



HAL
open science

On the complexity of computing determinants

Erich Kaltofen, Gilles Villard

► **To cite this version:**

Erich Kaltofen, Gilles Villard. On the complexity of computing determinants. [Research Report] LIP RR-2003-36, Laboratoire de l'informatique du parallélisme. 2003, 2+35p. hal-02102099

HAL Id: hal-02102099

<https://hal-lara.archives-ouvertes.fr/hal-02102099>

Submitted on 17 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

***Laboratoire de l'Informatique du
Parallélisme***



École Normale Supérieure de Lyon
Unité Mixte de Recherche CNRS-INRIA-ENS LYON
n° 5668




***On the complexity of computing
determinants***

Erich Kaltofen
North Carolina State University
and Gilles Villard

July 2003

Research Report N° 2003-36



**École Normale Supérieure de
Lyon**

46 Allée d'Italie, 69364 Lyon Cedex 07, France
Téléphone : +33(0)4.72.72.80.37
Télécopieur : +33(0)4.72.72.80.80
Adresse électronique : lip@ens-lyon.fr



On the complexity of computing determinants

Erich Kaltofen and Gilles Villard

July 2003

Résumé

En combinant la technique des pas de géant/pas de bébé de Kaltofen (1992) appliquée à l'algorithme de calcul du déterminant de Wiedemann (1986), avec l'utilisation de projections par blocs de vecteurs suivant Coppersmith (1994) et l'analyse correspondante de Villard (1997), nous obtenons de nouveaux algorithmes asymptotiquement rapides pour la résolution de différents problèmes sur des matrices denses.

La première catégorie de problèmes concerne les matrices A denses $n \times n$ à coefficients entiers. Nous exprimons les coûts de calcul pour des entiers exacts dans une base fixée, i.e. en termes d'opérations binaires, en notant $\|A\|$ le plus grand coefficient de la matrice en valeur absolue. Nous calculons le déterminant, le polynôme caractéristique, la forme normale de Frobenius et la forme normale de Smith de A en $(n^{3.2} \log \|A\|)^{1+o(1)}$ et $(n^{2.697263} \log \|A\|)^{1+o(1)}$ opérations binaires, l'ajustement “ $+o(1)$ ” de l'exposant tient compte de facteurs additionnels $C_1(\log n)^{C_2}(\log \|A\|)^{C_3}$ où C_1 , C_2 et C_3 sont des constantes positives réelles. La première complexité donnée ci-dessus, la plus lente asymptotiquement, est atteinte sans utiliser les algorithmes sous-cubiques pour le produit de matrices. Nos algorithmes sont probabilistes, cependant, le déterminant peut être certifié et conduit à une résolution Las Vegas.

La seconde catégorie de problèmes correspond au cas où la matrice A a ses coefficients dans un anneau commutatif abstrait, c'est-à-dire que l'on ne s'autorise pas les divisions. Nous présentons une approche déterministe pour calculer le déterminant, le polynôme caractéristique et la matrice adjointe de A en $n^{3.2+o(1)}$ additions, soustractions et multiplications dans l'anneau, sans utiliser de produits de matrices sous-cubiques. En faisant appel à la multiplication de matrices de Coppersmith et Winograd nous calculons le déterminant et l'adjointe en $O(n^{2.697263})$ opérations dans l'anneau, le polynôme caractéristique est obtenu en $O(n^{2.806515})$ opérations. Ces résultats reposent sur une nouvelle preuve de l'analyse de Villard (1997) pour l'algorithme de Wiedemann/Lanczos par blocs et sur une généralisation aux matrices polynomiales de l'algorithme d'Euclide à la Knuth/Schönhage/Moenck.

Mots-clés: Algèbre linéaire, déterminant, polynôme caractéristique, formes normales de matrices, algorithmes rapides, algorithmes probabilistes, algorithmes sans division.

Abstract

By combining Kaltofen's 1992 baby steps/giant steps technique for Wiedemann's 1986 determinant algorithm with Coppersmith's 1994 projections by a block of vectors in the Wiedemann approach and Villard's 1997 analysis of the block technique, we obtain new algorithms for dense matrix problems of asymptotically fast running time. The first category of problems is for a dense $n \times n$ matrix A with integer entries. We express the cost in terms of bit operations on the exact integers and denote by $\|A\|$ the largest entry in absolute value. Our algorithms compute the determinant, characteristic polynomial, Frobenius normal form and Smith normal form of A in $(n^{3.2} \log \|A\|)^{1+o(1)}$ and $(n^{2.697263} \log \|A\|)^{1+o(1)}$ bit operations, where the exponent adjustment by “ $+o(1)$ ” captures additional factors $C_1(\log n)^{C_2}(\log \|A\|)^{C_3}$ for positive real constants C_1, C_2, C_3 and where the first, asymptotically slower bit complexity does not require any of the sub-cubic matrix multiplication algorithms. Our algorithms are randomized, and we can certify the determinant of A in a Las Vegas fashion. The second category of problems deals with the setting where the matrix A has elements from an abstract commutative ring, that is, when no divisions in the domain of entries are possible. We present algorithms that deterministically compute the determinant, characteristic polynomial and adjoint of A with $n^{3.2+o(1)}$ ring additions, subtractions and multiplications, that without utilizing sub-cubic matrix multiplication algorithms. With the asymptotically fast matrix multiplication algorithms by Coppersmith and Winograd our method computes the determinant and adjoint in $O(n^{2.697263})$ ring operations and the characteristic polynomial in $O(n^{2.806515})$ ring operations. We achieve our results in part through new proofs for Villard's 1997 analysis of the block Wiedemann/Lanczos algorithm and a generalization of the Knuth/Schönhage/Moenck Euclidean remainder sequence algorithm to matrix polynomials.

1 Introduction

The computational complexity of many problems in linear algebra has been tied to the computational complexity of matrix multiplication. If the result is to be exact, for example the exact rational solution of a linear system, the lengths of the integers involved in the computation and the answer affect the running time of the used algorithms. A classical methodology is to compute the results via Chinese remaindering. Then the standard analysis yields a number of fixed radix, i.e. bit operations for a given problem that is essentially (within polylogarithmic factors) bounded by the number of field operations for the problem times the maximal scalar length in the output. The algorithms at times use randomization, because not all modular images may be usable. For the determinant

*This material is based on work supported in part by the National Science Foundation (USA) under Grants Nos. DMS-9977392, CCR-9988177 and CCR-0113121 (Kaltofen) and by CNRS (France) Actions Incitatives No 5929 et STIC LINBox 2001 (Villard).

Extended abstract appears in the *Computer Mathematics Proc. Fifth Asian Symposium (ASCM 2001)* edited by Kiyoshi Shirayanagi and Kazuhiro Yokoyama, Lecture Notes Series on Computing, vol. 9, World Scientific, Singapore, 2001, pages 13–27.

of an $n \times n$ integer matrix A one thus gets a running time of $(n^4 \log \|A\|)^{1+o(1)}$ bit operations [von zur Gathen and Gerhard 1999: Chapter 5.5], because the determinant can have at most $(n \log \|A\|)^{1+o(1)}$ digits; by $\|A\|$ we denote the largest entry in absolute value. Here and throughout this paper the exponent adjustment by “ $+o(1)$ ” captures additional factors $C_1(\log n)^{C_2}(\log \|A\|)^{C_3}$ for positive real constants C_1, C_2, C_3 (“soft- O ”). Via an algorithm that can multiply two $n \times n$ matrices in $O(n^\omega)$ scalar operations the time is reduced to $(n^{\omega+1} \times \log \|A\|)^{1+o(1)}$. By [Coppersmith and Winograd 1990] we can set $\omega = 2.375477$.

First, it was recognized that for the problem of computing the exact rational solution of a linear system the process of Hensel lifting can accelerate the bit complexity beyond the Chinese remainder approach [Dixon 1982], namely to cubic in n without using fast matrix multiplication algorithms. For the determinant of an $n \times n$ integer matrix A , an algorithm with $(n^{3.5} \log \|A\|^{1.5})^{1+o(1)}$ bit operations is given in [Eberly et al. 2000].* The algorithm by Eberly et al. computes the Smith normal form via the binary search technique of [Villard 2000].

Our algorithms combine three ideas.

- I) The first is an algorithm in [Wiedemann 1986] for computing the determinant of a sparse matrix over a finite field. Wiedemann finds the minimum polynomial for the matrix as a linear recurrence on a corresponding Krylov sequence. By preconditioning the input matrix, that minimum polynomial is the characteristic polynomial and the determinants of the original and preconditioned matrix have a direct relation.
- II) The second is from [Kaltofen 1992] where Wiedemann’s approach is applied to dense matrices whose entries are polynomials over a field. Kaltofen achieves speedup by employing Shank’s baby steps/giant steps technique for the computation of the linearly recurrent scalars (cf. [Paterson and Stockmeyer 1973]). For integer matrices the resulting randomized algorithm is of the Las Vegas kind—always correct, probably fast—and has worst case bit complexity $(n^{3.5} \log \|A\|)^{1+o(1)}$ and again can be speeded with sub-cubic time matrix multiplication [Kaltofen and Villard 2001]. A detailed description of this algorithm, with an early termination strategy in case the determinant is small (cf. [Emiris 1998; Brönnimann et al. 1999]), is presented in [Kaltofen 2002].
- III) By considering a bilinear map using two blocks of vectors rather than a single pair of vectors, Wiedemann’s algorithm can be accelerated [Coppersmith 1994; Kaltofen 1995; Villard 1997a,b]. Blocking can be applied to our algorithms for dense matrices and further reduces the bit complexity.

The above ingredients yield a randomized algorithm of the Las Vegas kind for computing the determinant of an $n \times n$ integral matrix A in $(n^{3+1/3} \times \log \|A\|)^{1+o(1)}$ expected bit operations, that with a standard cubic matrix mul-

*In [Eberly et al. 2000] the exponent for $\log \|A\|$ is 2.5, but the improvement to 1.5 based on fast Chinese remaindering [Aho et al. 1974] is immediate.

tiplication algorithm. If we employ fast FFT-based Padé approximation algorithms for matrix polynomials, for example the so-called half-GCD algorithm [von zur Gathen and Gerhard 1999] and fast matrix multiplication algorithms, we can further lower the expected number of bit operations. Under the assumption that two $n \times n$ matrices can be multiplied in $O(n^\omega)$ operations in the field of entries, and an $n \times n$ matrix by an $n \times n^\zeta$ matrix in $n^{2+o(1)}$ operations, we obtain an expected bit complexity for the determinant of

$$(n^\eta \log \|A\|)^{1+o(1)} \text{ with } \eta = \omega + \frac{1 - \zeta}{\omega^2 - (2 + \zeta)\omega + 2}. \quad (1)$$

The best known values $\omega = 2.375477$ [Coppersmith and Winograd 1990] and $\zeta = 0.2946289$ [Coppersmith 1997] yield $\eta = 2.697263$. For $\omega = 3$ and $\zeta = 0$ we have $\eta = 3 + 1/5$ as given in the abstract above (cf. [Kaltofen and Villard 2002; Pan 2002]).

Our techniques can be further combined with the ideas in [Giesbrecht 2001] to produce a randomized algorithm for computing the integer Smith normal form of an integer matrix. The method becomes Monte Carlo—always fast and probably correct—and has the same bit complexity (1). In addition, we can compute the characteristic polynomial of an integer matrix by Hensel lifting [Storjohann 2000b]. Again the method is Monte Carlo and has bit complexity (1). Both results utilize the fast determinant algorithm for matrix polynomials [Storjohann 2002, 2003].

The algorithm in [Kaltofen 1992] (see case II above) was originally put to a different use, namely that of computing the characteristic polynomial and adjoint of a matrix without divisions, counting additions, subtractions, and multiplications in the commutative ring of entries. Serendipitously, blocking (see case III above) can be applied to our original 1992 division-free algorithm, and we obtain a deterministic algorithm that computes the determinant of a matrix over a commutative ring in $n^{\eta+o(1)}$ ring additions, subtractions and divisions, where η is given by (1). The exponent $\eta = 2.697263$ seems to be the best that is known today for the division-free determinant problem. By the technique in [Baur and Strassen 1983] we obtain the adjoint of a matrix in the same division-free complexity. For the characteristic polynomial we can obtain a deterministic division-free complexity of $O(n^{2.806515})$ ring operations. The higher exponent here is a result of the lack of algorithms like those in [Storjohann 2002; Jeannerod and Villard 2002; Storjohann 2003] for the division-free model.

In [Kaltofen and Villard 2002] we have identified other algorithms for computing the determinant of an integer matrix. Those algorithms often perform at cubic bit complexity on what we call are propitious inputs, but they have a worst case bit complexity that is higher than our methods. One such method is Clarkson's algorithm [Clarkson 1992; Brönnimann and Yvinec 2000], where the number of mantissa bits in the intermediate floating point scalars that are necessary for obtaining a correct sign depends on the orthogonal defect of the matrix. If the matrix has a large first invariant factor, Chinese remaindering can be employed in connection with computing the solution of a random linear

system via Hensel lifting [Abbott et al. 1999] (cf. [Pan 1988]).

Notation: By $S^{m \times n}$ we denote the set of $m \times n$ matrices with entries in the set S . The set \mathbb{Z} are the integers. For $A \in \mathbb{Z}^{n \times n}$ we denote by $\|A\|$ the matrix height [Kaltofen and May 2003: Lemma 2]:

$$\|A\| = \|A\|_{\infty,1} = \max_{x \neq 0} \|Ax\|_{\infty} / \|x\|_1 = \max_{1 \leq i, j \leq n} |a_{i,j}|.$$

Hence the maximal bit length of all entries in A and their signs is, depending on the exact representation, at least $2 + \lfloor \log_2 \max\{1, \|A\|\} \rfloor$. In order to avoid zero factors or undefined logarithms, we shall simply define $\|A\| > 1$ whenever it is necessary.

Organization of the paper. Section 2 introduces Coppersmith's block Wiedemann algorithm and establishes all necessary mathematical properties of the computed matrix generators. In particular, we show the relation of the determinants of the generators with the (polynomial) invariant factors of the characteristic matrix (Theorem 4), which essentially captures the block version of the Cayley-Hamilton property. In addition, we characterize when short sequences are insufficient to determine the minimum generator. Section 3 deals with the computation of the block generator. We give the generalization of the Knuth/Schönhage/Moenck algorithm for polynomial quotient sequences to matrix polynomials and show that in our case by randomization all leading coefficients stay non-singular (Lemma 8). Section 4 presents our new determinant algorithm for integer matrices and gives the running time analysis when cubic matrix multiplication algorithms are employed (Theorem 10). Section 5 presents the division-free determinant algorithm. Section 6 contains the analysis for versions of our algorithms when fast matrix multiplication is introduced. The asymptotically best results are derived there. Section 7 presents the algorithms for the Smith normal form and the characteristic polynomial of an integer matrix. We give concluding thoughts in Section 8.

2 Generating polynomials of matrix sequences

Coppersmith [1994] first has introduced blocking to the Wiedemann method. In our description we also take into account the interpretation in [Villard 1997a,b], where the relevant literature from linear control theory is cited. Our algorithms rely on the notion minimum generating polynomials (generators) of matrix sequences. This notion is introduced below in Section 2.1. We also see how generators are related to block Hankel matrices and recall some basic facts concerning their computation. In Section 2.2 we then study determinants and Smith normal forms of generators and see how they will be used for solving our initial problem.

2.1 Generators and block Hankel matrices

For the “block” vectors $X \in \mathbb{K}^{n \times l}$ and $Y \in \mathbb{K}^{n \times m}$ consider the sequence of $l \times m$ matrices

$$B^{[0]} = X^{\text{Tr}}Y, B^{[1]} = X^{\text{Tr}}AY, B^{[2]} = X^{\text{Tr}}A^2Y, \dots, B^{[i]} = X^{\text{Tr}}A^iY, \dots \quad (2)$$

As in the unblocked Wiedemann method, we seek linear generating polynomials. A vector polynomial $\sum_{i=0}^d c^{[i]} \lambda^i$, where $c^{[i]} \in \mathbb{K}^m$, is said to linearly generate the sequence (2) from the right if

$$\forall j \geq 0: \sum_{i=0}^d B^{[j+i]} c^{[i]} = \sum_{i=0}^d X^{\text{Tr}} A^{i+j} Y c^{[i]} = 0^l. \quad (3)$$

For the minimum polynomial of A , $f^A(\lambda)$, and for the μ -th unit vector in \mathbb{K}^m , $e^{[\mu]}$, $f^A(\lambda)e^{[\mu]} \in \mathbb{K}[\lambda]^m$ is such a generator because it already generates the Krylov sequence $\{A^i Y^{[\mu]}\}_{i \geq 0}$, where $Y^{[\mu]}$ is the μ -th column of Y . We can now consider the set of all such right vector generators. This set forms a $\mathbb{K}[\lambda]$ -submodule of the $\mathbb{K}[\lambda]$ -module $\mathbb{K}[\lambda]^m$ and contains m linearly independent (over the field of rational functions $\mathbb{K}(\lambda)$) elements, namely all $f^A(\lambda)e^{[\mu]}$. Furthermore, the submodule has an (“integral”) basis over $\mathbb{K}[\lambda]$, namely any set of m linearly independent generators such that the degree in λ of the determinant of the matrix formed by those basis vector polynomials as columns is minimal. The matrices corresponding to all integral bases clearly are right equivalent with respect to multiplication from the right by any unimodular matrix in $\mathbb{K}[\lambda]^{m \times m}$, whose determinant is by definition of unimodularity a non-zero element in \mathbb{K} . Thus we can pick a matrix canonical form for this right equivalence, say the Popov form [Popov 1970] (see also [Kailath 1980: §6.7.2]) to get the following definition.

Definition 1 *The unique matrix generating polynomial for (2) in Popov form, denoted by $F_X^{A,Y} \in \mathbb{K}[\lambda]^{m \times m}$, is called the minimum matrix generating polynomial (generator).*

As we will show below, $\deg(\det F_X^{A,Y}) \leq n$. The computation of the minimum matrix generating polynomial from the matrix sequence (2) can be accomplished by several interrelated approaches. One is a sophisticated generalization the Berlekamp/Massey algorithm [Rissanen 1972; Dickinson et al. 1974; Copper-Smith 1994]. Another generalizes the theory of Padé approximation [Forney, Jr. 1975; Van Barel and Bultheel 1992; Beckermann and Labahn 1994; Giorgi et al. 2003]. The interpretation of the Berlekamp/Massey algorithm as a specialization of the extended Euclidean algorithm [Sugiyama et al. 1975; Dornstetter 1987] can be carried over to matrix polynomials [Coppersmith 1994; Thomé 2002] (see also Section 3 below). All approaches solve the classical Levinson-Durbin problem, which for matrix sequences becomes a block Toeplitz linear system [Kaltfen 1995]. The relation to Toeplitz/Hankel matrices turns out to be a useful device for establishing certain properties.

For a degree d and a length e we consider the $l \cdot e$ by $m \cdot (d+1)$ block Hankel matrix

$$\text{Hk}_{e,d+1}(A, X, Y) = \begin{bmatrix} B^{[0]} & B^{[1]} & \dots & B^{[d-1]} & B^{[d]} \\ B^{[1]} & B^{[2]} & & B^{[d]} & B^{[d+1]} \\ \vdots & & \ddots & \vdots & \vdots \\ B^{[e-1]} & \dots & \dots & \dots & B^{[d+e-1]} \end{bmatrix} \quad (4)$$

For any vector generator $\sum_{i=0}^d c^{[i]} \lambda^i \in \mathbb{K}^m[\lambda]$ we must have

$$\text{Hk}_{e,d+1} \cdot \begin{bmatrix} c^{[0]} \\ \vdots \\ c^{[d]} \end{bmatrix} = 0 \text{ for all } e > 0.$$

By considering the rank of (4) we can infer the reverse. If

$$\text{Hk}_{n,d+1} \cdot \begin{bmatrix} c^{[0]} \\ \vdots \\ c^{[d]} \end{bmatrix} = 0 \quad (5)$$

then $\sum_{i=0}^d c^{[i]} \lambda^i$ is a vector generator of (2). The claim follows from the fact that $\text{rank } \text{Hk}_{n,d+1} = \text{rank } \text{Hk}_{n+e',d+1}$ for all $e' > 0$. The latter is justified by observing that any row in the $(n+e')$ th block row of $\text{Hk}_{n+e',d+1}$ is linearly dependent on corresponding previous rows via the minimum polynomial f^A , which has degree $\deg(f^A) \leq n$.

We observe that $\text{rank}(\text{Hk}_{e,d}) \leq n$ for all $d > 0, e > 0$ by considering the factorization

$$\text{Hk}_{e,d} = \begin{bmatrix} X^{Tr} \\ X^{Tr}A \\ X^{Tr}A^2 \\ \vdots \\ X^{Tr}A^{e-1} \end{bmatrix} \cdot [Y \quad AY \quad A^2Y \quad \dots \quad A^{d-1}Y]$$

and noting that either matrix factor has rank at most n .

Therefore, with $d \geq \deg(F_X^{A,Y})$, all solutions to (5) are canonically generated over $\mathbb{K}[\lambda]$ by the columns of $F_X^{A,Y}(\lambda)$, which is in Popov form (see Definition 1). In this case, if the column degrees of the minimum generator are $\delta_1 \leq \dots \leq \delta_m$, the dimension of the right nullspace of $\text{Hk}_{e,d+1}$ in (5) over \mathbb{K} is $(d - \delta_1 + 1) + \dots + (d - \delta_m + 1)$. Hence $\text{rank}(\text{Hk}_{e,d+1}) = \delta_1 + \dots + \delta_m = \deg(\det F_X^{A,Y}) \leq n$ for $d \geq \deg F_X^{A,Y}$ and $e \geq n$. Since the last block column in $\text{Hk}_{e,d+1}$ with $d \geq \deg(F_X^{A,Y})$ is generated by previous block columns, shifting lower degree columns of $F_X^{A,Y}(\lambda)$ as necessary by multiplying with powers of λ , we have

$$\text{rank}(\text{Hk}_{e,d}) = \deg(\det F_X^{A,Y}) \text{ for } d \geq \deg F_X^{A,Y} \text{ and } e \geq n. \quad (6)$$

One may now define the minimum e_{\min} such that the matrix $\text{Hk}_{e_{\min},d}$ for $d = \deg F_X^{A,Y}$ has full rank $\deg(\det F_X^{A,Y})$. Any algorithm for computing the minimum generator requires the first $\deg(F_X^{A,Y}) + e_{\min}$ elements of the sequence (2).

We give an example over \mathbb{Q} [Turner 2002]. Let

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 2 & 0 & 0 & 0 \end{bmatrix}, \quad X = Y = \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$$

Then

$$B^{[0]} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, B^{[1]} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, B^{[2]} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, B^{[3]} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix},$$

$$B^{[4]} = \begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix}, B^{[5]} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, B^{[6]} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, B^{[7]} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

Therefore

$$\text{Hk}_{4,5}(A, X, Y) = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

and from

$$\text{nullspace Hk}_{4,5}(A, X, Y) = \text{span} \left(\begin{bmatrix} -2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right)$$

$$\text{we get } F_X^{A,Y}(\lambda) = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \lambda^4 + \begin{bmatrix} -2 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} \lambda^4 - 2 & 0 \\ 0 & 1 \end{bmatrix}.$$

Now let X as above and let $\bar{Y} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}^{Tr}$. Then

$$\bar{B}^{[0]} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \bar{B}^{[1]} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad \bar{B}^{[2]} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix},$$

$$\overline{B}^{[3]} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad \overline{B}^{[4]} = \begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix}, \quad \overline{B}^{[5]} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

Therefore

$$\text{Hk}_{4,3}(A, X, \overline{Y}) = \left[\begin{array}{cc|cc|cc} 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 1 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right],$$

and from

$$\text{nullspace } \text{Hk}_{4,3}(A, X, \overline{Y}) = \text{span} \left(\begin{bmatrix} 0 \\ -2 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right)$$

we get $F_X^{A, \overline{Y}}(\lambda) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \lambda^2 - \begin{bmatrix} 0 & 1 \\ 2 & 0 \end{bmatrix} = \begin{bmatrix} \lambda^2 & -1 \\ -2 & \lambda^2 \end{bmatrix}$. Note that in both cases the determinant of the minimum generator is $\lambda^4 - 2$, which is $\det(\lambda I - A)$.

The second above example, where $e_{\min} = 4 > \deg(F_X^{A, Y}) = 2$, shows that more than $2 \deg(F_X^{A, Y})$ sequence elements may be necessary to compute the generator, in contrast to the scalar Berlekamp/Massey theory: the last block row of $\text{Hk}_{4,3}(A, X, \overline{Y})$ is required to restrict the right nullspace to the two generating vectors.

However, for random projection block vectors X and Y both $\deg(F_X^{A, Y})$ and e_{\min} are small. Let us define

$$\nu = \max_{d \geq 1, e \geq 1, X \in \mathbb{K}^{n \times l}, Y \in \mathbb{K}^{n \times m}} \{\text{rank } \text{Hk}_{e,d}(A, X, Y)\}. \quad (7)$$

Indeed, the probabilistic analysis [Kaltofen 1995: Section 5], [Villard 1997b: Corollary 6.4] shows the existence of matrices $W \in \mathbb{K}^{n \times l}$ and $Z \in \mathbb{K}^{n \times m}$ such that the corresponding rank $\text{Hk}_{e_0, d_0}(A, W, Z) = \nu$ with $d_0 = \lceil \nu/m \rceil$ and $e_0 = \lceil \nu/l \rceil$. Moreover, ν is equal to the sum of the degrees of the first $\min\{l, m\}$ invariant factors of $\lambda I - A$ (see Theorem 4 below), and hence X, Y can be taken from any field extension of \mathbb{K} . Then due to the existence of W, Z , for symbolic entries in X, Y and therefore, by [DeMillo and Lipton 1978; Zippel 1979; Schwartz 1980], for random entries, the maximal rank is preserved for block dimensions e_0, d_0 . Note that the degree of the minimum matrix generating polynomial is now $\deg(F_X^{A, Y}) = d_0 < n/m + 1$ and the number of sequence elements required to compute the minimum generator is $d_0 + e_0 = \lceil \nu/l \rceil + \lceil \nu/m \rceil < n/l + n/m + 2$. If \mathbb{K} is a small finite field, Wiedmann's analysis has been generalized in [Villard 1997b; Brent et al. 2003].

As with the unblocked Wiedemann projections, unlucky projection block vectors X and Y may cause a drop in the determinantal degree $\deg(F_X^{A,Y})$. They may also increase the length of the sequence required to compute the generator $F_X^{A,Y}$.

2.2 Smith normal forms of matrix generating polynomials

In this section we study how the invariant structure of $F_X^{A,Y}$ partly reveals the structure of A and $\lambda I - A$. Our algorithms in Sections 4 and 5 pick random projections block vectors X, Y or use special projections and compute a generator from the first $d_0 + e_0$ elements of (2). Under the assumption that the rank of $\text{Hk}_{e,d} = \nu$ (see (7)) for sufficiently large d, e , we prove here that $\det(F_X^{A,Y})$ is the product of the first $\min\{l, m\}$ invariant factors of $\lambda I - A$. These are well-studied facts in the theory of realizations of multivariable control theory, for instance see Kailath [1980]. The basis is the matrix power series

$$X^{\text{Tr}}(\lambda I - A)^{-1}Y = X^{\text{Tr}}\left(\sum_{i \geq 0} \frac{A^i}{\lambda^{i+1}}\right)Y = \sum_{i \geq 0} \frac{B^{[i]}}{\lambda^{i+1}}.$$

Lemma 2 *One has the fraction description*

$$X^{\text{Tr}}(\lambda I - A)^{-1}Y = N(\lambda)D(\lambda)^{-1} \quad (8)$$

if and only if there exists $T \in \mathbb{K}[\lambda]^{m \times m}$ such that $D = F_X^{A,Y}T$.

Proof. For the necessary condition, since every polynomial numerator in $X^{\text{Tr}} \times (\lambda I - A)^{-1}Y$ has degree strictly less than the corresponding denominator, then every column of N has degree strictly less than that of the corresponding column of D . Thus it can be checked that the columns of D satisfy (3) and D must be a multiple of $F_X^{A,Y}$. Conversely, let $D = F_X^{A,Y}T$ in $\mathbb{K}[\lambda]^{m \times m}$ be an invertible matrix generator for (2). Using (3) for its m columns it can be seen that we have

$$X^{\text{Tr}}(\lambda I - A)^{-1}YD(\lambda) = N(\lambda) \in \mathbb{K}[\lambda]^{l \times m}$$

where the column degrees of N are lower than those of D . This yields the matrix fraction description (8). \square

Clearly, for $D = F_X^{A,Y}$, the minimum polynomial $f^A(\lambda)$ is a common denominator of the rational entries of the matrices on both sides of (8). If the least common denominator of the left side matrix is actually the characteristic polynomial $\det(\lambda I - A)$, then it follows from degree considerations that $\det F_X^{A,Y} = \det(\lambda I - A)$. Our algorithm uses the matrix preconditioners discussed in Section 4 and random or ad hoc projections (Section 5) to achieve this determinantal equality. We shall make the relationship between $\lambda I - A$ and $F_X^{A,Y}$ more explicit in Theorem 4 whose proof will rely on the structure of the matrix denominator D in (8) and on the following.

For a square matrix M over $\mathbb{K}[\lambda]$ we consider the Smith normal form [Newman 1972], which is an equivalent diagonal matrix over $\mathbb{K}[\lambda]$ with diagonal

elements $s_1(\lambda), \dots, s_\phi(\lambda), 1, \dots, 1, 0, \dots, 0$, where the s_i 's are the nontrivial invariant factors of M , that is, non-constant monic polynomials with the property that s_i is a (trivial or nontrivial) polynomial factor of s_{i-1} for all $2 \leq i \leq \phi$. Because the Smith normal form of the characteristic matrix $\lambda I - A$ corresponds to the Frobenius canonical form of A for similarity, the largest invariant factor of $\lambda I - A$, $s_1(\lambda)$, equals the minimum polynomial $f^A(\lambda)$.

Lemma 3 *Let $M \in \mathbb{K}[\lambda]^{\mu \times \mu}$ be non-singular and let $U \in \mathbb{K}[\lambda]^{\mu \times \mu}$ be unimodular such that*

$$MU = \begin{bmatrix} H & H_{12} \\ 0 & H_{22} \end{bmatrix} \quad (9)$$

where H is a square matrix, then the i -th invariant factor of H divides the i -th invariant factor of M .

Proof. Identity (9) may be rewritten as

$$MU = \begin{bmatrix} I & H_{12} \\ 0 & H_{22} \end{bmatrix} \begin{bmatrix} H & 0 \\ 0 & I \end{bmatrix}.$$

Since the invariant factors of two non-singular matrices divide the invariant factors of their product [Newman 1972: Theorem II.14], the largest invariant factors of $\text{diag}(H, I)$ that are those of H , divide the corresponding invariant factors of MU and thus M . \square

We can now see how the Smith form of $F_X^{A,Y}$ is related to that of $\lambda I - A$. Essentially the result may be obtained for instance following the lines of [Kailath 1980: §6.4.2], we give here a statement and a proof better suited to our purposes.

Theorem 4 *Let $A \in \mathbb{K}^{n \times n}$, $X \in \mathbb{K}^{n \times l}$, $Y \in \mathbb{K}^{n \times m}$ and let s_1, \dots, s_ϕ denote all invariant factors of $\lambda I - A$. The i -th invariant factor of $F_X^{A,Y}$ divides s_i . Furthermore, there exist matrices $W \in \mathbb{K}^{n \times l}$ and $Z \in \mathbb{K}^{n \times m}$ such that for all i , $1 \leq i \leq \min\{l, m, \phi\}$, the i -th invariant factor of $F_W^{A,Z}$ is equal to s_i and the $m - \min\{l, m, \phi\}$ remaining ones are equal to 1. Moreover, for fixed l and m ,*

$$\left. \begin{aligned} \deg_\lambda(\det(F_W^{A,Z}(\lambda))) &= \max_{X,Y} \deg_\lambda(\det(F_X^{A,Y}(\lambda))) \\ &= \deg(s_1) + \dots + \deg(s_{\min\{l,m,\phi\}}) \\ &= \nu, \text{ which is defined in (7).} \end{aligned} \right\} \quad (10)$$

Proof. We prove the first statement for a particular denominator matrix D of a fraction description of $X^{\text{Tr}}(\lambda I - A)^{-1}Y$. Indeed, if the i -th invariant factors of D divide s_i then, by Lemma 2 and using the product argument given in the proof of Lemma 3, the same holds by transitivity of division for $F_X^{A,Y}$. When Y has rank $r < m$, one may introduce an invertible transformation $Q \in \mathbb{K}^{m \times m}$ such that $YQ = [Y_1 \ 0]$ with $Y_1 \in \mathbb{K}^{n \times r}$. From there, if $X^{\text{Tr}}(\lambda I - A)^{-1}Y_1 = N_1 D_1^{-1}$ then

$$X^{\text{Tr}}(\lambda I - A)^{-1}Y = \begin{bmatrix} N_1(\lambda) & 0 \end{bmatrix} \begin{bmatrix} D_1(\lambda) & 0 \\ 0 & I \end{bmatrix}^{-1} Q^{-1}$$

and the invariant factors of the denominator matrix $Q \operatorname{diag}(D_1, I)$ are those of D_1 . We can thus without loss of generality assume that Y has full column rank. Let us now construct a fraction description of $X^{\operatorname{Tr}}(\lambda I - A)^{-1}Y$ with D as announced. Choose $Y_c \in \mathbb{K}^{n \times (n-m)}$ such that $T = \begin{bmatrix} Y & Y_c \end{bmatrix}$ is invertible in $\mathbb{K}^{n \times n}$ and let $D \in \mathbb{K}[\lambda]^{m \times m}$ be defined from a unimodular triangularization of $T^{-1}(\lambda I - A)$, that is:

$$T^{-1}(\lambda I - A)U(\lambda) = \begin{bmatrix} D(\lambda) & H_{12}(\lambda) \\ 0 & H_{22}(\lambda) \end{bmatrix} \quad (11)$$

with U unimodular. If V is the matrix formed by the first m columns of U we have the fraction descriptions $(\lambda I - A)^{-1}Y = VD^{-1}$ and $X^{\operatorname{Tr}}(\lambda I - A)^{-1}Y = (X^{\operatorname{Tr}}V)D^{-1}$. Thus D is a denominator matrix for $X^{\operatorname{Tr}}(\lambda I - A)^{-1}Y$. By (11) and Lemma 3, its i -th invariant factor divide the i -th invariant factor s_i of $\lambda I - A$ and the first assertion is proven.

To establish the rest of the theorem we work with the associated block Hankel matrix $\operatorname{Hk}_{e,d}(A, X, Y)$. By definition of the invariant factors we know that

$$\dim \operatorname{span}(X, A^{\operatorname{Tr}}X, (A^{\operatorname{Tr}})^2X, \dots) \leq \deg(s_1) + \dots + \deg(s_{\min\{l, \phi\}})$$

and

$$\dim \operatorname{span}(Y, AY, A^2Y, \dots) \leq \deg(s_1) + \dots + \deg(s_{\min\{m, \phi\}})$$

thus

$$\operatorname{rank} \operatorname{Hk}_{e,d}(A, X, Y) \leq \operatorname{rank} \left(\begin{bmatrix} X^{\operatorname{Tr}} \\ X^{\operatorname{Tr}}A \\ X^{\operatorname{Tr}}A^2 \\ \vdots \end{bmatrix} \cdot [Y \ AY \ A^2Y \ \dots] \right) \leq \bar{\nu},$$

where $\bar{\nu} = \deg(s_1) + \dots + \deg(s_{\min\{m, l, \phi\}})$. Hence, from the specializations W and Z of X and Y given in [Villard 1997b: Corollary 6.4], we get

$$\operatorname{rank} \operatorname{Hk}_{e_0, d_0}(A, W, Z) = \max_{X, Y, d, e} \operatorname{rank} \operatorname{Hk}_{e, d+1}(A, X, Y) = \bar{\nu} \quad (12)$$

with $d_0 = \lceil \bar{\nu}/m \rceil$ and $e_0 = \lceil \bar{\nu}/l \rceil$ and thus $\bar{\nu} = \nu$. Using (6) we also have

$$\deg_{\lambda}(\det(F_W^{A, Z}(\lambda))) = \max_{X, Y} \deg_{\lambda}(\det(F_X^{A, Y}(\lambda))) = \bar{\nu}. \quad (13)$$

With (12) and (13) we have proven the two maximality assertions. In addition, since the i -th invariant factor \bar{s}_i of $F_W^{A, Z}$ must divide s_i , the only way to get $\deg_{\lambda}(\det F_W^{A, Z}) = \nu$, is to take $\bar{s}_i = s_i$ for $1 \leq i \leq \min\{m, l, \phi\}$ and $\bar{s}_i = 1$ for $\min\{m, l, \phi\} < i \leq m$. \square

As already noticed, the existence of such W, Z establishes maximality of the matrix generator for symbolic X and Y and, by the Schwartz/Zippel lemma, for random projection matrices. In next sections we will use $\det F_W^{A, Z}(\lambda) = \det(\lambda I - A)$ for computing the determinant and the characteristic polynomial of matrices A such that $\phi \leq \min\{l, m\}$. For general matrices we will use $F_W^{A, Z}$ to determine the first $\min\{l, m\}$ invariant factors of A .

3 Normal matrix polynomial remainder sequences

As done for a scalar sequence in [Sugiyama et al. 1975; Brent et al. 1980; Dornstetter 1987] the minimum matrix generating polynomial of a sequence can be computed *via* a specialized matrix Euclidean algorithm [Coppersmith 1994; Thomé 2002]. Taking advantage of fast matrix multiplication algorithms requires to extend these approaches. In Section 3.1 we propose a matrix Euclidean algorithm which combines fast matrix multiplication with the recursive Knuth/Schönhage half-GCD algorithm [Knuth 1970; Schönhage 1971; Moenck 1973; von zur Gathen and Gerhard 1999]. This is applicable to computing the matrix minimum polynomial of a sequence $\{X^{Tr}AY\}_{i \geq 0}$ if the latter leads to a normal matrix polynomial remainder chain. We show in Section 3.2 that this is satisfied, with high probability, by our random integer sequences. This will be satisfied by construction by the sequence in the division-free computation. For simplicity we work in the square case $l = m$ thus with a sequence $\{B^{[i]}\}_{i \geq 0}$ of matrices in $\mathbb{K}^{m \times m}$.

3.1 Minimum polynomials and half Euclidean algorithm

If $F = \sum_{i=0}^d F^{[i]} \lambda^i \in \mathbb{K}[\lambda]^{m \times m}$ is a generating matrix polynomial for $\{B^{[i]}\}_{i \geq 0}$ then, as we have seen with (5), we have

$$\begin{bmatrix} B^{[0]} & B^{[1]} & \dots & B^{[d]} \\ B^{[1]} & B^{[2]} & \dots & B^{[d+1]} \\ \vdots & \vdots & \ddots & \vdots \\ B^{[d-1]} & B^{[d+1]} & \dots & B^{[2d-1]} \end{bmatrix} \begin{bmatrix} F^{[0]} \\ F^{[1]} \\ \vdots \\ F^{[d]} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \quad (14)$$

The left side matrix was denoted by $\text{Hk}_{d,d+1}$ in (4). We define \hat{B} in $\mathbb{K}[\lambda]^{m \times m}$ by $\hat{B} = \sum_{i=0}^{2d-1} B^{[2d-i-1]} \lambda^i$. Identity (14) is satisfied if and only if there exists matrices S and T of degree less than $d-1$ in $\mathbb{K}[\lambda]^{m \times m}$ such that

$$\lambda^{2d} S(\lambda) + \hat{B}(\lambda) F(\lambda) = T(\lambda). \quad (15)$$

Thus $\lambda^{2d} I$ and \hat{B} may be considered as the inputs of an extended Euclidean scheme. In the scalar case, the remainder sequence of the Euclidean algorithm is said to be normal when at each step the degree is decreased by 1 exactly. By the theorem of subresultants, the remainder sequence is normal if and only if the subresultants are non-zero [Brown and Traub 1971]. In an analogous way we will identify normal matrix remainder sequences related to the computation of matrix generating polynomials. We use these remainder sequences to establish a recursive algorithm based on fast matrix polynomial multiplication.

For two matrices $M = \sum_{i=0}^{2d} M^{[i]} \lambda^i$ and $N = \sum_{i=0}^{2d-1} N^{[i]} \lambda^i$ in $\mathbb{K}[\lambda]^{m \times m}$, if the leading matrix $N^{[2d-1]}$ is invertible in $\mathbb{K}^{m \times m}$ then one can divide M by N in an obvious way to get:

$$\begin{cases} M = NQ + R, \text{ with } \deg Q = 1, \deg R \leq 2d - 2, \\ Q = (N^{[2d-1]})^{-1} (M^{[2d]} \lambda + M^{[2d-1]} - N^{[2d-2]} (N^{[2d-1]})^{-1} M^{[2d]}). \end{cases} \quad (16)$$

If the leading matrix coefficient of R is invertible (matrix coefficient of degree $2d - 2$), then the process can be continued. The remainder sequence is normal if all matrix remainders have invertible leading matrices, if so we define:

$$\begin{cases} M_{-1} = M, M_0 = N \\ M_i = M_{i-2} - M_{i-1}Q_i, 1 \leq i \leq d \end{cases} \quad (17)$$

with $\deg M_i = 2d - 1 - i$. The above recurrence relations define matrices S_i and F_i in $\mathbb{K}[\lambda]^{m \times m}$ such that

$$M_{-1}(\lambda)S_i(\lambda) + M_0(\lambda)F_i(\lambda) = M_i(\lambda), 1 \leq i \leq d, \quad (18)$$

S_i has degree $i - 1$ and F_i has degree i . We also define $S_{-1} = I, S_0 = 0, F_{-1} = 0$ and $F_0 = I$. As shown below, the choice $M_{-1} = \lambda^{2d}I$ and $M_0 = \hat{B}$ leads to a minimum matrix generating polynomial $F = F_d$ for the sequence $\{B^{[i]}\}_{i \geq 0}$ (compare (18) and (15)).

Theorem 5 *Let \hat{B} be the matrix polynomial $\sum_{i=0}^{2d-1} B^{[2d-i-1]}\lambda^i \in \mathbb{K}^{m \times m}$. If for all $1 \leq k \leq d$ we have $\det(\text{Hk}_{k,k}) \neq 0$, then the half matrix Euclidean with $M_{-1} = \lambda^{2d}I$ and $M_0 = \hat{B}$ works as announced. In particular:*

i) M_i has degree $2d - 1 - i$ ($0 \leq i \leq d$) and its leading matrix $M_i^{[2d-1-i]}$ is invertible ($1 \leq i \leq d - 1$);

ii) F_i has degree i and its leading matrix $F_i^{[i]}$ is invertible ($0 \leq i \leq d$); S_i has degree $i - 1$ ($1 \leq i \leq d$).

The algorithm produces a minimum matrix generating polynomial $F_d(\lambda)$ for the sequence $\{B^{[i]}\}_{0 \leq i \leq 2d-1}$ and $F = (F_d^{[d]})^{-1}F_d(\lambda)$ is the unique one in Popov normal form.

Furthermore, if in the half matrix Euclidean algorithm the conditions i-ii are met for all i with $1 \leq i \leq d$, then $\det(\text{Hk}_{k,k}) \neq 0$ for all $1 \leq k \leq d$.

Proof. We prove the assertions by induction. For $i = 0$, since by assumption $B^{[0]}$ is invertible, M_0 satisfies i). By definition $F_0 = I$ and starting at $i = 1$, $S_1 = I$. Now assume that the properties are true for $i - 1$. Then, following (16),

$$Q_i = \tilde{Q}_i\lambda + \bar{Q}_i = \left(M_{i-1}^{[2d-i]}\right)^{-1} M_{i-2}^{[2d-i+1]}\lambda + \bar{Q}_i,$$

\tilde{Q}_i is invertible by i) at previous steps and \bar{Q}_i is in $\mathbb{K}^{m \times m}$. The leading matrix of F_i is

$$F_i^{[i]} = -F_{i-1}^{[i-1]}\tilde{Q}_i$$

thus F_i satisfies ii). The same argument holds for S_i ($i - 1 \geq 1$). By construction M_i has a degree lower than $2d - 1 - i$ hence, looking at the right side coefficient

matrices of (18), we know that

$$\underbrace{\begin{bmatrix} B^{[0]} & B^{[1]} & \dots & B^{[i]} \\ B^{[1]} & B^{[2]} & \dots & B^{[i+1]} \\ \vdots & \vdots & \ddots & \vdots \\ B^{[i]} & B^{[i+1]} & \dots & B^{[2i]} \end{bmatrix}}_{\text{Hk}_{i+1,i+1}} \begin{bmatrix} F_i^{[0]} \\ F_i^{[1]} \\ \vdots \\ F_i^{[i]} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ M_i^{[2d-1-i]} \end{bmatrix}. \quad (19)$$

By assumption of non-singularity of $\text{Hk}_{i+1,i+1}$ and since we have proved that $F_i^{[i]}$ is invertible, the columns in the right side matrix of (19) are linearly independent, thus $M_i^{[2d-1-i]}$ is invertible. This proves 1). Identity (18) for $i = d$ also establishes (14) which means that F_d is a matrix generating polynomial for $\{B^{[i]}\}_{0 \leq i \leq 2d-1}$ whose leading matrix $F_d^{[d]}$ is invertible. It follows that $F = (F_d^{[d]})^{-1} F_d(\lambda)$ is in Popov normal form. The minimality comes from the fact that $\text{Hk}_{d,d}$ is invertible and hence no vector generator (column of a matrix generator) can be of degree less than d .

We finally prove that invertible leading coefficient matrices in the Euclidean algorithm guarantee non-singularity for all $\text{Hk}_{k,k}$. To that end, we consider the range of $\text{Hk}_{i+1,i+1}$ in (19). Clearly, the block vector $[0 \ I_m]^{Tr}$ is in the range, since $M_i^{[2d-1-i]}$ is invertible. By induction hypothesis for $\text{Hk}_{i,i}$, we see that the first i block columns of $\text{Hk}_{i+1,i+1}$ can generate $[I_{mi} \ 0]^{Tr}$, where the block zero row at the bottom is achieved by subtraction of appropriate linear combinations of the previous block vector $[0 \ I_m]^{Tr}$. Hence the range of $\text{Hk}_{i+1,i+1}$ has full dimension. \square

For $B^{[i]} = X^{Tr} A^i Y$, $i \geq 0$, the next corollary shows that F is as expected.

Corollary 6 *Let A be in $\mathbb{K}^{n \times n}$, let $B^{[i]} = X^{Tr} A^i Y \in \mathbb{K}^{m \times m}$, $i \geq 0$, and let $\nu = md$ be the determinantal degree $\deg_\lambda(\det F_X^{A,Y})$. If the block Hankel matrix $\text{Hk}_{d,d}(A, X, Y)$ satisfies the assumption of Theorem 5 then $F = F_X^{A,Y}$.*

Proof. We know from (6) that ν is the maximum possible rank for the block Hankel matrices associated to the sequence, thus the infinite one $\text{Hk}_{\infty,d+1}$ satisfies

$$\text{rank } \text{Hk}_{\infty,d+1} = \text{rank} \left(\begin{bmatrix} B^{[0]} & B^{[1]} & \dots & B^{[d]} \\ B^{[1]} & B^{[2]} & \dots & B^{[d+1]} \\ \vdots & \vdots & \ddots & \vdots \end{bmatrix} \right) = \text{rank } \text{Hk}_{d,d+1} = \nu.$$

It follows that $\text{Hk}_{\infty,d+1}$ and $\text{Hk}_{d,d+1}$ have the same nullspace and F , which by Theorem 5 is a matrix generator for the truncated sequence $\{B^{[i]}\}_{0 \leq i \leq 2d-1}$, is a generator for the whole sequence. The argument used for the minimality of F remains valid hence $F = F_X^{A,Y}$. \square

Remark 7 In Theorem 5 and Corollary 6 we have only addressed the case where the target determinantal degree is an exact multiple md of the blocking factor m . This can be assumed with no loss of generality for the algorithms in sections 4 and 5 and the corresponding asymptotic costs in Section 6. Indeed, we will work there with $\nu = n$ and the input matrix A may be padded to $\text{diag}(A, I)$.

In the general case or in practice to avoid padding, the Euclidean algorithm leads to $\text{rank}(M_{d-1}^{[d]}) = \nu \bmod m \leq m$ and requires a special last division step. The minimum generator $F = F_X^{A,Y}$ has degree $d = \lceil \nu/m \rceil$, with column degrees $[\delta_1, \dots, \delta_m] = [d-1, \dots, d-1, d, \dots, d]$ where $d-1$ is repeated $m \lceil \nu/m \rceil - \nu$ times [Villard 1997b: Proposition 6.1]. \square

The above method may be combined with the Knuth [1970]/Schönhage [1971]/Moënck [1973] recursive approach. If ω is the exponent of matrix multiplication then, as soon as the block Hankel matrix has the required rank profile, $F_X^{A,Y}$ may be computed with $(n^\omega d)^{1+o(1)}$ operations in \mathbb{K} . The required FFT-based multiplication algorithms for matrix polynomials are described in [Cantor and Kaltofen 1991].

3.2 Normal matrix remainder sequences over the integers

The normality of the remainder sequence associated to a given matrix A essentially comes from the genericity of the projections. This may be partly seen in the scalar case for Lanczos algorithm from [Eberly and Kaltofen 1997: Lemma 4.1], [Eberly 2002] or [Kaltofen et al. 2000; Kaltofen and Lee 2003] and in the block case from [Kaltofen 1995: Proposition 3] or [Villard 1997b: Proposition 6.1].

We show here that the block Hankel matrix has generic rank profile for generic projections, and then the integer case follows by randomization. We let \mathcal{X} and \mathcal{Y} be two $n \times m$ matrices with indeterminates entries $\xi_{i,j}$ and $v_{i,j}$ for $1 \leq i \leq n$ and $1 \leq j \leq m$. Let also ν be the maximum determinantal degree defined by (10) in Theorem 4.

Lemma 8 *With $d = \lceil \nu/m \rceil$, the block Hankel matrix $\text{Hk}_{d,d}(A, \mathcal{X}, \mathcal{Y})$ has rank ν and its principal minors of order i are non-zero for $1 \leq i \leq \nu$.*

Proof. For simplifying the presentation we only detail the case where ν is a multiple of m (see Remark 7). Let $\text{Kr}_i(A, Z) \in \mathbb{K}^{n \times i}$ be the block Krylov matrix formed by the i first columns of $[Z \ AZ \ \dots \ A^{d-1}Z]$ for $1 \leq i \leq \nu$. The specialization $Z \in \mathbb{K}^{n \times m}$ of \mathcal{Y} given in [Villard 1997b: Proposition 6.1] satisfies

$$\text{rank Kr}_i(A, Z) = i, 1 \leq i \leq \nu. \quad (20)$$

We now argue, by specializing \mathcal{X} and \mathcal{Y} , that the target principal minors are non-zero. If $i \leq m$, using (20) one can find $X \in \mathbb{K}^{n \times i}$ such that the rank of $X^T \text{Kr}_i(A, Z)$ equals i . If $m < i \leq \nu$ then one can find $X \in \mathbb{K}^{n \times m}$ such

that $X^{Tr}Kr_i(A, Z) = [0 \ J_m]$ where J_m is the $m \times m$ reversion matrix. Hence $Hk_{d,d}(A, X, Z)$ has ones on its i th anti-diagonal and zeros above, the corresponding principal minor of order i is $(-1)^i$. \square

The polynomial $\prod_{k=1}^d \det(Hk_{k,k}(A, \mathcal{X}, \mathcal{Y}))$ is non-zero of degree no more $md(d+1)$ in $\mathbb{K}[\dots, \xi_{i,j}, \dots, v_{i,j}, \dots]$. If the entries of X and Y are chosen uniformly and independently from a finite set $S \subset \mathbb{Z}$ then, by the Schwartz/Zippel lemma and Theorem 5, the associated matrix remainder sequence is normal with probability at least $1 - md(d+1)/|S|$.

4 The block baby steps/giant steps determinant algorithm

We shall present our algorithm for integer matrices. Generalizations to other domains, such as polynomial rings, are certainly possible. The algorithm follows the Wiedemann paradigm [Wiedemann 1986: Chapter V] and uses a baby steps/giant steps approach for computing the sequence elements [Kaltofen 1992]. In addition, the algorithm blocks the projections [Coppersmith 1994]. A key ingredient is that from the theory of realizations described in Section 2, it is possible to recover the characteristic polynomial of a preconditioning of the input matrix.

Algorithm *Block Baby Steps/Giant Steps Determinant*

Input: a matrix $A \in \mathbb{Z}^{n \times n}$.

Output: an integer that is the determinant of A , or “failure;” the algorithm fails with probability no more than $1/2$.

Step 0. Let $h = \log_2 \text{Hd}(A)$, where $\text{Hd}(A)$ is a bound on the magnitude of the determinant of A , for instance, Hadamard’s bound (see, for example, [von zur Gathen and Gerhard 1999]). For purpose of guaranteeing the probability of a successful completion, the algorithm uses positive constants $\gamma_1, \gamma'_1 \geq 1$.

Choose a random prime integer $p_0 \leq \gamma'_1 h^{\gamma_1}$ and compute $\det(A) \bmod p_0$ by LU-decomposition over \mathbb{Z}_{p_0} .

If the result is zero, A is most likely singular, and the algorithm calls an algorithm for computing $x \in \mathbb{Z}^n \setminus \{0\}$ with $Ax = 0$, see Remark 12 on page 22 below.

Step 1. Precondition A such that with high probability $\det(\lambda I - A) = s_1(\lambda) \cdots s_{\min\{m, \phi\}}(\lambda)$, where s_1, \dots, s_ϕ are the invariant factors of $\lambda I - A$. We have two very efficient preconditioners at our disposal. The first is $A \leftarrow DA$ where D is a random diagonal matrix with the diagonal entries chosen uniformly and independently from a set S of integers [Chen et al. 2002:

Theorem 4.3]. The second from [Turner 2001] is $A \leftarrow EA$ where

$$E = \begin{bmatrix} 1 & w_1 & 0 & \dots & 0 \\ 0 & \ddots & \ddots & & \vdots \\ \vdots & \ddots & 1 & & w_{n-1} \\ 0 & \dots & 0 & & 1 \end{bmatrix}, \quad w_i \in S.$$

The product DA is slightly cheaper than EA , but recovery of $\det(A)$ requires division by $\det(D)$. Thus, all moduli that divide $\det(D)$ would have to be discarded from the Chinese remainder algorithm below for the first preconditioner. Both preconditioners achieve $s_1(\lambda) = \det(\lambda I - A)$ with probability $1 - O(n^2/|S|)$. Note that A is non-singular. We shall choose $S = \{i \mid -\lfloor \gamma'_2 n^{\gamma_2} \rfloor \leq i \leq \lceil \gamma'_2 n^{\gamma_2} \rceil\}$, where $\gamma_2 \geq 2, \gamma'_2 \geq 1$ are real constants.

Step 2. Let the blocking factors be $l = m = \lceil n^\sigma \rceil$ where $\sigma = 1/3$.

Select random $X, Y \in S^{n \times m}$.

We will compute the sequence $B^{[i]} = X^{Tr} A^i Y$ for all $0 \leq i < \lceil 2n/m \rceil = O(n^{1-\sigma})$ by utilizing our baby steps/giant steps technique [Kaltofen 1992].

Let the number of giant steps be $s = \lceil n^\tau \rceil$, where $\tau = 1/3$, and let the number of baby steps be $r = \lceil 2\lceil n/m \rceil / s \rceil = O(n^{1-\sigma-\tau})$.

Substep 2.1 for $j = 0, 1, \dots, r-1$ Do $V^{[j]} \leftarrow A^j Y$;

Substep 2.2 $Z \leftarrow A^r$;

Substep 2.3. For $k = 0, 2, \dots, s-1$ Do $(U^{[k]})^{Tr} \leftarrow X^{Tr} Z^k$;

Substep 2.4. For $j = 0, 1, \dots, r-1$ Do
For $k = 0, 1, \dots, s-1$ Do $B^{[kr+j]} \leftarrow (U^{[k]})^{Tr} V^{[j]}$.

Step 3. Compute the minimum matrix generator $F_X^{A,Y}(\lambda)$ from the initial sequence segment $\{B^{[i]}\}_{0 \leq i < \lceil 2n/m \rceil}$. Here we can use the method from Section 3, padding the matrix so that m divides n (see Remark 7 on page 14), and return failure whenever the coefficient $F_i^{[i]}$ of the matrix remainder polynomial is singular. For alternative methods, we refer to the Remark 9 below the algorithm.

Step 4. If $\deg(\det F_X^{A,Y}) < n$ return “failure” (this check may be redundant, depending on which method was used in Step 3). Otherwise, since $F_X^{A,Y}(\lambda)$ is in Popov form we know that its determinant is monic and by Theorem 4 we have $\det F_X^{A,Y}(\lambda) = \det(\lambda I - A)$. Return $\det(A) = \Delta(0)$. \boxtimes

Remark 9 As we have seen in Section 2.1 there are several alternatives for carrying out Step 3 [Rissanen 1972; Dickinson et al. 1974; Forney, Jr. 1975; Van Barel and Bultheel 1992; Beckermann and Labahn 1994; Kaltofen 1995; Coppersmith 1994; Thomé 2002; Giorgi et al. 2003]. In Step 4 we require that

$\det F_X^{A,Y}(\lambda) = \det(\lambda I - A)$. In order to achieve the wanted bit complexity, we must stop any of the algorithms after having processed the first $2\lceil n/m \rceil$ elements of (2). The used algorithm then must return a candidate matrix polynomial \tilde{F} . Clearly, if Step 4 exposes $\deg(\det \tilde{F}) < n$ one knows that the randomizations were unlucky. However, if $\deg(\det \tilde{F}) = n$ there still may be the possibility that $\tilde{F} \neq F_X^{A,Y}$ due to a situation where the first $2\lceil n/m \rceil$ elements do not determine the generator, as would be the case in the two examples given in Section 2. In order to achieve the Las Vegas model of randomized algorithmic complexity, verification of the computed generator is thus necessary here. For example, the used algorithm could do so by establishing that $\text{rank Hk}_{\lceil n/m \rceil, \lceil n/m \rceil}(A, X, Y) = n$. Our algorithm from Section 3 implicitly does so via Theorem 5 on page 13. One could do so explicitly by computing the rank of $\text{Hk}_{\lceil n/m \rceil, \lceil n/m \rceil}$ modulo a random prime number.

We remark that the *arithmetic* cost of verifying that the candidate for $F_X^{A,Y}$ is a generator for the block Krylov sequence $\{A^i Y\}_{i \geq 0}$ is the same as step 2. The reduction is seen by applying the transposition principle [Kaltofen 2000: Section 6]: note that computing all $B^{[i]}$ is the block diagonal left product

$$[(X^{Tr})_{1,*} \mid (X^{Tr})_{2,*} \mid \dots] \cdot \begin{bmatrix} \dots AY^i \dots & 0 & 0 & \dots & 0 \\ 0 & \dots AY^i \dots & 0 & \dots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \dots & \dots & \dots AY^i \dots \end{bmatrix},$$

where $(X^{Tr})_{i,*}$ denotes the i -th row of X^{Tr} . Computing $\sum_i A^i Y c^{[i]}$, where $c^{[i]} \in \mathbb{K}^{m \times m}$ are the coefficients of $F_X^{A,Y}$, is the block diagonal right product

$$\begin{bmatrix} \dots AY^i \dots & 0 & 0 & \dots & 0 \\ 0 & \dots AY^i \dots & 0 & \dots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \dots & \dots & \dots AY^i \dots \end{bmatrix} \cdot \begin{bmatrix} (c^{[0]})_{*,1} \\ (c^{[1]})_{*,1} \\ \vdots \\ (c^{[0]})_{*,2} \\ (c^{[1]})_{*,2} \\ \vdots \end{bmatrix},$$

where $(c^{[i]})_{*,j}$ denotes the j -th column of the matrix $c^{[i]}$. One may also develop an explicit baby steps/giant steps algorithm for computing $\sum_i A^i Y c^{[i]}$. However, because the integer lengths of the entries in $c^{[i]}$ are much larger than those of X and Y , we do not know how to keep the bit complexity low enough to allow verification of the candidate generator via verification as a block Krylov space generator. \square

We shall first give the bit complexity analysis for our block algorithm under the assumption that no subcubic matrix multiplication à la Strassen or subquadratic block Toeplitz solver/greatest common divisor algorithm à la Knuth/Schönhage is employed. We will investigate those best theoretically possible running times in Section 6.

Theorem 10 *Our algorithm computes the determinant of any non-singular matrix $A \in \mathbb{Z}^{n \times n}$ with $(n^{3+1/3} \log \|A\|)^{1+o(1)}$ bit operations. Our algorithm utilizes $(n^{1+1/3} + n \log \|A\|)^{1+o(1)}$ random bits and either returns the correct determinant or it returns “failure,” the latter with probability of no more than $1/2$.*

In our analysis, we will use modular arithmetic. The following lemma will be used to establish the probability of getting a good reduction with prime moduli.

Lemma 11 *Let $\gamma \geq 1$, $\gamma' \geq 1$ be positive real constants. Then for all integers $H \in \mathbb{Z}_{\geq 2}$ that with $h = \log_e(H)/0.84 \leq 1.72 \log_2(H)$ satisfy $h \geq 114$ and $\gamma' \leq h^\gamma$ the probability*

$$\text{Prob}(p \text{ divides } H \mid p \text{ a prime integer, } 2 \leq p \leq \gamma' h^\gamma) \leq \frac{5}{2} \frac{\gamma}{\gamma' h^{\gamma-1}}. \quad (21)$$

Proof. We have the following estimates for the distribution of prime numbers:

$$\prod_{\substack{p \text{ prime} \\ p \leq x}} p > e^{C_1 x}, \pi(x) = \sum_{\substack{p \text{ prime} \\ p \leq x}} 1 > \frac{C_2 x}{\log_e x}, \pi(x) < \frac{C_3 x}{\log_e x}$$

where C_1 , C_2 and C_3 are positive constants. Explicit values for C_1 , C_2 and C_3 have been derived. It is shown in [Rosser and Schoenfeld 1962] that we may choose $C_1 = 0.84$ for $x \geq 101$, $C_2 = 1$ for $x \geq 17$, and $C_3 = 1.25$ for $x \geq 114$.

Since we have $\prod_{p \leq h} p > e^{C_1 h} = H$, there are at most $\pi(h) < C_3 h / (\log_e h)$ distinct prime factors in H . The number of primes $\leq \gamma' h^\gamma$ is more than $C_2 \gamma' h^\gamma / (\gamma \log_e h + \log_e \gamma')$, because from our assumptions we have that $\gamma' h^\gamma \geq 114 > 17$. Therefore the probability for a random p to divide H is no more than, using $\log_e \gamma' \leq \gamma \log_e h$,

$$\frac{C_3 h / (\log_e h)}{C_2 \gamma' h^\gamma / (\gamma \log_e h + \log_e \gamma')} \leq \frac{C_3 h / (\log_e h)}{C_2 \gamma' h^\gamma / (2\gamma \log_e h)} \leq \frac{5}{2} \frac{\gamma}{\gamma' h^{\gamma-1}}. \quad \square$$

In the above Lemma 11 we have introduced the constant γ' so that it is possible to choose $\gamma = 1$ and have a positive probability of avoiding a prime divisor of H .

Proof of Theorem 10. The unblocked version of the algorithm is fully analyzed in [Kaltofen 2002] with the additional modification of early termination when the determinant is small. That analysis uses a residue number system (Chinese remaindering) for representing long integers, which we adopt for the blocked algorithm. This adds the bit cost of generating a stream of sufficiently large random primes (including p_0 in Step 0).

Step 0 has by $h = O(n \log(n \|A\|))$, which follows from Hadamard’s bound, the bit complexity $(n^3 + n^2 \log \|A\|)^{1+o(1)}$, the latter term constituting taking every entry of A modulo p_0 . The failure probability of Step 0, that is when $\det(A) \equiv 0 \pmod{p_0}$ for non-singular A , is bounded by Lemma 11. Thus, for

$H = \det(A)$ and appropriate choice of γ_1 and γ'_1 in Step 0 all non-singular matrices will pass with probability no less than $9/10$.

Step 1 increases $\log \|DA\|$ or $\log \|EA\|$ to no more than $O((\log n)^2 \log \|A\|)$ and has bit cost $(n^3 \log \|A\|)^{1+o(1)}$.

Steps 3, and 4 are performed modulo sufficiently many primes p_l so that $\det(A)$ can be recovered via Chinese remaindering. Using $p_l \geq 2$, we obtain the very loose count

$$1 \leq l \leq 2 \log_2(\text{Hd } A) = 2h = O(n \log(n \|A\|)), \quad (22)$$

the factor 2 accounting for recovery of negative determinants. Modular arithmetic becomes necessary for the avoidance of length growth in the scalars in $F_X^{A,Y}$ during Steps 3 and 4. We shall first estimate the probability of success, and then the bit complexity. The probabilistic analysis will also determine the size of the prime moduli.

The algorithm fails if

- I) the preconditioners D or E in Step 1 do not yield $\det(\lambda I - A) = s_1(\lambda) \cdots s_{\min\{m, \phi\}}(\lambda)$, that with probability $\leq O(1/n^{\gamma_2-2})$. As for Step 0, we select the constant γ_2, γ'_2 so that the preconditioners fail with probability $\leq 1/10$.
- II) the projections X, Y in Step 2 do not yield $\text{rank Hk}_{\lceil n/m \rceil, \lceil n/m \rceil}(A, X, Y) = n$. Since for $X = \mathcal{X}$ and $Y = \mathcal{Y}$ with variables $\xi_{i,j}, v_{i,j}$ as entries full rank is achieved (see Section 2), we can consider an $n \times n$ non-singular submatrix $\Gamma(\mathcal{X}, \mathcal{Y})$ of $\text{Hk}_{\lceil n/m \rceil, \lceil n/m \rceil}(A, \mathcal{X}, \mathcal{Y})$. By [DeMillo and Lipton 1978; Zippel 1979; Schwartz 1980] we get

$$\text{Prob}(\det \Gamma(X, Y) = 0 \mid X, Y \in S^{n \times m}) \leq \frac{\deg(\det \Gamma)}{|S|} \leq \frac{2n}{|S|} \leq \frac{1}{\gamma'_2 n^{\gamma_2-1}}.$$

If we use the matrix polynomial remainder sequence algorithm of Section 3 for Step 3, we also fail if $\prod_{1 \leq k < \lceil n/m \rceil} \det(\text{Hk}_{k,k}(A, X, Y)) = 0$, that with probability no more than $n(n/m + 1)/|S| \leq (n^{1-\sigma} + 1)/(2\gamma'_2 n^{\gamma_2-1})$.

Again, the constant γ_2, γ'_2 are chosen so that the probability is $\leq 1/10$.

- III) the computation modulo one of the moduli p_l fails for Step 3 or 4. Then p_l divides $\det \Gamma(A, X, Y)$. We have $\log |\det(\Gamma(A, X, Y))| = (n^2/m \log \|A\|)^{1+o(1)}$. Therefore we select the random moduli in the range

$$2 \leq p_l \leq \gamma'_3 (n^{2-\sigma} \log \|A\|)^{(1+o(1))\gamma_3} = q \quad (23)$$

where $\sigma = 1/3$ and $\gamma_3 \geq 2, \gamma'_3 \geq 1$ are constants. Note that in (23) the exponent $(1 + o(1))$ captures derivable polylogarithmic factors $C_1(\log n)^{C_2} \times (\log \|A\|)^{C_3}$, where C_1, C_2, C_3 are explicit constants. By Lemma 11 the probability that any one of the $\leq 2h$ moduli fails, i.e. divides $\det(\Gamma(A, X, Y))$, is no more than $2h / (n^{2-\sigma} \log \|A\|)^{(1+o(1))(\gamma_3-1)}$. By the Hadamard estimate (22) we can make this probability no larger than $1/10$ via selecting the constants γ_3, γ'_3 sufficiently large.

If we also must avoid divisors of $\prod_{1 \leq k < \lceil n/m \rceil} \det(\text{Hk}_{k,k}(A, X, Y))$ for the matrix polynomial remainder sequence algorithm, the range (23) increases to $p_l \leq \gamma'_3 (n^{3-2\sigma} \log \|A\|)^{(1+o(1))\gamma_3}$.

- IV) the algorithms fails to compute sufficiently many random prime moduli $p_l \leq q$ (see (23)). There is now a deterministic algorithm of bit complexity $(\log p_l)^{12+o(1)}$ for primality testing [Agrawal et al. 2002], which is not required but simplifies the theoretical analysis here. We pick $k = 4h \log q$ positive integers $\leq q$. The probability for each to be prime is $\geq 1/\log q = \psi$ (provided $q \geq 17$ [Rosser and Schoenfeld 1962]). By Chernoff bounds for the tail of the binomial distribution, the probability that fewer than $2h = (1 - 1/2)\psi k$ are prime is $\leq e^{-(1/2)^2 \psi k/2} = 1/e^{h/2}$. Thus for $h \geq 5$ the probability of failing to find $2h$ primes is $\leq 1/10$.

The cases I-IV together with Step 0 add up to a failure probability of $\leq 1/2$. We conclude by estimating the number of bit operations for Steps 2-4.

Step 2 computes $B^{[i]} \bmod p_l$ for $0 \leq i < 2\lceil n/m \rceil$ and $1 \leq l \leq 2h$ as follows. First, all $B^{[i]}$ are computed as exact integers. For substeps 2.1 and 2.2 that requires $O(n^3 \log r)$ arithmetic operations on integers of length $(r \log \|A\|)^{1+o(1)}$, in total $(n^{4-\sigma-\tau} \log \|A\|)^{1+o(1)}$ bit operations (recall that $\sigma = \tau = 1/3$). Substeps 2.3 and 2.4 require $O(sm n^2)$ arithmetic operations on integers of length $(rs \log \|A\|)^{1+o(1)}$, again $(n^{3+\tau} \log \|A\|)^{1+o(1)}$ bit operations. Then all $O(n/m \times m^2)$ entries of all $B^{[i]}$ are taken modulo p_l with l in the range (22) and p_l in (23). Straight-forward remaindering would yield a total of $(nmhrs \log \|A\|)^{1+o(1)}$ bit operations, which is $(n^3 (\log \|A\|)^2)^{1+o(1)}$. The complexity can be reduced to $(n^3 \log \|A\|)^{1+o(1)}$ via a tree evaluation scheme [Heindel and Horowitz 1971; Aho et al. 1974: Algorithm 8.4].[†]

Steps 3 and 4 are performed modulo all $O(h)$ prime moduli p_l . For each prime the cost of extended Euclidean algorithm on matrix polynomials is $O(m^3 (n/m)^2)$ residue operations. Overall, the bit complexity of Steps 3 and 4 is again $(n^{3+\sigma} \log \|A\|)^{1+o(1)}$. The number of required random bits in D or E , X and Y , and case IV above is immediate. \square

It is possible to derive explicit values for the constants $\gamma_1, \gamma'_1, \gamma_2, \gamma'_2, \gamma_3$, and γ'_3 so that Theorem 10 holds. However, any implementation of the algorithm would select reasonably small values. For example, all prime moduli would be chosen 32 or 64 bit in length. Since the method is Las Vegas, such choice only effects the probability of not obtaining a result.

If Step 3 uses a Knuth/Schönhage half-GCD approach with FFT-based polynomial arithmetic for the Euclidean algorithm on matrix polynomials of Section 3, the complexity for each modulus reduces to $(m^2 n)^{1+o(1)}$ residue operations. Thus, the overall complexity of Steps 3 and 4 reduces to $(n^{2+2\sigma} \times \log \|A\|)^{1+o(1)}$ bit operations. For $\sigma = 3/5$ and $\tau = 1/5$ the bit complexity of the algorithm then is $(n^{3+1/5} \log \|A\|)^{1+o(1)}$ (cf. [Kaltofen and Villard 2002] and [Pan 2002]).

[†]Note that this speedup comes at a cost of an extra log-factor.

Remark 12 In order to state a Las Vegas bit complexity for the determinant of a general square matrix, we need to consider the cost of certifying singularity in Step 0 on page 16 above. In order to meet the complexity of Theorem 10 on page 19 above we can use the algorithm in [Dixon 1982]. Reduction to a non-singular subproblem can be accomplished by methods in [Kaltofen and Saunders 1991], and the rank is determined in a Monte Carlo manner via a random prime modulus; see also [Villard 1988: page 102].

5 Improved division-free complexity

Our baby steps/giant steps algorithm with blocking of Section 4 can be employed to improve the division-free complexity of the determinant of [Kaltofen 1992]. Here we consider a matrix $A \in R^{n \times n}$, where R is a commutative ring with a unit element. At task is to compute the determinant of A by ring additions, subtractions and multiplications. Blocking can improve the number of ring operations from $n^{3.5+o(1)}$ [Kaltofen 1992] to $n^{3+1/3+o(1)}$, that without subcubic matrix multiplication or subquadratic Toeplitz/GCD algorithms, and best possible from $O(n^{3.0281})$ [Kaltofen 1992][‡] to $O(n^{2.6973})$. Our algorithm combines the blocked determinant algorithm with the elimination of divisions technique of [Kaltofen 1992]. Our computational model is either a straight-line program/arithmetic circuit or an algebraic random access machine [Kaltofen 1988]. Further problems are to compute the characteristic polynomial and the adjoint matrix of A .

The main idea of [Kaltofen 1992] follows [Strassen 1973] and for the input matrix A computes the determinant of the polynomial matrix $L(z) = M + z(A - M)$, where $M \in \mathbb{Z}^{n \times n}$ is an integral matrix whose entries are independent of the entries in A . For $\Delta(z) = \det(L(z))$ we have $\det(A) = \Delta(1)$. All intermediate elements are represented as polynomials in $R[z]$ or as truncated power series in $R[[z]]$ and the “shift” matrix M determines them in such a manner that whenever a division by a polynomial or truncated power series is performed the constant coefficients are ± 1 . For the algorithm in Section 4 we not only pick M but also concrete projection block vectors $X \in \mathbb{Z}^{n \times m}$ and $Y \in \mathbb{Z}^{n \times m}$. No randomization is necessary, as M is a “good” input matrix ($\phi = m$) and X and Y are “good” projections, we have $\det F_X^{L(z), Y}(\lambda) = \det(\lambda I - L(z))$.

The matrices M , X and Y are block versions of the ones constructed in [Kaltofen 1992]. Suppose that the blocking factor m is a divisor of n , the dimension of A . This we can always arrange by padding A to $\begin{bmatrix} A & 0 \\ 0 & I \end{bmatrix}$. Let $d = n/m$ and let

$$a_i = \binom{i}{\lfloor i/2 \rfloor}, \quad c_i = -(-1)^{\lfloor (d-i+1)/2 \rfloor} \binom{\lfloor (d+i)/2 \rfloor}{i},$$

[‡]The proceedings paper gives an exponent 3.188; the smaller exponent is in a postnote added to the version posted on www.kaltofen.us/bibliography.

and let

$$C = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \ddots & 0 \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & & 0 & 1 \\ c_0 & c_1 & \dots & c_{d-2} & c_{d-1} \end{bmatrix}, v = \begin{bmatrix} a_0 \\ a_1 \\ \vdots \\ a_{d-1} \end{bmatrix}.$$

We show in [Kaltofen 1992] that for the sequence $a_i = e_1^{Tr} C^i v$, where $e_1^{Tr} = [1 \ 0 \ \dots \ 0] \in \mathbb{Z}^{1 \times d}$ is the first d -dimensional unit (row) vector, then the Berlekamp/Massey algorithm divides by only ± 1 . We define

$$M = \begin{bmatrix} C & 0 & \dots & 0 \\ 0 & C & \ddots & 0 \\ \vdots & 0 & \ddots & \vdots \\ 0 & \dots & 0 & C \end{bmatrix} \in \mathbb{Z}^{n \times n},$$

$$X = \begin{bmatrix} e_1 & 0 & \dots & 0 \\ 0 & e_1 & \ddots & 0 \\ \vdots & 0 & \ddots & \vdots \\ 0 & \dots & & e_1 \end{bmatrix} \in \mathbb{Z}^{n \times m}, Y = \begin{bmatrix} v & 0 & \dots & 0 \\ 0 & v & \ddots & 0 \\ \vdots & 0 & \ddots & \vdots \\ 0 & \dots & & v \end{bmatrix} \in \mathbb{Z}^{n \times m}.$$

By construction, the algorithm for computing the determinant of Section 4 performed now with the matrices X, M, Y results in a minimum matrix generator

$$F_X^{M,Y}(\lambda) = (\lambda^d - c_{d-1}\lambda^{d-1} - \dots - c_0)I_m,$$

where I_m is an $m \times m$ identity matrix. Furthermore, this generator can be computed from the sequence of block vectors $B^{[i]} = a_i I_m$ by a matrix Euclidean algorithm (see Section 3) in which all leading coefficient matrices are equal to $\pm I_m$.

The arithmetic cost for executing the block baby steps/giant steps algorithm on the polynomial matrix $L(z) = M + z(A - M)$ is related to the bit complexity of Section 4. Now the intermediate lengths are the degrees in z of the computed polynomials in $R[z]$. Therefore, the matrices $X^{Tr} L(z)^i Y \in R[z]^{m \times m}$ can be computed for all $0 \leq i < 2d$ in $n^{3+1/3+o(1)}$ ring operations. In the matrix Euclidean algorithm for Step 3 we perform truncated power series arithmetic modulo z^{n+1} . The arithmetic cost is $(d^2 m^3 n)^{1+o(1)}$ ring operations for the classical Euclidean algorithm with FFT-based power series arithmetic. For the latter, we employ a division-free FFT-based polynomial multiplication algorithm [Cantor and Kaltofen 1991]. Finally, for obtaining the characteristic polynomial, we may slightly extend Step 4 on page 17 and compute the entire determinant of $F_X^{L(z),Y}(\lambda)$ division-free in truncated power series arithmetic over $R[z, \lambda] \bmod (z^{n+1}, \lambda^{n+1})$. For this last step we can use our original division-free algorithm [Kaltofen 1992] and achieve arithmetic complexity $(m^{3.5} n^2)^{1+o(1)}$. We have proven the following theorem.

Theorem 13 *Our algorithm computes the characteristic polynomial of any matrix $A \in R^{n \times n}$ with $(n^{3+1/3})^{1+o(1)}$ ring operations in R . By the results in [Baur and Strassen 1983] the same complexity is obtained for the adjoint matrix, which can be symbolically defined as $\det(A)A^{-1}$.*

6 Using fast matrix multiplication

As stated in the Introduction, by use of sub-cubic matrix multiplication algorithms the worst case bit complexity of the block algorithms in Section 4 and 5 can be brought below cubic complexity in n . We note that taking the n^2 entries of the input matrix modulo n prime residues is already a cubic process in n ; our algorithms therefore proceed differently.

Now let ω be the exponent for fast matrix multiplication. By [Coppersmith and Winograd 1990] we may set $\omega = 2.375477$. The considerations in this section are of a purely theoretical nature.

Substep 2.1 in Section 4 is done by repeated doubling as in

$$\left[A^{2^\mu} Y \ A^{2^{\mu+1}} Y \ \dots \ A^{2^{\mu+1}-1} Y \right] = A^{2^\mu} \left[Y \ AY \ \dots \ A^{2^\mu-1} Y \right] \text{ for } \mu = 0, 1, \dots$$

Therefore the bit complexity for Substeps 2.1 and 2.2 is $(n^\omega r \log \|A\|)^{1+o(1)}$ with an exponent $\omega + 1 - \sigma - \tau$ for n . Note that σ and τ determine the blocking factor and number of giant steps, and will be chosen later so as to minimize the complexity.

Substep 2.3 both splits the integer entries in $U^{[k]}$ into chunks of length $(r \log \|A\|)^{1+o(1)}$, which is the bit length of the entries in Z . There are at most $s^{1+o(1)}$ such chunks. Thus each block vector times matrix product $(U^{[k]})^{Tr} Z$ is a rectangular matrix product of dimensions $(ms)^{1+o(1)} \times n$ by $n \times n$. We now appeal to fast methods for rectangular matrices [Coppersmith 1997] (we seem not to need the results in [Huang and Pan 1998]), which show how to multiply an $n \times n$ matrix by an $n \times \nu$ matrix in $n^{\omega-\theta+o(1)} \nu^{\theta+o(1)}$ arithmetic operations (by blocking the $n \times n$ matrix into $(t \times t)$ -sized blocks and the $n \times \nu$ matrix into $(t \times t^\zeta)$ -sized blocks such that $n/t = \nu/t^\zeta$ and that the individual block products only take $t^{2+o(1)}$ arithmetic steps each), where $\theta = (\omega - 2)/(1 - \zeta)$ with $\zeta = 0.2946289$. There are s such products on integers of length $(r \log \|A\|)^{1+o(1)}$, so the bit complexity for Substep 2.3 is $(sn^{\omega-\theta} (ms)^\theta r \log \|A\|)^{1+o(1)}$ with an exponent $\omega + 1 - \sigma + (\sigma + \tau - 1)\theta$ for n .

Step 3 for each individual modulus can be performed by the method presented in Section 3 in $(m^\omega n/m)^{1+o(1)}$ residue operations. For all $\leq 2h$ moduli we get a total bit complexity for Step 3 of $(m^{\omega-1} n^2 \log \|A\|)^{1+o(1)}$ with an exponent $2 + \sigma(\omega - 1)$ for n .

The bit complexities of Substep 2.4 and Step 4 are dominated by the complexities of other steps.

All of the above bit costs lead to total bit complexity of $(n^\eta \log \|A\|)^{1+o(1)}$ where the exponent η depends on the use matrix multiplication exponents ω and ζ . Table 1 displays the optimal values of η for selected exponents together

	ω	ζ	η	σ	τ
1	ω	ζ	$\omega + \frac{1-\zeta}{\omega^2-(2+\zeta)\omega+2}$	$1 - \frac{\omega-(1+\zeta)}{\omega^2-(2+\zeta)\omega+2}$	$\frac{\omega-2}{\omega^2-(2+\zeta)\omega+2}$
2	2.375477	0.2946289	2.697263	0.506924	0.171290
3	ω	0	$\omega + \frac{1}{(\omega-1)^2+1}$	$1 - \frac{\omega-1}{(\omega-1)^2+1}$	$\frac{\omega-2}{(\omega-1)^2+1}$
4	3	0	$3 + \frac{1}{5}$	$\frac{3}{5}$	$\frac{1}{5}$
5	$\log_2(7)$	0	3.041738	0.576388	0.189230
6	2.375477	0	2.721267	0.524375	0.129836
7	2	0	$2 + \frac{1}{2}$	$\frac{1}{2}$	0

Table 1: Determinantal bit/division-free complexity exponent η .

with the exponents for the blocking factor and giant stepping that achieve the optimum. Line 1 is the symbolic solution, Line 2 gives the best exponent that we have achieved. Line 3 is the solution without appealing to faster rectangular matrix multiplication schemes. Line 4 corresponds to the comments before Remark 12 on page 22, and Line 5 uses Strassen's original subquadratic matrix multiplication algorithm. Line 6 exhibits the slowdown without faster rectangular matrix multiplication algorithms. Line 7 is our complexity for a hypothetical quadratic matrix multiplication algorithm.

An issue arises whether the singularity certification in Step 0 of our algorithm can be accomplished at a matching or lower bit complexity than the ones given above for the determinant. We refer to possible approaches in [Mulders and Storjohann 2000; Storjohann 2003].

The above analysis applies to our algorithm in Section 5 and yields for the determinant and adjoint matrix a division-free complexity of $O(n^{2.697263})$ ring operations. To our knowledge, this is the best-known to-date. A complication arises in Step 4 when the entire characteristic polynomial is to be computed without divisions. The computation of $\det F_X^{L(z),Y}(\lambda) \bmod (z^{n+1}, \lambda^{n+1})$ seems to require $O(m^\omega)$ operations modulo (z^{n+1}, λ^{n+1}) . We note that for $\lambda = z = 0$ the generator matrix polynomial evaluates to I_m , so asymptotically fast LU-decomposition algorithms are applicable [Aho et al. 1974]. Step 4 now needs $(n^{\sigma\omega+2})^{1+o(1)}$ ring operations, reducing the division-free complexity for the characteristic polynomial to $n^{\chi+o(1)}$, $\chi = \omega + \frac{2(1-\zeta)}{\omega^2-(1+\zeta)\omega+1-\zeta}$ ring operations. We obtain with $\omega = 2.375477$ and $\zeta = 0.2946289$ at $\sigma \approx 0.339517$ and $\tau \approx 0.229446$ a division-free complexity for the characteristic polynomial of $O(n^{2.806515})$ ring operations.

A Maple 7 worksheet that contains our exponent calculations is posted at <http://www.kaltofen.us/bibliography>.

7 Integer characteristic polynomial and normal forms

As already seen in Sections 5 and 6 over an abstract ring R , our determinant algorithm also computes the adjoint matrix and the characteristic polynomial. In the case of integer matrices, although differently from the algebraic setting, the algorithm of Section 4 may also be extended to solving other problems. We briefly mention two extensions in the following. For $A \in \mathbb{Z}^{n \times n}$ we shall first see that the algorithm leads to the characteristic polynomial of a preconditioning of A and consequently to the Smith normal form of A . We shall then see how $F_X^{A,Y}$ may be used for computing the Frobenius normal form of A and hence its characteristic polynomial. Note that the exponents in our bit complexity are of the same order than those discussed for the determinant problem in Table 1.

7.1 Smith normal form of integer matrices

A randomized Monte Carlo algorithm for computing the Smith normal form $S \in \mathbb{Z}^{n \times n}$ of an integer matrix $A \in \mathbb{Z}^{n \times n}$ of rank r may be designed by combining the algorithm of Section 4 with the approach of [Giesbrecht 2001]. Here we improve on the best previously known randomized algorithm of [Eberly et al. 2000]. The current estimate for a deterministic computation of the form is $(n^{\omega+1} \log \|A\|)^{1+o(1)}$ [Storjohann 1996].

The Smith normal form over \mathbb{Z} is defined in a way similar to what we have seen in Section 2.2 for polynomial matrices. The Smith form S is an equivalent diagonal matrix in $\mathbb{Z}^{n \times n}$, with diagonal elements $s_1, s_2, \dots, s_r, 0, \dots, 0$ such that s_i divides s_{i-1} for $2 \leq i \leq r$. The s_i 's are the invariant factors of A [Newman 1972].

Giesbrecht's approach reduces the computation of S to the computation of the characteristic polynomials of matrices $D_1^{(i)} T^{(i)} D_2^{(i)} A$ for $l = (\log n + \log \log \|A\|)^{1+o(1)}$ random choices of diagonal matrices $D_1^{(i)}$ and $D_2^{(i)}$ and of Toeplitz matrices $T^{(i)}$, $1 \leq i \leq l$. The invariant factors may be computed from the coefficients of these characteristic polynomials. The preconditioning $B \leftarrow D_1^{(i)} T^{(i)} D_2^{(i)} A$ ensures that the minimum polynomial f^B of B is squarefree [Giesbrecht 2001: Theorem 1.4] (see also [Chen et al. 2002] for such preconditionings). Hence if \bar{f}^B denotes the largest divisor of f^B such that $\bar{f}^B(0) \neq 0$, we have $r = \text{rank } B = \deg \bar{f}^B$ which is $-1 + \deg f^B$ if A is singular. By Theorem 4, for random X and Y we shall have, with high probability, $\Delta(\lambda) = \det(F_X^{B,Y}(\lambda)) = \lambda^{k_1} f^B(\lambda) = \lambda^{k_2} \bar{f}^B(\lambda)$ for two positive integers k_1 and k_2 that depend on the rank and on the blocking factor m . The needed characteristic polynomials $\lambda^{n-r} \bar{f}^B$ and then the Smith form are thus obtained from the determinants of l matrix generating polynomials.

To ensure a high probability of success, the computations are done with $D_1^{(i)}, D_2^{(i)}$ and $T^{(i)}$ chosen over a ring extension $R_{\mathbb{Z}}$ of degree $O((\log n)^2)$ of \mathbb{Z} , in combination with Chinese remaindering modulo $(n \log \|A\|)^{1+o(1)}$ primes [Giesbrecht 2001: Theorem 4.2]. For one choice of $B^{(i)}$, the cost overhead compared to Step 4 in Section 4 is the one for computing the entire determinant of

the $m \times m$ matrix polynomial $F_X^{B^{(i)}, Y}$ of degree $d = \lceil n/m \rceil$. Over a field, by [Storjohann 2002: Proposition 24] or [Storjohann 2003: Proposition 41] such a determinant is computed in $(m^\omega d)^{1+o(1)}$ arithmetic operations. Using the $(n \log \|A\|)^{1+o(1)}$ primes and the fact that the ring extension $R_{\mathbb{Z}}$ has degree $O((\log n)^2)$, $\det F_X^{B^{(i)}, Y} \in R_{\mathbb{Z}}[\lambda]$ is thus computed in $(n^{2+\sigma(\omega-1)} \log \|A\|)^{1+o(1)}$ bit operations.

From there we see that the cost for computing the l characteristic polynomials, which is the dominant cost for computing the Smith form, corresponds to the estimate already taken into account for Steps 3 of the determinant algorithm. Hence the values of η in Table 1 remain valid for the computation of the Smith normal form using a randomized Monte Carlo algorithm.

7.2 Integer characteristic polynomial and Frobenius normal form

As used above, a direct application of Section 4 leads to the characteristic polynomial of a preconditioning of A . For computing the characteristic polynomial of A itself, we extend our approach using the Frobenius normal form and the techniques of [Storjohann 2000b]. The Frobenius normal form of $A \in \mathbb{Z}^{n \times n}$ is a block diagonal matrix in $\mathbb{Z}^{n \times n}$ similar to A . Its diagonal blocks are the companion matrices for the invariant factors $s_1(\lambda), \dots, s_\phi(\lambda)$ of $\lambda I - A$. Hence the characteristic polynomial $\det(\lambda I - A) = \prod_{i=1}^{\phi} s_i(\lambda)$ is directly obtained from the normal form. Our result is a randomized Monte Carlo algorithm which improves on previous complexity estimates for computing the characteristic polynomial or the Frobenius normal form over \mathbb{Z} [Storjohann 2000a: Table 10.1]. The certified randomized algorithm of [Giesbrecht and Storjohann 2002] uses $(n^{\omega+1} \log \|A\|)^{1+o(1)}$ bit operations.

By Theorem 4 on page 10, if we avoid the preconditioning step (Step 1) in the determinant algorithm of on page 16 in Section 4, the computation leads to $F_X^{A, Y}(\lambda)$ and to

$$\det(F_X^{A, Y}(\lambda)) = \prod_{i=1}^{\min\{m, \phi\}} s_i(\lambda).$$

The first invariant factor $s_1(\lambda)$ is the minimum polynomial f^A of A , hence $\det(F_X^{A, Y})$ is a multiple of f^A and a factor of the characteristic polynomial in $\mathbb{Z}[\lambda]$. Following the cost analysis of the previous Section 7.1 for the determinant of the matrix generating polynomial, the exponents in Table 1 are thus valid for the computation of $\det(F_X^{A, Y})$. The square free part f_{sqfr}^A of $\det(F_X^{A, Y})$ may be deduced in $(n^2 \log \|A\|)^{1+o(1)}$ bit operations [Gerhard 2001: Theorem 11].

From the Frobenius normal form of A modulo a random prime p , f_{sqfr}^A allows a multifactor Hensel lifting for reconstructing the form over \mathbb{Z} [Storjohann 2000b]. With high probability, $\lambda I - A$ also has ϕ invariant factors modulo p . We denote them by $\bar{s}_1, \dots, \bar{s}_\phi$. They can be decomposed into ϕ products

$$\bar{s}_i = \bar{t}_1^{e_{i1}} \dots \bar{t}_m^{e_{im}}, 1 \leq i \leq \phi,$$

for a GCD-free family $\{\bar{t}_1, \dots, \bar{t}_m\}$ of square free polynomials in $\mathbb{F}_p[\lambda]$ and for indices $(e_{i1}, \dots, e_{im}) \in \mathbb{Z}_{>0}^m$, $1 \leq i \leq \phi$. This decomposition is computed in $(n^2 \log p)^{1+o(1)}$ bit operations [Bach and Shallit 1996: Section 4.8]. With high probability we also have

$$\bar{t}_1 \bar{t}_2 \dots \bar{t}_m = f_{\text{sqfr}}^A \pmod{p}.$$

The latter factorization can be lifted, for instance using the algorithm of [von zur Gathen and Gerhard 1999: §15.5], into a family $\{t_1, \dots, t_m\}$ of polynomials modulo a sufficiently high power k of p . With high probability, the invariant factors of $\lambda I - A$ over \mathbb{Z} and the Frobenius form of A may finally be obtained as the following combinations of the t_i 's:

$$s_i = t_1^{e_{i1}} \dots t_m^{e_{im}} \pmod{p^k}, 1 \leq i \leq \phi,$$

with coefficients reduced in the symmetric range.

In addition to the computation of $F_X^{A,Y}(\lambda)$, the dominant cost is the cost of the lifting. Any divisor of the characteristic polynomial has a coefficient size in $(n \log \|A\|)^{1+o(1)}$ (for instance see [Giesbrecht and Storjohann 2002: Lemma 2.1]) hence one can take $k = (n \log \|A\|)^{1+o(1)}$. The polynomials t_1, \dots, t_m are thus computed in $(n^2 \log \|A\|)^{1+o(1)}$ bit operations [von zur Gathen and Gerhard 1999: Theorem 15.18]. We may conclude that the values of the exponent of η in Table 1 are valid for the randomized computation of the Frobenius normal form and the characteristic polynomial of an integer matrix.

Victor Pan has brought to our attention Theorem 5.4 in [Pan 2002]. There a Las Vegas bit complexity of $(n^{16/5} \log \|A\|)^{1+o(1)}$ is stated for the Frobenius factors of a matrix $A \in \mathbb{Z}^{n \times n}$ by a different method. Pan has told us that the result is actually Monte Carlo. We have been unable to verify that Pan's algorithm has the stated bit complexity.

8 Concluding Remarks

Our baby steps/giant steps and blocking techniques apply to entry domains other than the integers, like polynomial rings and algebraic number rings. We would like to add that if the entries are polynomials over a possibly finite field, there are additional new techniques possible [Storjohann 2002; Jeannerod and Villard 2002; Mulders and Storjohann 2003; Storjohann 2003]. In [Storjohann 2003: Section 18] it is suggested that the results for polynomial matrices can be adapted to matrices with integral entries, thus yielding a Las Vegas algorithm that computes $\det(A)$ where $A \in \mathbb{Z}^{n \times n}$ in $(n^\omega \log \|A\|)^{1+o(1)}$ bit operations, when $n \times n$ matrices are multiplied in $O(n^\omega)$ algebraic operations; Storjohann writes that this will be published in a future paper. The best known division-free complexity of the determinant remains at $O(n^{2.697263})$ as stated in Sections 5 and 6. Furthermore, the best known bit-complexity of the characteristic polynomial of an integer matrix is to our knowledge the one in Section 7.2, namely $(n^{2.697263} \log \|A\|)^{1+o(1)}$.

For the classical matrix multiplication exponent $\omega = 3$, the bit complexity of integer matrix determinants is thus proportional to $n^{\eta+o(1)}$ as follows: $\eta = 3 + \frac{1}{2}$ [Kaltofen 1992; Eberly et al. 2000; Kaltofen 2002], $\eta = 3 + \frac{1}{3}$ (Theorem 10 on page 19), $\eta = 3 + \frac{1}{5}$ (Line 4 in Table 1 on page 25), $\eta = 3$ [Storjohann 2003: Section 18]. Together with the algorithms discussed in Section 1 on page 3 that perform well on propitious inputs, such a multitude of results poses a problem for the practitioner: which of the methods can yield faster procedures in computer algebra systems? With William J. Turner we have implemented the baby steps/giant steps algorithm of [Kaltofen 1992, 2002] in Maple 6 with mixed results in comparison to Gaussian elimination and Chinese remaindering. The main problem seems the overhead hidden in the $n^{o(1)}$ -factor. For example, for $n_1 = 10000$ one has $(\log_2 n_1)/n_1^{1/3} > 0.616$, which means that saving a factor of $n^{1/3}$ at the cost of a factor $\log_2 n$ may for practical considerations be quite immaterial. In addition, one also needs to consider other properties, such as the required intermediate space and whether the algorithm is easily parallelized. We believe that the latter may be the most important advantage in practice of our block approach (cf. [Coppersmith 1994; Kaltofen 1995]).

The reduction of the bit complexity of an algebraic problem below that of its known algebraic complexity times the bit length of the answer should raise important considerations for the design of generic algorithms with abstract coefficient domains [Jenks et al. 1988] and for the interpretation of algebraic lower bounds for low complexity problems [Strassen 1990]. We demonstrate that the interplay between the algebraic structure of a given problem and the bits of the intermediately computed numbers can lead to a dramatic reduction in the bit complexity of a fundamental mathematical computation task.

Acknowledgement

We thank William J. Turner for his observations made on the practicality of our method, Mark Giesbrecht for reporting to us the value of the smallest exponent in [Eberly et al. 2000] prior to its publication, and Elwyn Berlekamp for comments on the Berlekamp/Massey algorithm.

References

Note: many of the authors' publications cited below are accessible through links in their webpages listed under the title.

Abbott, J., Bronstein, M., and Mulders, T. Fast deterministic computation of determinants of dense matrices. In Dooley, S., editor, *ISSAC 99 Proc. 1999 Internat. Symp. Symbolic Algebraic Comput.*, pages 181–188, New York, N. Y., 1999. ACM Press. ISBN 1-58113-073-2.

Agrawal, Manindra, Kayal, Neeraj, and Saxena, Nitin. PRIMES is in

- P. Manuscript, 2002. Available from <http://www.cse.iitk.ac.in/news/primality.pdf>.
- Aho, A., Hopcroft, J., and Ullman, J. *The Design and Analysis of Algorithms*. Addison and Wesley, Reading, MA, 1974.
- Bach, E. and Shallit, J. *Algorithmic Number Theory Volume 1: Efficient Algorithms*. The MIT Press, Cambridge, Massachusetts, USA, 1996.
- Baur, W. and Strassen, V. The complexity of partial derivatives. *Theoretical Comp. Sci.*, 22:317–330, 1983.
- Beckermann, B. and Labahn, G. A uniform approach for fast computation of matrix-type Padé approximants. *SIAM J. Matrix Anal. Applic.*, 15(3): 804–823, July 1994.
- Brent, Richard P., Gao, Shuhong, and Lauder, Alan G. B. Random Krylov spaces over finite fields. *SIAM J. Discrete Math.*, 16(2):276–287, 2003.
- Brent, R. P., Gustavson, F. G., and Yun, D. Y. Y. Fast solution of Toeplitz systems of equations and computation of Padé approximants. *J. Algorithms*, 1:259–295, 1980.
- Brönnimann, H., Emiris, I., Pan, V., and Pion, S. Sign determination in residue number systems. *Theoretical Comput. Sci.*, 210(1):173–197, 1999. Special issue on real numbers and computers.
- Brönnimann, H. and Yvinec, M. Efficient exact evaluation of signs of determinant. *Algorithmica*, 27:21–56, 2000.
- Brown, W. S. and Traub, J. F. On Euclid’s algorithm and the theory of subresultants. *J. ACM*, 18:505–514, 1971.
- Cantor, D. G. and Kaltofen, E. On fast multiplication of polynomials over arbitrary algebras. *Acta Inform.*, 28(7):693–701, 1991.
- Chen, L., Eberly, W., Kaltofen, E., Saunders, B. D., Turner, W. J., and Villard, G. Efficient matrix preconditioners for black box linear algebra. *Linear Algebra and Applications*, 343–344:119–146, 2002. Special issue on *Structured and Infinite Systems of Linear Equations*, edited by P. Dewilde, V. Olshevsky and A. H. Sayed.
- Clarkson, Kenneth L. Safe and efficient determinant evaluation. In *Proc. 33rd Annual Symp. Foundations of Comp. Sci.*, pages 387–395, Los Alamitos, California, 1992. IEEE Computer Society Press.
- Coppersmith, D. Solving homogeneous linear equations over $\text{GF}(2)$ via block Wiedemann algorithm. *Math. Comput.*, 62(205):333–350, 1994.
- Coppersmith, D. Rectangular matrix multiplication revisited. *J. Complexity*, 13:42–49, 1997.

- Coppersmith, D. and Winograd, S. Matrix multiplication via arithmetic progressions. *J. Symbolic Comput.*, 9(3):251–280, 1990. Special issue on complexity theory.
- DeMillo, R. A. and Lipton, R. J. A probabilistic remark on algebraic program testing. *Information Process. Letters*, 7(4):193–195, 1978.
- Dickinson, Bradley W., Morf, Martin, and Kailath, Thomas. A minimal realization algorithm for matrix sequences. *IEEE Trans. Automatic Control*, AC-19(1):31–38, February 1974.
- Dixon, J. Exact solution of linear equations using p -adic expansions. *Numer. Math.*, 40(1):137–141, 1982.
- Dornstetter, J. L. On the equivalence between Berlekamp’s and Euclid’s algorithms. *IEEE Trans. Inf. Theory*, IT-33(3):428–431, 1987.
- Eberly, W. Avoidance of look-ahead in Lanczos by random projections, 2002. Manuscript in preparation.
- Eberly, W., Giesbrecht, M., and Villard, Gilles. On computing the determinant and Smith form of an integer matrix. In *Proc. 41st Annual Symp. Foundations of Comp. Sci.*, pages 675–685, Los Alamitos, California, 2000. IEEE Computer Society Press.
- Eberly, W. and Kaltofen, E. On randomized Lanczos algorithms. In Küchlin [1997], pages 176–183. ISBN 0-89791-875-4.
- Emiris, I. Z. A complete implementation for computing general dimensional convex hulls. *Int. J. Comput. Geom. Appl.*, 8(2):223–254, 1998.
- Forney, Jr., G. David. Minimal bases of rational vector spaces, with applications to multivariable linear systems. *SIAM J. Control*, 13(3):493–520, May 1975.
- von zur Gathen, J. and Gerhard, J. *Modern Computer Algebra*. Cambridge University Press, Cambridge, New York, Melbourne, 1999. ISBN 0-521-64176-4.
- Gerhard, Jürgen. Fast modular algorithms for squarefree factorization and Hermite integration. *Applic. Algebra Engin. Commun. Comput.*, 11(3):203–226, 2001.
- Giesbrecht, M. Fast computation of the Smith form of a sparse integer matrix. *Computational Complexity*, 10:41–69, 2001.
- Giesbrecht, Mark and Storjohann, Arne. Computing rational forms of integer matrices. *J. Symbolic Comput.*, 34(3):157–172, 2002.
- Giorgi, Pascal, Jeannerod, Claude-Pierre, and Villard, Gilles. On the complexity of polynomial matrix computations. In Sendra [2003], pages 135–142. ISBN 1-58113-641-2.

- Heindel, L. E. and Horowitz, E. On decreasing the computing time for modular arithmetic. In *Conference Record, IEEE 12th Annual Symp. on Switching and Automata Theory*, pages 126–128, 1971.
- Huang, Xiaohan and Pan, Victor Y. Fast rectangular matrix multiplication and applications. *J. Complexity*, 14:257–299, 1998.
- Jeannerod, Claude-Pierre and Villard, Gilles. Straight-line computation of the polynomial matrix inverse. rapport de recherche 47, LIP (Laboratoire de l’Informatique du Parallélisme), ENSL (Ecole Normale Supérieure de Lyon), France, <http://www.ens-lyon.fr/LIP/Pub/rr2003.html>, 2002.
- Jenks, R. D., Sutor, R. S., and Watt, S. M. Scratchpad II: An abstract datatype system for mathematical computation. In Rice, J. R., editor, *Mathematical Aspects of Scientific Software*, volume 14 of *The IMA Volumes in Mathematics and its Application*, pages 157–182. Springer Verlag, New York, 1988.
- Kailath, T. *Linear systems*. Prentice Hall, 1980.
- Kaltofen, E. Greatest common divisors of polynomials given by straight-line programs. *J. ACM*, 35(1):231–264, 1988.
- Kaltofen, E. On computing determinants of matrices without divisions. In Wang, P. S., editor, *Proc. 1992 Internat. Symp. Symbolic Algebraic Comput. (ISSAC’92)*, pages 342–349, New York, N. Y., 1992. ACM Press.
- Kaltofen, E. Analysis of Coppersmith’s block Wiedemann algorithm for the parallel solution of sparse linear systems. *Math. Comput.*, 64(210):777–806, 1995.
- Kaltofen, E. Challenges of symbolic computation my favorite open problems. *J. Symbolic Comput.*, 29(6):891–919, 2000. With an additional open problem by R. M. Corless and D. J. Jeffrey.
- Kaltofen, Erich. An output-sensitive variant of the baby steps/giant steps determinant algorithm. In Mora [2002], pages 138–144. ISBN 1-58113-484-3.
- Kaltofen, Erich and Lee, Wen-shin. Early termination in sparse interpolation algorithms. *J. Symbolic Comput.*, 36(3–4):365–400, 2003. Special issue Internat. Symp. Symbolic Algebraic Comput. (ISSAC 2002). Guest editors: M. Giusti & L. M. Pardo.
- Kaltofen, E., Lee, W.-s., and Lobo, A. A. Early termination in Ben-Or/Tiwari sparse interpolation and a hybrid of Zippel’s algorithm. In Traverso, C., editor, *Proc. 2000 Internat. Symp. Symbolic Algebraic Comput. (ISSAC’00)*, pages 192–201, New York, N. Y., 2000. ACM Press. ISBN 1-58113-218-2.
- Kaltofen, Erich and May, John. On approximate irreducibility of polynomials in several variables. In Sendra [2003], pages 161–168. ISBN 1-58113-641-2.

- Kaltofen, E. and Saunders, B. D. On Wiedemann's method of solving sparse linear systems. In Mattson, H. F., Mora, T., and Rao, T. R. N., editors, *Proc. AAEECC-9*, volume 539 of *Lect. Notes Comput. Sci.*, pages 29–38, Heidelberg, Germany, 1991. Springer Verlag.
- Kaltofen, E. and Villard, G. On the complexity of computing determinants. In Shirayanagi, Kiyoshi and Yokoyama, Kazuhiro, editors, *Proc. Fifth Asian Symposium on Computer Mathematics (ASCM 2001)*, volume 9 of *Lecture Notes Series on Computing*, pages 13–27, Singapore, 2001. World Scientific. ISBN 981-02-4763-X. Invited contribution; extended abstract.
- Kaltofen, E. and Villard, G. Computing the sign or the value of the determinant of an integer matrix, a complexity survey. *J. Computational Applied Math.*, 2002. To appear, 17 pages. Special issue on Congrès International Algèbre Linéaire et Arithmétique: Calcul Numérique, Symbolique et Parallèle, held in Rabat, Morocco, May 2001.
- Knuth, D. E. The analysis of algorithms. In *Congrès int. Math., Nice, France*, volume 3, pages 269–274, 1970.
- Küchlin, W., editor. *ISSAC 97 Proc. 1997 Internat. Symp. Symbolic Algebraic Comput.*, New York, N. Y., 1997. ACM Press. ISBN 0-89791-875-4.
- Moenck, R. T. Fast computation of GCDs. In *Proc. 5th ACM Symp. Theory Comp.*, pages 142–151, 1973.
- Mora, T., editor. *ISSAC 2002 Proc. 2002 Internat. Symp. Symbolic Algebraic Comput.*, New York, N. Y., 2002. ACM Press. ISBN 1-58113-484-3.
- Mulders, T. and Storjohann, A. Certified dense linear system solving. Technical Report 355, Institute of Computer Systems, ETH Zurich, <http://www.inf.ethz.ch/research/publications/>, January 2000. Document posted on the Internet lists Nr. 356 and date December 2000.
- Mulders, T. and Storjohann, A. On lattice reduction for polynomial matrices. *J. Symbolic Comput.*, 35(4):377–401, 2003.
- Newman, M. *Integral Matrices*. Academic Press, 1972.
- Pan, Victor. Computing the determinant and the characteristic polynomial of a matrix via solving linear systems of equations. *Information Process. Lett.*, 28:71–75, June 1988.
- Pan, Victor Y. Randomized acceleration of fundamental matrix computations. In *Proc. STACS 2002*, volume 2285 of *Lect. Notes Comput. Sci.*, pages 215–226, Heidelberg, Germany, 2002. Springer Verlag.
- Paterson, M. S. and Stockmeyer, L. J. On the number of nonscalar multiplications necessary to evaluate polynomials. *SIAM J. Comp.*, 2:60–66, 1973.

- Popov, V. M. Some properties of control systems with irreducible matrix transfer functions. In *Lecture Notes in Mathematics*, volume 144, pages 169–180. Springer Verlag, Berlin, 1970.
- Rissanen, Jorma. Realizations of matrix sequences. Technical Report RJ-1032, IBM Research, Yorktown Heights, New York, 1972.
- Rosser, J. Barkley and Schoenfeld, Lowell. Approximate formulas of some functions of prime numbers. *Illinois J. Math.*, 6:64–94, 1962.
- Schönhage, A. Schnelle Berechnung von Kettenbruchentwicklungen. *Acta Inform.*, 1:139–144, 1971. In German.
- Schwartz, J. T. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27:701–717, 1980.
- Sendra, J.R., editor. *ISSAC 2003 Proc. 2003 Internat. Symp. Symbolic Algebraic Comput.*, New York, N. Y., 2003. ACM Press. ISBN 1-58113-641-2.
- Storjohann, Arne. Near optimal algorithms for computing Smith normal forms of integer matrices. In Lakshman Y. N., editor, *ISSAC 96 Proc. 1996 Internat. Symp. Symbolic Algebraic Comput.*, pages 267–274, New York, N. Y., 1996. ACM Press.
- Storjohann, Arne. *Algorithms for matrix canonical forms*. Dissertation, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, December 2000a.
- Storjohann, Arne. Computing the Frobenius form of a sparse integer matrix. Paper to be submitted, April 2000b.
- Storjohann, Arne. Higher-order lifting. In Mora [2002], pages 246–254. ISBN 1-58113-484-3.
- Storjohann, Arne. High-order lifting and integrality certification. *J. Symbolic Comput.*, 36(3-4):613–648, 2003. Special issue Internat. Symp. Symbolic Algebraic Comput. (ISSAC 2002). Guest editors: M. Giusti & L. M. Pardo.
- Strassen, V. Vermeidung von Divisionen. *J. reine u. angew. Math.*, 264:182–202, 1973. In German.
- Strassen, V. Algebraic complexity theory. In van Leeuwen, J., editor, *Handbook of Theoretical Computer Science, Algorithms and Complexity*, volume A, pages 633–672. Elsevier Science Publ., Amsterdam, 1990.
- Sugiyama, Y., Kasahara, M., Hirasawa, S., and Namekawa, T. A method for solving key equation for decoding Goppa codes. *Information & Control*, 27: 87–99, 1975.
- Thomé, E. Subquadratic computation of vector generating polynomials and improvements of the block Wiedemann method. *J. Symbolic Comput.*, 33(5): 757–775, May 2002.

- Turner, William J. A note on determinantal divisors and matrix preconditioners. Paper to be submitted, October 2001.
- Turner, William J. *Black box linear algebra with the LINBOX library*. PhD thesis, North Carolina State Univ., Raleigh, North Carolina, August 2002. 193 pages.
- Van Barel, M. and Bultheel, A. A general module theoretic framework for vector M-Padé and matrix rational interpolation. *Numerical Algorithms*, 3:451–462, 1992.
- Villard, G. *Calcul Formel et Parallélisme : Résolution de Systèmes Linéaires*. PhD thesis, Institut National Polytechnique de Grenoble, France, December 1988.
- Villard, G. Further analysis of Coppersmith’s block Wiedemann algorithm for the solution of sparse linear systems. In Küchlin [1997], pages 32–39. ISBN 0-89791-875-4.
- Villard, G. A study of Coppersmith’s block Wiedemann algorithm using matrix polynomials. Rapport de Recherche 975 IM, Institut d’Informatique et de Mathématiques Appliquées de Grenoble, www.imag.fr, April 1997b.
- Villard, Gilles. Computing the Frobenius normal form of a sparse matrix. In Ganzha, V. G., Mayr, E. W., and Vorozhtsov, E. V., editors, *CASC 2000 Proc. the Third International Workshop on Computer Algebra in Scientific Computing*, pages 395–407. Springer Verlag, 2000.
- Wiedemann, D. Solving sparse linear equations over finite fields. *IEEE Trans. Inf. Theory*, IT-32:54–62, 1986.
- Zippel, R. Probabilistic algorithms for sparse polynomials. In *Proc. EUROSAM ’79*, volume 72 of *Lect. Notes Comput. Sci.*, pages 216–226, Heidelberg, Germany, 1979. Springer Verlag.