cdot \text{diag}(g(C_{f_1}), \dots, g(C_{f_k})) \cdot U^{-1}.$$

Furthermore, each $g(C_{f_i})$ can be evaluated quickly once $g \bmod f_i$ has been computed, and $g(S)$ can be computed with $O(n^2)$ operations in \mathbb{K} when $\deg g \leq n$.

LEMMA 6.1. *Let $C \in \mathbb{K}^{s \times s}$ be the companion matrix of $h \in \mathbb{K}[x]$, where $\deg h = s$. Also, let $g \in \mathbb{K}[x]$ have degree less than s . Then we can compute $g(C)$ with $O(s^2)$ operations in \mathbb{K} , or in parallel time $O(\log s)$ on $O(sM(s))$ processors.*

Proof. The vector space $\mathbb{K}^{s \times 1}$ has a $\mathbb{K}[x]$ -module structure, where any $f \in \mathbb{K}[x]$ acts on $\mathbb{K}^{s \times 1}$ as $f(C) \in \mathbb{K}^{s \times s}$. As a $\mathbb{K}[x]$ -module, $\mathbb{K}^{s \times 1}$ is isomorphic to $\mathbb{K}[x]/(h)$ — simply map $v = (b_0, \dots, b_{s-1}) \in \mathbb{K}^{s \times 1}$ to $\varphi(v) = \sum_{0 \leq i < s} b_i x^i \bmod h$. It is easy to show that $x\varphi(v) \equiv \varphi(Cv) \bmod h$, and therefore that $f \cdot \varphi(v) = \varphi(f(C)v)$ for any $f \in \mathbb{K}[x]$. In particular, $\varphi(e_i) = x^{i-1}$, where $e_i \in \mathbb{K}^{s \times 1}$ is zero except for a one in the i th row, and hence $\varphi(g(C)e_i) \equiv g \cdot x^{i-1} \bmod h$, for $1 \leq i \leq s$. Since $g(C)e_i$ is the i th column of $g(C)$, we can find $g(C)$ by computing $g \cdot x^i \bmod h$ for $0 \leq i < s$. This can be accomplished with $O(s^2)$ operations in \mathbb{K} by realizing that computing $C \cdot v$ for any $v \in \mathbb{K}^{s \times 1}$, and hence multiplying by x modulo h , requires only $O(s)$ operations in \mathbb{K} (C has only $2s$ non-zero entries). It can be done in time $O(\log s)$ on $O(sM(s))$ processors using the polynomial multiplication algorithm of Cantor & Kaltofen (1991). \square

THEOREM 6.2. *Given $g \in \mathbb{K}[x]$ of degree r and $T \in \mathbb{K}^{n \times n}$, we can compute $g(T) \in \mathbb{K}^{n \times n}$ with a Las Vegas type probabilistic algorithm requiring $O(\text{MM}(n) \log n + nM(n) + M(r))$ operations in \mathbb{K} . Here \mathbb{K} is any field with at least n^2 elements.*

Proof. First compute the Frobenius form $S \in \mathbb{K}^{n \times n}$ of T , and an invertible $U \in \mathbb{F}^{n \times n}$ such that $U^{-1}TU = S$. This requires $O(\text{MM}(n) \log n + nM(n))$ operations in \mathbb{K} using Theorem 4.1. For each invariant factor $f_i \in \mathbb{K}[x]$ of T , determine $g_i \equiv g \bmod f_i$, where $\deg g_i < \deg f_i$, for $1 \leq i \leq k$. This can be done with $O(M(r) + nM(n))$ operations in \mathbb{K} by first reducing g modulo f_1 , since all the other invariant factors divide f_1 . Next compute the matrices $g_i(C_{f_i}) = g(C_{f_i})$ for $1 \leq i \leq k$. Using Lemma 6.1 above, this can be accomplished with $O(n^2)$ operations in \mathbb{K} . Now $g(S) = \text{diag}(g(C_{f_1}), \dots, g(C_{f_k}))$ and $g(T) = Ug(S)U^{-1}$. \square

A more general theorem can be obtained using arithmetic circuits or straight-line programs (see Aho *et al.* 1974, Section 1.5 or von zur Gathen 1986 for a formal definition). This model gives a very natural computational representation of rational functions, nicely formalizing and generalizing such notions as sparseness and ease of computation. Briefly, an algebraic circuit or straight-line program φ over a field \mathbf{K} consists of a finite list s_1, s_2, \dots, s_m of statements, each of the form $a_i := b_i \text{ op}_i c_i$, where a_i is a variable (assigned only at step i), $\text{op}_i \in \{+, -, \times, /\}$, and each of b_i, c_i is either equal to a_j for some $1 \leq j < i$, or is a constant in \mathbf{K} , or is the *input* x . The output is the value of a_m upon evaluation of each of the steps in sequence (this semantics glosses over a number of subtleties — see von zur Gathen 1986). The sequential cost of φ is m . To define the parallel cost, note that φ defines a directed acyclic graph on the nodes x, s_1, \dots, s_k , where there is an edge from s_i to s_j (or x to s_j) if s_j references a_i (resp. x). The parallel cost is the depth, or maximum path length from the node x to the node s_m .

Now let φ be algebraic circuit over \mathbf{K} . Let F be an extension ring of \mathbf{K} , so φ can be thought of as computing a function from $F \rightarrow F$. If φ has length m and depth d , then we can evaluate φ at a point $a \in F$ on a sequential arithmetic RAM over F with $O(m)$ operations in F , or on an arithmetic PRAM over F in time $O(d)$ on $O(m)$ processors. We obtain for straight-line programs a theorem similar to Theorem 6.2.

THEOREM 6.3. *Let φ be a straight-line program over \mathbf{K} as described above, with length m and depth d . Assume that \mathbf{K} has characteristic p and $\#\mathbf{K} \geq n^2$. For any $T \in \mathbf{K}^{n \times n}$, we can evaluate φ at T in time as follows:*

- (i) *with $O(\text{MM}(n) \log n + (n + m)\text{M}(n))$ operations in \mathbf{K} sequentially;*
- (ii) *in parallel time $O(\log^4 n + d \log n)$ when $p \geq n$ or $p = 0$, or $O(\log^5 n / \log p + d \log n)$ when $2 \leq p < n$, on $O(\text{MM}(n) \log n + (n + m)\text{M}(n))$ processors.*

The algorithm is of the Las Vegas type, returns the correct answer with probability at least $1/4$ and otherwise reports “failure”.

Proof. First compute the Frobenius form $S \in \mathbf{K}^{n \times n}$ of T , and an invertible $U \in \mathbf{F}^{n \times n}$ such that $U^{-1}TU = S$. Assume that $S = \text{diag}(C_{f_1}, \dots, C_{f_k})$, where $f_1, \dots, f_k \in \mathbf{K}[x]$ are the invariant factors of T , and C_{f_1}, \dots, C_{f_k} are their respective companion matrices. Then evaluate φ at $x \bmod f_1$ in the ring $\mathbf{K}[x]/(f_1)$. If this evaluation fails, then we report that the circuit is undefined at T . Otherwise, assume the result of this evaluation yields $h \bmod f_1$ for some $h \in \mathbf{K}[x]$. Next compute $h_i \equiv h \bmod f_i$ for $1 \leq i \leq k$. The final result is then $U^{-1} \text{diag}(h_1(C_{f_1}), \dots, h_k(C_{f_k}))U \in \mathbf{K}^{n \times n}$. Correctness is clear from the ring-isomorphism between $\mathbf{K}[T]$ and $\mathbf{K}[x] \bmod f_1$.

The cost of the algorithm is measured as follows. Computing the Frobenius form requires $O(\text{MM}(n) \log n + n\text{M}(n))$ operations in \mathbf{K} using Theorem 4.1. In parallel, we can use Theorem 5.4, which requires $O(\log^4 n)$ time on $O(\text{MM}(n))$ processors when $p \geq n$ or $p = 0$, or time $O(\log^5 n / \log p)$ on $O(\text{MM}(n))$ processors when $2 \leq p < n$. We can evaluate φ at x in $\mathbf{K}[x]/(f_1)$ with $O(m \cdot \text{M}(n))$ operations in \mathbf{K} sequentially, or time $O(d \log n)$ on $O(m \cdot \text{M}(n))$ processors in parallel. The time required by the the remaining steps is dominated by the time required by these first two. \square

As an application of this straight-line program evaluation technique we consider computing high powers of matrices, an operation performed in some cryptographic systems (see, for example, Chuang & Dunham 1990). Using linear recurrences and companion matrices to compute powers of matrices is certainly not new, dating at least to Ranum (1911).

COROLLARY 6.4. *Given $s \geq 0$ and $T \in \mathbb{K}^{n \times n}$, we can compute $T^s \in \mathbb{K}^{n \times n}$ with a Las Vegas type probabilistic algorithm requiring $O(\text{MM}(n) \log n + (n + \log s)\text{M}(n))$ or $\tilde{O}(\text{MM}(n) + n \log s)$ operations in \mathbb{K} , where \mathbb{K} is any field with at least n^2 elements.*

Proof. The repeated squaring method of computing high powers of elements in any group yields a straight-line program of depth and length $O(\log s)$. Apply Theorem 6.3 to obtain the result. \square

We now summarize the above results for fields with fewer than n^2 elements. These follow immediately from the use of Theorems 4.2 (sequential) and 5.5 (parallel) instead of Theorems 4.1 and 5.4 respectively.

THEOREM 6.5. *Let \mathbb{K} be any field with $q < n^2$ elements and characteristic p , $T \in \mathbb{K}^{n \times n}$, and \mathbb{F} the smallest extension field of \mathbb{K} containing n^2 elements. Las Vegas type probabilistic algorithms exist as follows:*

- (i) *given $g \in \mathbb{K}[x]$ of degree r we can evaluate $g(T) \in \mathbb{K}^{n \times n}$ with $O((\text{MM}(n) \log n + n\text{M}(n)) \cdot \text{M}(\log_q n) \log \log_q n + \text{M}(r))$ or $\tilde{O}(\text{MM}(n) + r)$ operations in \mathbb{K} sequentially. In parallel we can do this in time $O(\log^4 n \log \log_q n + \log r)$ when $p = 0$ or $p \geq n$, or in time $O(\log^5 n \log \log_q n / \log p + \log r)$ when $2 \leq p < n$, on $O(\text{MM}(n)\text{M}(\log_q n) + \text{M}(r))$ processors.*
- (ii) *given a straight-line program φ with length m and depth d , we can evaluate φ at T with $O((\text{MM}(n) \log n + n\text{M}(n)) \cdot \text{M}(\log_q n) \log \log_q n + m\text{M}(n))$ operations in \mathbb{K} . In parallel we can do this in time $O(\log^4 n \log \log_q n + d \log n)$ when $p = 0$ or $p \geq n$, or in time $O(\log^5 n \log \log_q n / \log p + d \log n)$ when $2 \leq p < n$, on $O(\text{MM}(n)\text{M}(\log_q n) + m \cdot \text{M}(n))$ processors.*

A Lower Bound on Evaluating Polynomials at Matrices. In this subsection we show that evaluating a non-linear polynomial at a matrix is at least as hard as matrix multiplication. Specifically, if $g \in \mathbb{K}[x]$ has degree $r \geq 2$ and we can evaluate $g(T)$ at any $T \in \mathbb{K}^{n \times n}$ with at most $t(n)$ operations in \mathbb{K} , then we can multiply two $n \times n$ matrices over \mathbb{K} with $O(t(n))$ operations in \mathbb{K} . We make a technical assumption about the cost function $t: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$:

$$\forall a \in \mathbb{R}_{>0} \quad \forall b \in \mathbb{R}_{\geq 1} \quad b^2 t(a) \leq t(ab) \leq b^{3.5} t(a).$$

The problem of evaluating a polynomial at a matrix has a trivial lower bound of $\Omega(n^2)$ operations in \mathbb{K} and, as discussed above, a deterministic upper bound of $O(n^{3.5})$ operations in \mathbb{K} , making this assumption reasonable.

We will assume here that $\#\mathbb{K} \geq r - 1$. A similar argument is presented in Giesbrecht (1993) for smaller fields. Let $A, B \in \mathbb{K}^{n \times n}$. We show how to use an algorithm for evaluating g at a matrix to multiply A and B . Consider the matrix

$$T = \begin{pmatrix} I_n & A & 0_n \\ 0_n & cI_n & B \\ 0_n & 0_n & I_n \end{pmatrix} \in \mathbb{K}^{3n \times 3n},$$

where $c \in \mathbb{K}$ and $0_n, I_n \in \mathbb{K}^{n \times n}$ are the zero and identity matrices respectively.

LEMMA 6.6. *For $j \geq 1$,*

$$T^j = \begin{pmatrix} I_n & \left(\sum_{0 \leq i \leq j-1} c^i \right) \cdot A & \left(\sum_{0 \leq i \leq j-2} (j-i-1)c^i \right) \cdot AB \\ 0_n & c^j I_n & \left(\sum_{0 \leq i \leq j-1} c^i \right) \cdot B \\ 0_n & 0_n & I_n \end{pmatrix} \in \mathbb{K}^{3n \times 3n}.$$

Proof. We proceed by induction on j . Clearly the lemma is true for $j = 1$. Assume the lemma holds for $j - 1$ (we assume $j \geq 2$). Then

$$\begin{aligned} T^j &= T \cdot T^{j-1} = \begin{pmatrix} I_n & A & 0_n \\ 0_n & cI_n & B \\ 0_n & 0_n & I_n \end{pmatrix} \begin{pmatrix} I_n & \left(\sum_{0 \leq i \leq j-2} c^i \right) \cdot A & \left(\sum_{0 \leq i \leq j-3} (j-i-2) c^i \right) \cdot AB \\ 0_n & c^{j-1} I_n & \left(\sum_{0 \leq i \leq j-2} c^i \right) \cdot B \\ 0_n & 0_n & I_n \end{pmatrix} \\ &= \begin{pmatrix} I_n & \left(\sum_{0 \leq i \leq j-1} c^i \right) \cdot A & \left(\sum_{0 \leq i \leq j-2} (j-i-1) c^i \right) \cdot AB \\ 0_n & c^j I_n & \left(\sum_{0 \leq i \leq j-1} c^i \right) \cdot B \\ 0_n & 0_n & I_n \end{pmatrix}. \end{aligned}$$

□

Let $C \in \mathbb{K}^{n \times n}$ be the top right $n \times n$ block of $g(T)$. If $g = \sum_{0 \leq i \leq r} b_i x^i$, then

$$C = \sum_{0 \leq i \leq r} b_i \sum_{0 \leq j \leq i-2} c^j AB = \left(\sum_{0 \leq i \leq r-2} c^j \left(\sum_{i+2 \leq j \leq r} (j-i-1) b_i \right) \right) \cdot AB = h(c) \cdot AB,$$

where $h = \sum_{0 \leq i \leq r-2} z^i \left(\sum_{i+2 \leq j \leq r} (j-i-1) b_j \right) \in \mathbb{K}[z]$ and z is an indeterminate. Since the coefficient of z^{r-2} in h is $b_r \neq 0$, we know $\deg h = r - 2$. Now, we have assumed that $\#\mathbb{K} > r - 2$, so there exists a $c \in \mathbb{K}$ such that $h(c) \neq 0$ because h has at most $r - 2$ roots. Such a c can be found by evaluating h at $r - 1$ distinct points in \mathbb{K} , and using fast multi-point evaluation (see Aho *et al.* 1974, Section 8.5), can be found with $O(M(r) \log r)$ operations in \mathbb{K} . Now $AB = h(c)^{-1} C$, which can be computed with an additional $O(t(n))$ operations in \mathbb{K} .

THEOREM 6.7. *Let $g \in \mathbb{K}[x]$ have degree $r \geq 2$. Suppose we can evaluate $g(T)$ for any matrix $T \in \mathbb{K}^{n \times n}$ with $t(n)$ operations in \mathbb{K} . If $\#\mathbb{K} \geq r - 1$ we can multiply two $n \times n$ matrices with $O(t(n) + M(r) \log r)$ operations in \mathbb{K} .*

In fact, for any given n , we need only compute a single $c \in \mathbb{K}$ such that $h(c) \neq 0$. After such a c has been found, evaluating g at any matrix in $\mathbb{K}^{n \times n}$ requires only $O(t(n))$ operations in \mathbb{K} .

7. Computing the Rational Jordan Form. The rational Jordan form of a matrix $T \in \mathbb{K}^{n \times n}$ is a generalization of the usual Jordan form. Whereas a Jordan form of T exists only if \mathbb{K} contains all eigenvalues of T , the rational Jordan form is a matrix in $\mathbb{K}^{n \times n}$. The eigenvalues of T , which may generate an algebraic extension of \mathbb{K} of exponential degree over \mathbb{K} , are replaced by the companion matrices of their minimal polynomials in $\mathbb{K}[x]$. For any monic irreducible $g \in \mathbb{K}[x]$ of degree r , and any

$m > 0$, define the *rational Jordan block* $J_g^{(m)} \in \mathbb{K}^{mr \times mr}$ as

$$J_g^{(m)} = \begin{pmatrix} \begin{array}{cc|ccc} C_g & I_r & & & \\ & & C_g & I_r & \\ & & & \ddots & \\ & & & & I_r \\ & & & & & C_g \end{array} & \mathbf{0} \end{pmatrix} \in \mathbb{K}^{mr \times mr}, \quad (7.4)$$

where I_r is the $r \times r$ identity matrix, and $C_g \in \mathbb{K}^{r \times r}$ is the companion matrix of g . It is easily shown that both the minimal and characteristic polynomials of $J_g^{(m)}$ are equal to $g^m \in \mathbb{K}[x]$ (see 7.1 below). Any $T \in \mathbb{K}^{n \times n}$ is similar to its *rational Jordan form*, a matrix $Q = \text{diag}(Q_1, Q_2, \dots, Q_l) \in \mathbb{K}^{n \times n}$, where Q_1, \dots, Q_l are rational Jordan blocks. Notice that when all eigenvalues of T lie in \mathbb{K} , then Q is in fact the usual Jordan form of T . The rational Jordan form is unique up to the order of the (rational) Jordan blocks, as is the usual Jordan form. Rational forms akin to the rational Jordan form were investigated at least as early as Frobenius (1911). Kaltofen *et al.* (1990) exhibit fast probabilistic parallel algorithms for a somewhat different rational Jordan form, as well as a ‘‘symbolic Jordan form’’ of a matrix. Roch & Villard (1994) demonstrate fast deterministic parallel algorithms, in the complexity class NC^3 , for the symbolic Jordan form of a matrix.

Let \mathbb{K} be any field with at least n^2 elements. The problem of finding the rational Jordan form $Q \in \mathbb{K}^{n \times n}$ of $T \in \mathbb{K}^{n \times n}$ is reducible, with $O(\text{MM}(n) \log n + nM(n))$ operations in \mathbb{K} , to factoring a single polynomial in $\mathbb{K}[x]$ of degree at most n , namely the minimal polynomial of T which is found from the Frobenius form of T . The necessary observation is that once we know the Frobenius form of T and the complete factorization of its minimal polynomial f_1 , we can immediately determine the rational Jordan form of T .

LEMMA 7.1. *Let $g \in \mathbb{K}[x]$ be monic and irreducible of degree r , and $m \geq 1$. The rational Jordan block $J_g^{(m)} \in \mathbb{K}^{mr \times mr}$ has minimal polynomial g^m .*

Proof. If $C_g \in \mathbb{K}^{r \times r}$ is the companion matrix of g , then $\mathbb{K}[C_g]$ is a field isomorphic to $\mathbb{L} = \mathbb{K}[z]/(g)$ for an indeterminate z (see Lidl & Niederreiter 1983, Section 2.5). This isomorphism is realized by mapping $C_g \mapsto z \bmod g$, and can be extended to an isomorphism of matrices in $\mathbb{K}[C_g]^{m \times m}$ with $\mathbb{L}^{m \times m}$. Observe that $J_g^{(m)} \in \mathbb{K}[C_g]^{m \times m}$, and that its image $\mathcal{J}_g^{(m)} \in \mathbb{L}^{m \times m}$ is a Jordan block, with its eigenvalue $\gamma = z \bmod g$ on the diagonal and ones on the superdiagonal. The minimal polynomial of $\mathcal{J}_g^{(m)}$ is $(x - \gamma)^m \in \mathbb{L}[x]$, so the minimal polynomial of $J_g^{(m)}$ in $\mathbb{K}[x]$ is the multiple of $(x - \gamma)^m$ in $\mathbb{K}[x] \setminus \{0\}$ of lowest degree. This is precisely g^m . \square

LEMMA 7.2. *Let $f \in \mathbb{K}[x]$ be monic of degree n and $C_f \in \mathbb{K}^{n \times n}$ the companion matrix of f . If $f = g_1^{m_1} \dots g_s^{m_s}$ for distinct monic irreducible $g_1, \dots, g_s \in \mathbb{K}[x]$, and $m_1, \dots, m_s > 0$, then C_f has rational Jordan form $B = \text{diag}(J_{g_1}^{(m_1)}, \dots, J_{g_s}^{(m_s)})$.*

Proof. It is well known that the minimal polynomial of a block diagonal matrix is the least common multiple of the minimal polynomials of its blocks. Thus, $\min(B; \mathbb{K}^{n \times 1}) =$

$\text{lcm}(g_1^{e_1}, \dots, g_s^{e_s}) = g_1^{e_1} \cdots g_s^{e_s} = f$, since the minimal polynomial of $J_{g_i}^{(e_i)}$ is $g_i^{m_i}$ for $1 \leq i \leq s$ by Lemma 7.1. Since $\deg f = n$ and f is the minimal polynomial of B , B is similar to the companion matrix C_f of f . \square

THEOREM 7.3. *Suppose $T \in \mathbb{K}^{n \times n}$ has invariant factors $f_1, \dots, f_k \in \mathbb{K}[x]$, where f_i has companion matrix $C_{f_i} \in \mathbb{K}^{d_i \times d_i}$, with $\deg f_i = d_i$, for $1 \leq i \leq k$. If $B_i \in \mathbb{K}^{d_i \times d_i}$ is the rational Jordan form of C_{f_i} for $1 \leq i \leq k$, then $Q = \text{diag}(B_1, \dots, B_k) \in \mathbb{K}^{n \times n}$ is the rational Jordan form of T .*

Proof. By Lemma 7.2 there exists a non-singular $P_i \in \mathbb{K}^{d_i \times d_i}$ such that $B_i = P_i^{-1}C_{f_i}P_i$ for $1 \leq i \leq k$. The matrix $P = \text{diag}(P_1, \dots, P_k) \in \mathbb{K}^{n \times n}$ satisfies $Q = P^{-1}TP$, and Q is in rational Jordan form. \square

THEOREM 7.4. *Let \mathbb{K} be any field with at least n^2 elements. Given $T \in \mathbb{K}^{n \times n}$, there is a Las Vegas reduction from the problem of finding the rational Jordan form $Q \in \mathbb{K}^{n \times n}$ of T and an invertible $P \in \mathbb{K}^{n \times n}$ such that $Q = P^{-1}TP$ to the problem of factoring a single polynomial in $\mathbb{K}[x]$ of degree at most n (namely the minimal polynomial $f_1 \in \mathbb{K}[x]$ of T). This reduction requires $O(\text{MM}(n) \log n + n\text{M}(n))$ or $O^*(\text{MM}(n))$ operations in \mathbb{K} .*

Proof. We find the Frobenius form $S \in \mathbb{K}^{n \times n}$ of T and an invertible $U_0 \in \mathbb{K}^{n \times n}$ such that $U_0^{-1}TU_0 = S$ using the algorithm **FrobeniusForm**. We factor the first invariant factor $f_1 \in \mathbb{K}[x]$ (determined by the Frobenius form and identical to the minimal polynomial of T) and construct the rational Jordan form $Q \in \mathbb{K}^{n \times n}$ of T using Theorem 7.3. An invertible $U_1 \in \mathbb{K}^{n \times n}$ such that $U_1^{-1}QU_1 = S$ is constructed using **FrobeniusForm** (we know Q is similar to S by its construction). Then $Q = U_1U_0^{-1}TU_0U_1^{-1}$ and $P = U_0U_1^{-1}$ satisfies $Q = P^{-1}TP$. \square

To find the rational Jordan form of $T \in \mathbb{K}^{n \times n}$ when \mathbb{K} has fewer than n^2 elements we embed \mathbb{K} into a slightly larger field \mathbb{F} to find the Frobenius form of T , using Theorem 4.2. As in Theorem 4.2, $\mathbb{F} = \mathbb{K}[x]/(\psi)$, where $\psi \in \mathbb{K}[x]$ is monic and irreducible of degree $\lceil 2 \log_q n \rceil$.

THEOREM 7.5. *Let \mathbb{K} be a field with $q = \#\mathbb{K} < n^2$. Given $T \in \mathbb{K}^{n \times n}$, there is a Las Vegas reduction from the problem of finding the rational Jordan form $Q \in \mathbb{K}^{n \times n}$ of T , and an invertible $P \in \mathbb{F}^{n \times n}$ such that $Q = P^{-1}TP$, to factoring a single polynomial in $\mathbb{K}[x]$ of degree at most n (namely the minimal polynomial $f_1 \in \mathbb{K}[x]$ of T). This reduction requires $O((\text{MM}(n) \log n + n\text{M}(n)) \cdot \text{M}(\log_q n) \log \log_q n)$ or $O^*(\text{MM}(n))$ operations in \mathbb{K} .*

Using the Las Vegas algorithm of Berlekamp (1970) for factoring polynomials over finite fields we obtain an algorithm for computing the rational Jordan form of a matrix, given only that matrix.

COROLLARY 7.6. *Let \mathbb{K} be a finite field of size q and $T \in \mathbb{K}^{n \times n}$. The rational Jordan form of T can be computed by a Las Vegas algorithm with $O(\text{MM}(n) \log n + n\text{M}(n) + n \log q)$ or $O^*(\text{MM}(n) + n \log q)$ operations in \mathbb{K} if $\#\mathbb{K} \geq n^2$, and $O((\text{MM}(n) \cdot \log n + n\text{M}(n)) \cdot \text{M}(\log_q n) \log \log_q n + n \log q)$ or $O^*(\text{MM}(n) + n \log q)$ operations in \mathbb{K} if $\#\mathbb{K} < n^2$.*

Acknowledgement. The author wishes to thank Joachim von zur Gathen and Erich Kaltofen for their enthusiasm and help while writing this paper.

REFERENCES

- A. V. AHO, J. E. HOPCROFT, AND J. D. ULLMAN, *The Design and Analysis of Computer Algorithms*, Addison-Wesley, Reading MA, 1974.
- W. BAUR AND V. STRASSEN, *The complexity of partial derivatives*, *Theor. Comp. Science*, 22 (1982), pp. 317–330.
- E. R. BERLEKAMP, *Factoring polynomials over large finite fields*, *Math. Comp.*, 24 (1970), pp. 713–735.
- D. BINI AND V. PAN, *Improved parallel polynomial division*, in *Proc. 33rd IEEE Symp. Foundat. Comput. Sci.*, Pittsburgh, USA, 1992, pp. 131–136.
- A. BORODIN AND I. MUNRO, *The Computational Complexity of Algebraic and numeric Problems*, American Elsevier, New York NY, 1975.
- R. P. BRENT AND H. T. KUNG, *Fast algorithms for manipulating formal power series*, *J. Assoc. Comput. Mach.*, 25 (1978), pp. 581–595.
- J. BUNCH AND J. HOPCROFT, *Triangular factorization and inversion by fast matrix multiplication*, *Math. Comp.*, 28 (1974), pp. 231–236.
- D. CANTOR AND E. KALTOFEN, *Fast multiplication of polynomials over arbitrary algebras*, *Acta Informatica*, 28 (1991), pp. 693–701.
- C. CHUANG AND J. DUNHAM, *Matrix extensions of the RSA algorithm*, in *Proc. Crypto 1990*, 1990, pp. 137–153.
- D. COPPERSMITH AND S. WINOGRAD, *Matrix multiplication via arithmetic progressions*, *J. Symb. Comp.*, 9 (1990), pp. 251–280.
- A. DANILEVSKY, *The numerical solution of the secular equation*, *Matem. sbornik*, 44 (1937), pp. 169–171. In Russian.
- W. EBERLY, *Efficient parallel independent subsets and matrix factorizations*, in *Proc. 3rd IEEE Symp. Parallel and Distributed Processing*, Dallas, USA, 1991, pp. 204–211.
- D. FADDEEV AND V. FADDEEVA, *Computational methods of linear Algebra*, W. H. Freeman and Company, San Francisco, 1963.
- G. FROBENIUS, *Über den Rang einer Matrix*, *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften zu Berlin*, (1911), pp. 54–65.
- F. R. GANTMACHER, *The Theory of Matrices, Vol. I*, Chelsea Publishing Co., New York NY, 1990.
- M. GIESBRECHT, *Nearly Optimal Algorithms for Canonical Matrix Forms*, PhD thesis, University of Toronto, 1993. 196 pp.
- K. HOFFMAN AND R. KUNZE, *Linear Algebra*, Prentice-Hall, Englewood Cliffs, N.J., 1971.
- E. KALTOFEN, *Efficient solution of sparse linear systems*, 1992. Lecture notes, Dept. of Computer Science, Rensselaer Polytechnic Institute, Troy, NY, USA.
- E. KALTOFEN, M. S. KRISHNAMOORTHY, AND B. D. SAUNDERS, *Fast parallel computation of Hermite and Smith forms of polynomial matrices*, *SIAM J. Algebraic and Discrete Methods*, 8 (1987), pp. 683–690.
- , *Parallel algorithms for matrix normal forms*, *Linear Algebra and its Applications*, 136 (1990), pp. 189–208.
- E. KALTOFEN AND V. PAN, *Processor-efficient parallel solution of linear systems over an abstract field*, in *Proc. 3rd Annual ACM Symp Parallel Algorithms and Architectures*, ACM Press, 1991, pp. 180–191.
- , *Processor-efficient parallel solution of linear systems II: The general case*, in *Proc. 33rd IEEE Symp. Foundations of Comp. Sci.*, Pittsburgh, USA, 1992, pp. 714–723.
- R. KANNAN, *Polynomial-time algorithms for solving systems of linear equations over polynomials*, *Theoretical Computer Science*, 39 (1985), pp. 69–88.
- R. KANNAN AND A. BACHEM, *Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix*, *SIAM J. Comp.*, 8 (1979), pp. 499–507.
- R. M. KARP AND V. RAMACHANDRAN, *Parallel algorithms for shared-memory machines*, in *Handbook of Theoretical Computer Science*, J. van Leeuwen, ed., vol. A: Algorithms and Complexity, MIT Press/Elsevier, Cambridge, Mass., USA, 1990, pp. 869–932.
- W. KELLER-GEHRIG, *Fast algorithms for the characteristic polynomial*, *Theor. Computer Science*, 36 (1985), pp. 309–317.
- R. LIDL AND H. NIEDERREITER, *Finite Fields*, vol. 20 of *Encyclopedia of Mathematics and its Applications*, Addison-Wesley, Reading MA, 1983.
- H. LÜNEBURG, *On Rational Normal Form of Endomorphisms: a Primer to Constructive Algebra*, Wissenschaftsverlag, Mannheim, 1987.
- H. J. NUSSBAUMER, *Fast polynomial transform algorithms for digital convolutions*, *IEEE Trans. ASSP*, 28 (1980), pp. 395–398.
- P. OZELLO, *Calcul Exact Des Formes De Jordan et de Frobenius d'une Matrice*, PhD thesis, Uni-

- versité Scientifique Technologique et Medicale de Grenoble, 1987.
- V. PAN AND J. REIF, *Efficient parallel solutions of linear systems*, in Proc. 17th Annual Symp. Theory of Computing, Providence, Rhode Island, USA, 1985, pp. 143–152.
- M. S. PATERSON AND L. STOCKMEYER, *On the number of nonscalar multiplications necessary to evaluate polynomials*, SIAM J. Comput., 2 (1973), pp. 60–66.
- A. RANUM, *The general term of a recurring series*, Bull. American Math. Society, 17 (1911), pp. 457–461.
- J. ROCH AND G. VILLARD, *Calcul formel et parallélisme: étude de la forme normal de Jordan*, Tech. Report 7, Institut IMAG, Laboratoire LMC, 1994.
- A. SCHÖNHAGE, *Schnelle Multiplikation von Polynomen über Körpern der Charakteristik 2*, Acta Informatica, 7 (1977), pp. 395–398.
- A. SCHÖNHAGE AND V. STRASSEN, *Schnelle Multiplikation großer Zahlen*, Computing, 7 (1971), pp. 281–292.
- J. T. SCHWARTZ, *Fast probabilistic algorithms for verification of polynomial identities*, J. Assoc. Computing Machinery, 27 (1980), pp. 701–717.
- V. SHOUP, *Fast construction of irreducible polynomials over finite fields*, in Proc. 4th Annual ACM-SIAM Symp. on Discrete Algorithms, Austin, USA, 1993, pp. 484–492.
- B. L. VAN DER WAERDEN, *Algebra*, vol. 1, Frederick Ungar Publishing Co., New York, 7th ed., 1970.
- G. VILLARD, *Calcul formel et parallélisme: étude de la forme normal de Smith*, Tech. Report 8, Institut IMAG, Laboratoire LMC, 1994.
- J. VON ZUR GATHEN, *Parallel arithmetic computations: a survey*, in Proc. 12th Int. Symp. Math. Foundations of Computer Science, vol. 233 of Springer Lecture Notes in Computer Science, Bratislava, 1986, pp. 93–112.
- , *Parallel linear algebra*, in Synthesis of Parallel Algorithms, J. Reif, ed., Morgan Kaufmann, San Mateo, USA, 1993, pp. 573–617.
- D. WIEDEMANN, *Solving sparse linear equations over finite fields*, IEEE Transactions on Information Theory, IT-32 (1986), pp. 54–62.
- Y. ZALCSTEIN AND M. GARZON, *An NC^2 algorithm for testing similarity of matrices*, Information Processing Letters, (1987), pp. 253–254.