ASYMPTOTICALLY FAST POLYNOMIAL MATRIX ALGORITHMS FOR MULTIVARIABLE SYSTEMS

Claude-Pierre Jeannerod and Gilles Villard

CNRS, INRIA, Laboratoire LIP, École normale supérieure de Lyon
46, Allée d'Italie, 69364 Lyon Cedex 07, France

Abstract

We present the asymptotically fastest known algorithms for some basic problems on univariate polynomial matrices: *rank, nullspace, determinant, generic inverse, reduced form* [8, 9, 16, 17]. We show that they essentially can be reduced to two computer algebra techniques, *minimal basis computations* and *matrix fraction expansion/reconstruction*, and to polynomial matrix multiplication. Such reductions eventually imply that all these problems can be solved in about the same amount of time as polynomial matrix multiplication. The algorithms are deterministic, or randomized with certified output in a Las Vegas fashion.

1 Introduction

We aim at drawing attention to today's asymptotically fastest known algorithms for computing with polynomial matrices. In particular, we shall focus on the following problems: compute the *rank*, a right or *left nullspace*, the *determinant*, the *inverse* and a column- or *row-reduced form* of a given polynomial matrix. Polynomial matrices are quite common in the analysis of multivariable linear systems and Kailath's treatise *Linear Systems* [10] is a good illustration of this.

Recently, algorithms have been designed [8, 9, 16, 17] that allow to compute solutions to these problems in essentially the same amount of time as for multiplying two polynomial matrices together. More precisely, given a field K—for example the complex numbers, the rationals or a finite field—and given a polynomial matrix $A \in \mathsf{K}[x]^{n\times n}$ whose entries have degree in x bounded by d, randomized algorithms allow to compute $\operatorname{rank} A$, $\ker A$, $\det A$ and to row-reduce A in $O^{\sim}(n^{\omega}d)$ operations in K. The inverse matrix A^{-1} when A is generic can be computed in $O^{\sim}(n^3 d)$ operations in K by a deterministic algorithm. The estimate $O^{\sim}(n^{\omega}d)$ is the best known asymptotic upper bound for multiplying two matrices in $\mathsf{K}[x]^{n\times n}$ of degree d [5, 3], where $2 \le \omega < 2.376$ is the exponent of matrix multiplication over K [4,

URLs: `http://perso.ens-lyon.fr/claude-pierre.jeannerod`
`http://perso.ens-lyon.fr/gilles.villard`

International Journal of Control Submission — 15/2/2006

Chapter 15]. Using schoolbook matrix multiplication, we have $\omega = 3$ and the bound $O^{\sim}(n^{\omega}d)$ becomes $O^{\sim}(n^3 d)$. Furthermore, the soft-O notation $O^{\sim}$ simply indicates some missing logarithmic factors of the form $\alpha(\log n)^{\beta}(\log d)^{\gamma}$ for three positive real numbers α, β, γ.

Throughout the paper if not specified the algorithms are deterministic. The randomized algorithms are of the Las Vegas kind—always correct, probably fast. We mean that in time $O^{\sim}(n^{\omega}d)$ they either return a correct result or return "failure", the latter with probability no more than, say, $1/2$; a correct result will be obtained after repetition. When we assume for inversion that the input matrix is *generic*, this means that the solution is computed in time $O^{\sim}(n^3 d)$ provided some minors of a suitable linearization of this matrix are non-zero. A precise definition of these minors can be found in [9].

By achieving the complexity estimate $O^{\sim}(n^{\omega}d)$, these algorithms improve upon all the complexity estimates that were known previously.

In this paper, evidence is given that the key tools for such improvements are:

- Minimal bases of $\mathsf{K}[x]$-modules;
- Expansion/reconstruction of polynomial matrix fractions.

The former has the same flavour as in [6] while for the fractions we heavily rely on the concepts in [10, Chapter 6]. Two kinds of minimal bases, namely *approximant bases* and *nullspace bases*, are studied in Section 2. There we will see that such bases are small enough to be computed fast, that is, in $O^{\sim}(n^{\omega}d)$ operations in K. Polynomial matrix fractions are matrices $F \in \mathsf{K}(x)^{n\times n}$, where $\mathsf{K}(x)$ is the field of rational functions over K. By expansion of F, we thus mean a power series expansion $F = \sum_{i=0}^{\infty} F_i x^i \in \mathsf{K}[[x]]^{n\times n}$, and by reconstruction (or factorization) of F we mean a left or right quotient of polynomial matrices like $F = A^{-1}B$ or $F = BA^{-1}$. It turns out that all we need is truncated expansions and reconstructed quotients that can be computed fast, as seen in Section 3. The key idea here is that a truncated expansion of sufficiently high order—with respect to the input problem—may lead to an exact solution over $\mathsf{K}[x]$. This is well-known in computer algebra, at least for scalar rational functions [7, §5.7] but, as far as we know, the extension to the matrix case is more recent [8, 9, 16, 17].

Minimal bases and matrix fractions are interesting not only because they can be computed fast, but also—and, perhaps, mainly—because computing a minimal basis and expanding/reconstructing a matrix fraction are problems to which we can reduce other problems like *rank*, *left nullspace*, *determinant*, *generic inverse* and *row-reduced form*. The goal of Section 4 is precisely to show this: there the above problems are thus seen as applications of the techniques studied in Sections 2 and 3.

If we assume given a $O^{\sim}(n^{\omega}d)$ algorithm for multiplying two n by n polynomial matrices of degree d, combining the reductions of Section 4 with the cost estimates of Sections 2 and 3 then yields $O^{\sim}(n^{\omega}d)$ solutions to all our problems under consideration. Of course, we could have introduced a cost function $\mathsf{MM}(n, d)$ for polynomial matrix multiplication and derived more precise complexity estimates for each of the problems, in terms of functions

of $\mathsf{MM}(n,d)$ (as done for instance in [8, Section 1] or [17, Section 1]). However, we prefer for this paper to stick to the more readable $O\tilde{}(n^\omega d)$ bound, which already gives a good sense of the link with polynomial matrix multiplication.

A first task remaining would be to relax the regularity assumptions made for *inversion* (the input should be generic and of dimensions a power of two, see Section 4.1) and for *row-reduction* (the input should be non-singular, see Section 4.3). But even these generic situations are enough for our purpose here of showing how to rely on minimal bases and matrix fraction expansions/reconstructions.

Also, recently, other problems on polynomial matrices than those treated in this paper have been shown to have about the same complexity as polynomial matrix multiplication. An example is the problem of computing the *Smith normal form* and thus also the determinant, whose solution in [16] gives us Theorem 3.1. However—and this is the second task remaining—, the list of problems that can be solved in about the same number of operations as for polynomial matrix multiplication still has to be augmented. The question is particularly interesting for the problem of computing the *characteristic polynomial* and the *Frobenius normal form*, for which the best known solutions [11, 12] have cost $O\tilde{}(n^{2.7}d)$ still greater than $O\tilde{}(n^\omega d)$.

Notation and basic reminders. Here and hereafter log denotes the logarithm in base two and I_n the n by n identity matrix. For a matrix A over $\mathsf{K}[x]$, we denote its value at $x=0$ by $A(0)$. For $d \in \mathbb{N}$ and a matrix F over $\mathsf{K}[[x]]$, $F \equiv 0 \mod x^d$ means that each entry of F is a multiple of x^d, and $F \mod x^d$ means that we truncate F into a polynomial matrix of degree less than d. By *size* of a polynomial matrix over $\mathsf{K}[x]$ we mean the number of elements of K that are necessary to represent it. For example, $M \in \mathsf{K}[x]^{n\times p}$ of degree d has size at most $np(d+1) = O(npd)$. A polynomial matrix is said to be *non-singular* when it is square and when its determinant is a non identically zero polynomial. Two matrices $A, R \in \mathsf{K}[x]^{n\times n}$ are *unimodularly left equivalent* when there exists $U \in \mathsf{K}[x]^{n\times n}$ such that $\det U$ is a non-zero constant—that is, U is *unimodular*— and when $UA = R$.

2 Minimal approximant bases and minimal nullspace bases

Our solutions for solving a class of polynomial matrix problems in about the same number of operations in K as for multiplying two polynomial matrices will fundamentally rely on computing minimal bases of $\mathsf{K}[x]$-modules. The target complexity estimate $O\tilde{}(n^\omega d)$ is reached since the bases we use are small, with size $O(n^2 d)$ is most cases, and can be computed fast (see Theorem 2.2 below).

Definition 2.1 *Let $\mathcal{M}$ be a $\mathsf{K}[x]$-submodule of $\mathsf{K}[x]^n$ of dimension $\mathcal{D}$. A basis $N_1, \ldots, N_{\mathcal{D}} \in \mathsf{K}[x]^n$ of $\mathcal{M}$ with degrees $\delta_1 \leq \cdots \leq \delta_{\mathcal{D}}$ is called a* minimal basis *if any other basis of $\mathcal{M}$ with degrees $d_1 \leq \cdots \leq d_{\mathcal{D}}$ satisfies $d_i \geq \delta_i$ for $1 \leq i \leq \mathcal{D}$. The degrees δ_i are called the* minimal indices *of $\mathcal{M}$.*

In applications to multivariable systems, this definition follows the study of minimal polynomial bases of vector spaces in [6]. The two important examples of such bases that

we use in this paper are minimal approximant bases and minimal nullspace bases. Approximant bases are defined for a power series matrix over $\mathsf{K}[[x]]$ whereas nullspace bases are computed as special approximant bases for a polynomial matrix over $\mathsf{K}[x]$.

2.1 Minimal approximant bases

Given a formal power series $F \in \mathsf{K}[[x]]^{n\times p}$ and $d \in \mathbb{N}$, we take for $\mathcal{M}$ the set of all approximants of order d to F:

$$\mathcal{M} = \{v \in \mathsf{K}[x]^{1\times n} \ : \ vF \equiv 0 \mod x^d\}.$$

The minimal bases of $\mathcal{M}$ are called *minimal approximant bases for F of order d*. Since $\mathcal{M}$ has dimension n, such bases form non-singular $n \times n$ polynomial matrices. These polynomial matrices further have degree less than d and size $O(n^2 d)$.

Theorem 2.2 [8]. *Let $F \in \mathsf{K}[[x]]^{n\times p}$ with $p = O(n)$, and $d \in \mathbb{N}$. A minimal approximant basis for F of order d can be computed in $O\tilde{}(n^\omega d)$ operations in K.*

Our notion of minimal approximant bases is directly inspired by [1] with some adaptations for fully reflecting the polynomial matrix point of view. The cost estimate of Theorem 2.2 is a matrix polynomial generalization of the recursive Knuth/Schönhage half-gcd algorithm for scalar polynomials [13, 15] (see also [7, §11.1]), that takes into account fast polynomial matrix multiplication.

For a matrix A over $\mathsf{K}[x]$, we denote by d_i its ith *row degree*, that is, the highest degree of all the entries of the ith row of A. The *leading (row) coefficient matrix* of A is the constant matrix whose ith row consists of the coefficients of x^{d_i} in the ith row of A. We recall from [10, §6.3.2] that a full row rank polynomial matrix is *row-reduced* when its leading (row) coefficient matrix also has full rank. As a consequence of their minimality, minimal approximant bases have the following properties, which will be used in Section 2.2 when specializing approximants for power series matrices to approximants for polynomial matrices.

Property 2.3 *Let N be a minimal approximant basis for F of order d. Then,*

I. *N is row-reduced;*

II. *If $v \in \mathcal{M}$ has degree δ with $\delta < d$, then there is a unique $u \in \mathsf{K}[x]^{1\times n}$ such that $v = uN$, and N has at least one row of degree no more than δ.*

Property 2.3.I is a consequence of minimality [10, Theorem 6.5-10]. Property 2.3.II is the fact that the rows of N form a basis, together with the predictable-degree property [10, Theorem 6.3-13].

2.2 Minimal nullspace bases

Given a polynomial matrix $A \in \mathsf{K}[x]^{n\times p}$ of rank r, we now take

$$\mathcal{M} = \{v \in \mathsf{K}[x]^{1\times n} \ : \ vA = 0\}.$$

This is a $\mathsf{K}[x]$-submodule of $\mathsf{K}[x]^n$ of dimension $n-r$. Its minimal bases are called *minimal nullspace bases for* A and form full rank $(n-r)\times n$ polynomial matrices. The minimal indices $\delta_1 \leq \cdots \leq \delta_{n-r}$ (see Definition 2.1) are called the (left) *Kronecker indices* of A [10, §6.5.4]. For any given threshold degree δ, we further define

$$\kappa = \max\{1 \leq i \leq n-r \ : \ \delta_i \leq \delta\}. \tag{1}$$

A corresponding family of κ linearly independent vectors of degrees $\delta_1, \cdots, \delta_\kappa$ is a family of *minimal nullspace vectors* of degree at most δ. The theorem below says that if $F = A$ is a polynomial matrix then any minimal approximant basis for A of sufficiently high order actually contains a family of minimal nullspace vectors for A.

Theorem 2.4 *Let $A \in \mathsf{K}[x]^{n\times p}$ be of degree d. Let N be a minimal approximant basis for A of order $\delta + d + 1$. Then exactly κ rows of N have degree at most δ; these rows are in the left nullspace of A and their degrees are the Kronecker indices $\delta_1, \ldots, \delta_\kappa$.*

Proof. Let $N_{i,*}$ be the ith row of N and let d_i be its degree. With no loss of generality we may assume that $d_1 \leq d_2 \leq \cdots \leq d_n$. First, we show by induction that $N_{i,*} \in \ker A$ and $d_i = \delta_i$ for $1 \leq i \leq \kappa$. (Here and hereafter $\ker A$ means the *left* nullspace of A.) For $i = 1$, let $v_1 \in \ker A$ of degree δ_1. Since v_1 is an approximant for A of *any* order, it follows from Property 2.3.II and from the ordering on the d_i's that $d_1 \leq \delta_1$. Therefore $\deg(N_{1,*}A) \leq d_1 + d < \delta + d + 1$, and $N_{1,*}A \equiv 0 \mod x^{\delta+d+1}$ yields $N_{1,*}A = 0$. Since by definition of δ_1 all polynomial vectors in $\ker A$ have degree at least δ_1, we get $d_1 \geq \delta_1$. We thus have shown that $N_{1,*} \in \ker A$ and $d_1 = \delta_1$. Now, for $i \in \{2, \ldots, \kappa\}$, assume that the first $i-1$ rows of N are in $\ker A$ and have respective degrees $\delta_1, \ldots, \delta_{i-1}$. Since $i-1 < \dim \ker A$, one can take $v_i \in \ker A$ of degree δ_i and linearly independent of these rows. Let $u_i = [u_{i1}, \ldots, u_{in}]$ be the vector such that $v_i = u_i N$, given by Property 2.3.II. There exists $j_0 \in \{i, \ldots, n\}$ such that $u_{ij_0} \neq 0$, for otherwise v_i would be a linear combination of the first $i-1$ rows of N. By Property 2.3.I and the predictable-degree property [10, Theorem 6.3-13], we have $\delta_i = \max\{\deg u_{ij} + d_j \ : \ u_{ij} \neq 0\}$, and therefore $\delta_i \geq d_{j_0} \geq d_i$. Here, $d_i \leq \delta_i$ yields $N_{i,*} \in \ker A$ in the same way as for $i = 1$. Now, N being non-singular, $N_{i,*}$ is in fact a nullspace vector that is linearly independent of the nullspace vectors $N_{1,*}, \ldots, N_{i-1,*}$ whose degrees are the first $i-1$ successive minimal indices $\delta_1, \ldots, \delta_{i-1}$. Using the minimality of δ_i, this implies $d_i \geq \delta_i$. By induction, we thus have shown that the first κ rows of N are nullspace vectors of degrees $\delta_1, \ldots, \delta_\kappa$; these degrees are at most δ by definition of κ. To conclude, let us check that N must have at most κ rows of degree at most δ. Otherwise $d_{\kappa+1} \leq \delta$ due to the ordering on the d_i's. But $d_{\kappa+1} \leq \delta$ implies, as before, $N_{\kappa+1} \in \ker A$. When $\kappa = n - r$, this contradicts $\dim \ker A = \kappa$, for the first $\kappa + 1$ rows of N are linearly independent. When $\kappa < n - r$, this implies $\delta_{\kappa+1} \leq d_{\kappa+1}$ in the same way as for the induction above; hence $\delta_{\kappa+1} \leq \delta$, which contradicts the definition of κ. ∎

For some applications, a *shifted degree* may be introduced (see [2] and the references therein), and some aspects of Theorem 2.4 may be generalized accordingly (see [2, Theorem 4.2] or [17, Lemma 6.3]).

Notice that if the Kronecker indices of A are all bounded by d then an *entire* minimal nullspace basis for A can already be computed fast: by Theorem 2.4, it suffices to compute a minimal approximant basis for A of order $2d+1$ and, by Theorem 2.2, this computation can be done in time $O^{\sim}(n^{\omega}d)$ when $p = O(n)$.

However, in the general case the Kronecker indices may be unbalanced, and range between 0 and nd (see [17, Theorem 3.3]). Computing a nullspace basis fast is then far less immediate. The method we shall give in Section 4.4 relies on the complexity result given in Theorem 2.5 below.

Theorem 2.5 [17]. *For $m \leq n$, let $A \in \mathsf{K}[x]^{(n+m)\times n}$ be of full column rank and degree at most d. If $\delta \in \mathbb{N}$ satisfies*

$$\delta m = O(nd), \tag{2}$$

then a family of minimal nullspace vectors of degree at most δ can be computed by a randomized Las Vegas (certified) algorithm in $O^{\sim}(n^{\omega}d)$ operations in K.

Random values are introduced through a random compression matrix $P \in \mathsf{K}[x]^{n\times m}$ that allows to compute minimal vectors more efficiently using the matrix $AP \in \mathsf{K}[x]^{(n+m)\times m}$ rather than directly from $A \in \mathsf{K}[x]^{(n+m)\times n}$ (see [17, Proposition 5.4]). The compression is essential to the complexity estimation when δ is large (Section 4.4 heavily relies on this). Indeed, a direct combination of Theorems 2.2 and 2.4 only leads to $O^{\sim}(n^{\omega}(d+\delta))$, which for $\delta = nd$ is $O^{\sim}(n^{\omega+1}d) \gg O^{\sim}(n^{\omega}d)$. The cost estimate $O^{\sim}(n^{\omega}d)$ relies on the compromise (2) between the minimal nullspace vector degree bound δ and the row dimension of matrix A. For example, when $m = 1$ one can compute a nullspace vector of degree as large as $O(nd)$, whereas when $m = n$ one may compute up to n nullspace vectors of degree $O(d)$.

3 Matrix fraction expansion and reconstruction

Matrix fraction expansion and *reconstruction* will be key tools especially for the row reduction and the nullspace problems. Fraction reconstruction is a useful tool in computer algebra (*e.g.* see [7, §5.7] for scalar polynomials), that is directly connected to *coprime factorization* (see below, and [10, Chapter 6] or [14] and the references therein).

For a polynomial matrix A that is non-singular at $x = 0$ and a polynomial matrix B, the techniques of [16, Proposition 17] reduce the computation of parts of the power series expansion

$$A^{-1}B = \sum_{i=0}^{\infty} F_i x^i$$

to polynomial matrix multiplication. By parts of the expansion, we mean a given number of consecutive matrix coefficients F_i. This is summarized in the following theorem.

Theorem 3.1 [16]. *For $m \leq n$, let $A \in \mathsf{K}[x]^{n\times n}$ non-singular and $B \in \mathsf{K}[x]^{n\times m}$, both of degree at most d. Let further $h \in \mathbb{N}$ be such that $h = O(nd)$. If $\det A(0) \neq 0$ and if $\delta \in \mathbb{N}$ satisfies*

$$\delta m = O(nd), \tag{3}$$

then the δ coefficients $F_h, F_{h+1}, \ldots, F_{h+\delta-1} \in \mathsf{K}^{n\times m}$ of the expansion of $A^{-1}B$ at $x = 0$ can be computed in $O\tilde{}(n^\omega d)$ operations in K.

Similarly to Theorem 2.5, the cost estimate $O\tilde{}(n^\omega d)$ relies on the compromise (3) between approximation order δ and the column dimension of matrix B. For instance, for a vector $B = b \in \mathsf{K}[x]^{n\times 1}$ and $h = 0$, one can expand $A^{-1}b$ up to order $O(nd)$, whereas with $B = I_n$ and $h = 0$, one gets the expansion of A^{-1} up to order $O(d)$. In Section 4.3, we shall use this result with $B = I_n$ and $h = (n-1)d+1$ in order to get a high-order slice of length $O(d)$ of the expansion of A^{-1}.

Remark 3.2 *Notice also that the regularity assumption $\det A(0) \neq 0$ in Theorem 3.1 is not restrictive provided that the polynomial matrix A is non-singular. In subsequent sections Theorem 3.1 is applied for algorithms with polynomial matrices in input and output. The regularity can be ensured with high probability using random shifts, thus yielding randomized algorithms for any $A(0)$. Typically, with a randomly chosen $x_0 \in \mathsf{K}$, we shift x in the input like $x \leftarrow x + x_0$ to get a regular input polynomial matrix at zero and, at the end of the computation, we shift x back like $x \leftarrow x - x_0$ to recover the output polynomial matrix (see [16, 8, 17]).*

A rational matrix $H \in \mathsf{K}(x)^{n\times m}$ is *strictly proper* if $\lim_{x\to\infty} H(x) = 0 \in \mathsf{K}^{n\times m}$. In most applications, difficulties arise when $A^{-1} \in \mathsf{K}(x)^{n\times n}$ is *not* strictly proper. However, one can define another fraction that is always strictly proper and shares some invariants with A^{-1}. Before seeing this, let us recall some useful facts about *greatest common divisors* of two polynomial matrices.

Definition 3.3 *Let A and B be polynomial matrices with the same number of rows. A (left)* matrix gcd *of A and B is any polynomial matrix G such that G has full column rank and $[A \;\; B]U = [G \;\; 0]$ for some unimodular polynomial matrix U.*

Definition 3.3 is for instance from [10, Lemma 6.3-3]. If $[A \;\; B]$ has full row rank then all the (left) gcd's of A and B are non-singular and unimodularly (right) equivalent (see [10, Lemma 6.3-4]). A non-singular $A \in \mathsf{K}[x]^{n\times n}$ is said to be (left) *coprime* with $B \in \mathsf{K}[x]^{n\times m}$ if any (left) gcd of A and B is unimodular; in this case, *the* (left) gcd may be chosen as being the identity matrix I_n. Similar definitions hold for right gcd's and right coprimeness.

Theorem 3.4 [8]. *Let $A \in \mathsf{K}[x]^{n\times n}$ of degree d, with $\det A(0) \neq 0$. For $A^{-1} = \sum_{i=0}^{\infty} F_i x^i$ and $h > (n-1)d$, let $H \in \mathsf{K}(x)^{n\times n}$ be given by $H = \sum_{i=0}^{\infty} F_{h+i} x^i$. Then $H = A^{-1}(AH) = (HA)A^{-1}$ is strictly proper, and AH and HA are polynomial matrices that are respectively left and right coprime with A.*

Proof. Let $B = AH$. By definition of H, we have $I_n = A(A^{-1} \mod x^h) + x^h B$. (In [16], this is essentially (17) with B and T respectively set to I_n and A.) Together with the fact that A and H have only non-negative powers of x, this implies that B is a polynomial matrix. Let us now verify that the fraction $H = A^{-1}B$ is strictly proper. Writing A^* for the adjoint of A, we have $A^{-1} = (\det A)^{-1}A^*$, where $\deg(\det A) \geq 0$ and $\deg A^* \leq (n-1)d < h$. Applying these degree bounds to $H = x^{-h}A^{-1} - x^{-h}(A^{-1} \mod x^h)$ yields $\lim_{x\to\infty} H(x) = 0$. For left coprimeness, notice that

$$\begin{bmatrix} A & x^h B \end{bmatrix} \underbrace{\begin{bmatrix} I_n & (A^{-1} \mod x^h) \\ 0 & I_n \end{bmatrix} \begin{bmatrix} 0 & I_n \\ I_n & -A \end{bmatrix}}_{U} = \begin{bmatrix} I_n & 0 \end{bmatrix}.$$

The matrix U above is unimodular and I_n is therefore a left gcd of A and $x^h B$. Now let G be a left gcd of A and B: by definition, it is a left divisor of A and B, and therefore also of A, $x^h B$ and their left gcd I_n. Hence $GP = I_n$ for some polynomial matrix P. This means that G is unimodular and, eventually, that A and B are left coprime. Taking $B = HA$, one can show in the same way that HA is over $\mathsf{K}[x]$ and right coprime with A. ∎

For our application in Section 4.3, we will need only the first, say δ, coefficients of the expansion of H as in Theorem 3.4. These coefficients thus correspond to a slice of order h and length δ of the expansion of A^{-1} and, to recover them, we shall use Theorem 3.1 with $B = I_n$.

Matrix power series expansion will be used in conjunction with matrix *(irreducible) fraction reconstruction* or, equivalently, *(coprime) factorization*. We show below that minimal approximant bases are appropriate tools for solving these problems.

Definition 3.5 *A (left)* factorization of degree δ *of a rational matrix* $H \in \mathsf{K}(x)^{n\times n}$ *is a representation* $H = V^{-1}U$ *with* U *and* V *two polynomial matrices of degree at most* δ*. This factorization is said to be* coprime *when* U *and* V *are (left) coprime.*

A similar definition holds for right factorizations. Hence, given $H \in \mathsf{K}(x)^{n\times n}$ (part of the expansion of H in practice), the reconstruction or factorization problem is to recover two n by n matrices U and V over $\mathsf{K}[x]$ such that $V^{-1}U = H$. If H is defined at $x = 0$ and given by its formal expansion $F \in \mathsf{K}[[x]]^{n\times n}$, this problem reduces to computing a suitable $[U \;\; V] \in \mathsf{K}[x]^{n\times 2n}$ such that

$$\begin{bmatrix} U & V \end{bmatrix} \begin{bmatrix} -I_n \\ F \end{bmatrix} = 0.$$

Theorem 3.6 *Let* $H \in \mathsf{K}(x)^{n\times n}$ *be strictly proper, with expansion* $F \in \mathsf{K}[[x]]^{n\times n}$ *at* $x = 0$*. Assume that* H *admits a right factorization of degree* δ_R *and a left factorization of degree* δ_L*. Let* $N \in \mathsf{K}[x]^{2n\times 2n}$ *be a minimal approximant basis for* $[-I_n \;\; F^T]^T$ *of order* $\delta_L+\delta_R+1$*. Then exactly* n *rows of* N *have degree at most* δ_L*; these rows form a matrix* $[U \;\; V] \in \mathsf{K}[x]^{n\times 2n}$ *such that* $V^{-1}U$ *is a left coprime factorization of* H*, with* V *row-reduced.*

Proof. Let BA^{-1} be a right factorization of H of degree δ_R, and let $T^{-1}S$ be a left factorization of H of degree δ_L. Let $M = [-I_n \quad F^T]^T A = [-A^T \quad B^T]^T$. Let us first see that N is also a minimal approximant basis for M by verifying that $[-I_n \quad F^T]^T$ and M have the same approximants $v \in \mathsf{K}[x]^{1\times n}$. If $v[-I_n \quad F^T]^T \equiv 0 \mod x^k$ for some $k \geq 0$ then obviously $vM = v[-I_n \quad F^T]^T A \equiv 0 \mod x^k$. Conversely, $v[-I_n \quad F^T]^T A \equiv 0 \mod x^k$ implies $v[-I_n \quad F^T]^T = x^k w A(0)^{-1}$ for a $w \in \mathsf{K}[[x]]^{1\times n}$ hence $v[-I_n \quad F^T]^T \equiv 0 \mod x^k$. Furthermore, M has rank n and $[S \quad T]$ is a basis of its left nullspace. Now let $\delta = \delta_L$ and $d = \delta_R$. Since by assumption both S and T have degree at most δ, the n left Kronecker indices of M are at most δ. Hence κ in (1) satisfies $\kappa = n$. Therefore, when applied to M, Theorem 2.4 says that N has exactly n rows of degree at most δ, and that these rows are in the left nullspace of M. Denoting by $[U \quad V]$ the submatrix of N corresponding to these rows, we thus have $U, V \in \mathsf{K}[x]^{n\times n}$ of degree at most δ such that $UA = VB$. Let us now show that $V^{-1}U$ is a left coprime factorization of H. First, V must be non-singular; for otherwise, using $\det A \neq 0$ and $UA = VB$, there exists a non-zero vector w such that $wU = wV = 0$, which contradicts the fact that $[U \quad V]$ has full rank. Hence the left factorization $H = V^{-1}U$. Since by Theorem 2.4 the row degrees of $[U \quad V]$ are exactly the Kronecker indices of M, the nullspace basis $[U \quad V]$ is in fact minimal and thus, by [10, Theorem 6.5-10], irreducible. Now, irreducibility implies that U and V are left coprime. Otherwise there exists a polynomial matrix G such that $U = GU'$, $V = GV'$ and $\deg(\det G) > 0$. At the finite zeroes of $\det G$ the rank of $[U \quad V]$ is then not full anymore, which contradicts irreducibility. We have thus shown that $V^{-1}U$ is a left coprime factorization of H and it remains to prove that its denominator V is row-reduced. This last point is a consequence of the fact that $[U \quad V]$ is itself row-reduced, as a submatrix of the row-reduced matrix N (see Property 2.3.i). Indeed, since by assumption $H = V^{-1}U$ is strictly proper, the row degrees of U are strictly smaller than those of V [10, Lemma 6.3-10], and the leading (row) coefficient matrix of $[U \quad V]$ has the form $[0 \quad L]$ where L is the leading (row) coefficient matrix of V, which is then non-singular. ∎

As an immediate consequence of Theorem 3.6 and Theorem 2.2, coprime factorizations can be computed fast when the input matrix fractions admit left and right factorizations of degree $O(d)$. This corollary, given below, will be applied in Section 4.3 to the particular matrix fraction H of Theorem 3.4.

Corollary 3.7 *Let $H \in \mathsf{K}(x)^{n\times n}$ be as in Theorem 3.6 with $\delta_L = O(d)$ and $\delta_R = O(d)$. Given the first $\delta_L + \delta_R + 1$ coefficients of the expansion of H at $x = 0$, one can compute a left coprime factorization of H in $O\tilde{}(n^\omega d)$ operations in K.*

4 Applications

In this section, we show how the techniques presented in Sections 2 and 3 can be used to solve the following problems asymptotically fast:

- $\mathsf{Inv}_{n,d}$: given a non-singular $A \in \mathsf{K}[x]^{n\times n}$ of degree d, compute A^{-1}.

- $\mathsf{Det}_{n,d}$: given $A \in \mathsf{K}[x]^{n\times n}$ of degree d, compute $\det A$.
- $\mathsf{RowRed}_{n,d}$: given $A \in \mathsf{K}[x]^{n\times n}$ of degree d, compute a row-reduced form of A.
- $\mathsf{NullSpace}_{n,d}$: given $A \in \mathsf{K}[x]^{n\times n}$ of degree d, compute the rank r of A and a full rank $N \in \mathsf{K}[x]^{(n-r)\times n}$ such that $NA = 0$.
- $\mathsf{Factor}_{n,d}$: given a right factorization of degree d of $H \in \mathsf{K}(x)^{n\times n}$, compute a left factorization of H.

Our approach here is to reduce each of the above five problems to (collections of) the problems below, for which $O\tilde{}(n^\omega d)$ solutions are known:

- $\mathsf{MatMul}_{n,d}$: given $A, B \in \mathsf{K}[x]^{n\times n}$ of degree at most d, compute the product AB.

  $\hookrightarrow$ *For solutions in time $O\tilde{}(n^\omega d)$ see [5], [3].*
- $\mathsf{PartialNullSpace}_{m,\delta}$: given $\delta = O(nd/m)$ with n, d fixed, and given $A \in \mathsf{K}[x]^{(n+m)\times n}$ of degree d and with $m \le n$, compute a family of minimal nullspace vectors of A of degree at most δ.

  $\hookrightarrow$ *Solved in time $O\tilde{}(n^\omega d)$ by Theorem 2.5 (randomized Las Vegas).*
- $\mathsf{MatFracExp}_{m,\delta}$: given $\delta = O(nd/m)$ with n, d, h fixed such that $h = O(nd)$, and given $A \in \mathsf{K}[x]^{n\times n}, B \in \mathsf{K}[x]^{n\times m}$ of degree at most d, with $m \le n$ and $A(0)$ non-singular, compute the δ coefficients $F_h, F_{h+1}, \ldots, F_{h+\delta-1}$ of the expansion of $A^{-1}B$ at $x = 0$.

  $\hookrightarrow$ *Solved in time $O\tilde{}(n^\omega d)$ by Theorem 3.1.*
- $\mathsf{MatFracRec}_{n,d}$: given $\delta_L, \delta_R = O(d)$ and the first $\delta_L + \delta_R + 1$ coefficients of the expansion at $x = 0$ of $H \in \mathsf{K}(x)^{n\times n}$ as in Theorem 3.6, compute a left coprime factorization of H with row-reduced denominator.

  $\hookrightarrow$ *Solved in time $O\tilde{}(n^\omega d)$ by Corollary 3.7.*

Assuming that n is a power of two and given a problem $\mathsf{P}_{n,d}$ or $\mathsf{P}_{m,\delta}$ such as any of those just introduced, we define the collections of problems we shall rely on as

$$\mathsf{P}^*_{n,d} := \left\{\text{solve } O(2^i) \text{ problems } \mathsf{P}_{n/2^i, 2^i d}\right\}_{0\le i<\log n}. \tag{4}$$

Such collections can be solved at about the same cost as polynomial matrix multiplication, as shown below. Here subscripts n, d and m, δ should be added to P and P^* depending on the underlying problem.

Lemma 4.1 *For all $\mathsf{P} \in \{\mathsf{MatMul}, \mathsf{PartialNullSpace}, \mathsf{MatFracExp}, \mathsf{MatFracRec}\}$, one can solve P^* in $O\tilde{}(n^\omega d)$ operations in K.*

Proof. This an immediate consequence of (4) and of the upper bound $O\tilde{}(n^\omega d)$ on the cost of each of these four problems. ∎

4.1 Polynomial matrix inversion ($\mathsf{Inv}_{n,d}$)

Given $A \in \mathsf{K}[x]^{n\times n}$ non-singular of degree d, the problem is to compute $A^{-1} \in \mathsf{K}(x)^{n\times n}$.

Assuming that A is generic and that n is a power of two, we recall from [9] how $\mathsf{Inv}_{n,d}$ reduces to $\mathsf{PartialNullSpace}^*_{n,d}$ plus some polynomial matrix multiplications. The algorithm in [9, p.75] essentially consists in computing in $\log n$ steps a non-singular matrix $U \in \mathsf{K}[x]^{n\times n}$ and a diagonal matrix $B \in \mathsf{K}[x]^{n\times n}$ such that

$$UA = B. \tag{5}$$

The inverse of A is then recovered as $A^{-1} = B^{-1}U$. The first step is as follows. Let $A = [A_L \ A_R]$ where $A_L, A_R \in \mathsf{K}[x]^{n\times(n/2)}$ and let $\underline{N}, \overline{N} \in \mathsf{K}[x]^{(n/2)\times n}$ be minimal nullspace bases for, respectively, A_L, A_R. This gives the first block-elimination step towards the diagonalization of A:

$$A = \begin{bmatrix} A_L & A_R \end{bmatrix} \quad \rightarrow \quad NA = \begin{bmatrix} \overline{N} \\ \underline{N} \end{bmatrix} \begin{bmatrix} A_L & A_R \end{bmatrix} = \begin{bmatrix} \overline{N}A_L & \\ & \underline{N}A_R \end{bmatrix}. \tag{6}$$

When A is generic of degree d, it turns out that all the minimal indices of both $\underline{N}$ and $\overline{N}$ are equal to d [9, Fact 1] and that $\overline{N}A_L$ and $\underline{N}A_R$ are $n/2 \times n/2$ polynomial matrices of degree exactly $2d$ on which we iterate.

We show in [9] that the property "dimension $\times$ degree $= nd$" generically carries from one iteration to the other: at step i, starting from 2^{i-1} blocks of dimensions $(n/2^{i-1}) \times (n/2^{i-1})$ and degree $2^{i-1}d$, we compute 2^{i-1} pairs $(\underline{N}_i^{(j)}, \overline{N}_i^{(j)})$ of minimal nullspace bases of dimensions $(n/2^i) \times (n/2^{i-1})$ and whose minimal indices are all equal to $2^{i-1}d$; this amounts to solving instances of $\mathsf{PartialNullSpace}_{n/2^i,2^{i-1}d}$. Let $(U, B) = (I_n, A)$ before the first step. Step i also requires to update the matrix transform as $U \leftarrow \mathrm{diag}[N_i^{(j)}]_j \times U$ and the right hand side as $B \leftarrow \mathrm{diag}[N_i^{(j)}]_j \times B$. Because of the special block-structure of the polynomial matrices involved, it can be shown that these updates reduce to solving $O(2^{2i})$ problems $\mathsf{MatMul}_{n/2^{i-1},2^{i-1}d}$.

Overall, the $\log n$ block-diagonalization steps thus reduce to $\mathsf{PartialNullSpace}^*_{n,d}$ and to

$$\left\{\text{solve } O(2^{2i}) \text{ problems } \mathsf{MatMul}_{n/2^i,2^id}\right\}_{0\le i<\log n}. \tag{7}$$

By Lemma 4.1 and (7), we therefore obtain a solution to $\mathsf{Inv}_{n,d}$ in $O^{\tilde{}}(n^3d)$ operations in K.

Since by Cramer's rule each entry of A^{-1} has the form $p/(\det A)$ where $p \in \mathsf{K}[x]$ may have degree at large as $(n-1)d$, the size of A^{-1} is of the order of n^3d. The above inversion algorithm, defined for A generic and n a power of two, is therefore essentially optimal.

4.2 Determinant computation ($\mathsf{Det}_{n,d}$)

Given $A \in \mathsf{K}[x]^{n\times n}$ of degree d, the problem is to compute $\det A \in \mathsf{K}[x]$.

The *Smith normal form*, and hence the determinant, can be computed with $O^{\sim}(n^{\omega}d)$ operations in K by a randomized Las Vegas algorithm. We refer to [16] for a detailed presentation of the approach. The algorithm proceeds in $\log n$ steps. The dimension/degree compromise (3) is ensured at every step for reducing the problem especially to $\mathsf{MatFracExp}^*_{n,d}$.

We present here an alternative solution for generic matrices that is derived from Section 4.1 above. We assume that A is generic with n is a power of two. It has been shown in [8] that the diagonal entries $b_{i,i}$ of the diagonal matrix B in (5) are non-zero constant multiples of $\det A$. Since $\det A(0)$ is generically non-zero, we have

$$\det A = \frac{\det A(0)}{b_{i,i}(0)} b_{i,i} \quad \text{for all } 1 \leq i \leq n.$$

The problem $\mathsf{Det}_{n,d}$ thus reduces essentially to computing the determinant of the constant matrix $A(0)$ and to the computation of, say, $b_{1,1}$. It is well-known that over K computing the determinant reduces to matrix multiplication [4, Section 16.4] (that is, $\mathsf{Det}_{n,0}$ reduces to $\mathsf{MatMul}_{n,0}$ using our notations). Concerning $b_{1,1}$, we perform $\log n$ steps as for inversion but, since $b_{1,1}$ is the upper-left corner of B, we use instead of (6) the simpler step

$$A = \begin{bmatrix} A_L & A_R \end{bmatrix} \quad \rightarrow \quad \overline{N} A_L. \tag{8}$$

As in (6), $\overline{N}$ is a minimal nullspace basis for A_R. Step i now consists in computing a single minimal nullspace basis of dimensions $(n/2^i) \times (n/2^{i-1})$ and minimal indices $2^{i-1}d$, and then in multiplying this basis with the left half of an $n/2^{i-1}$ by $n/2^{i-1}$ block of degree $2^{i-1}d$, as in (8). Hence, computing $b_{1,1}$ by performing these $\log n$ steps reduces to solving $\mathsf{PartialNullSpace}^*_{n,d}$ and $\mathsf{MatMul}^*_{n,d}$. By Lemma 4.1, this gives a solution to $\mathsf{Det}_{n,d}$ in $O^{\sim}(n^{\omega}d)$ operations in K.

4.3 Row reduction ($\mathsf{RowRed}_{n,d}$)

Given $A \in \mathsf{K}[x]^{n \times n}$ of degree d, the problem is to compute $R \in \mathsf{K}[x]^{n \times n}$ that is row-reduced and unimodularly left equivalent to A.

Theorem 3.8 in [8] establishes a deterministic algorithm with $O^{\sim}(n^{\omega}d)$ operations in K when $A(0)$ is non-singular. Using Remark 3.2, the same complexity bound is valid for row-reducing non-singular matrices $A \in \mathsf{K}[x]^{n \times n}$ with a randomized Las Vegas algorithm.

We assume below that $A(0)$ is non-singular. Recall from [10, §6.3.2] that R is a row-reduced form of A when R is row-reduced and $R = UA$ for some unimodular polynomial matrix U. The solution in [8] works by expansion/reconstruction of the matrix fraction H as in Theorem 3.4 with $h = (n-1)d + 1$.

First, we expand H up to order $2d+1$. This is done by solving $\mathsf{MatFracExp}_{n,2d+1}$ once, taking $B = I_n$ and $h = (n-1)d + 1 = O(nd)$. From Theorem 3.4 we know that H is a strictly proper matrix fraction which admits left and right factorizations $A^{-1}(AH)$ and $(HA)A^{-1}$. Strict properness further implies that the degrees of both AH and HA must be less than the degree of A [10, Lemma 6.3-10], and are thus bounded by d as well.

Therefore, these left and right factorizations of H are factorizations of degree d and, using Theorem 3.6, we can reconstruct H from its expansion up to order $2d+1$ as $H = R^{-1}S$. This reconstruction corresponds to solving problem $\mathsf{MatFracRec}_{n,d}$ once. On one hand, we know by Theorem 3.6 that R is row-reduced. On the other hand, $A^{-1}(AH)$ and $R^{-1}S$ are left coprime factorizations of the same fraction, which implies that there exists a unimodular U such that $UA = R$ [10, Theorem 6.5-4]. It follows that R is indeed a row-reduced form of A. By Lemma 4.1, this reduction to $\mathsf{MatFracExp}_{n,2d+1}$ and $\mathsf{MatFracRec}_{n,d}$ gives a solution to $\mathsf{RowRed}_{n,d}$ in $O\tilde{}(n^\omega d)$ operations in K.

4.4 Small nullspace computation ($\mathsf{NullSpace}_{n,d}$)

Given $A \in \mathsf{K}[x]^{n\times n}$ of degree d, the problem is to compute the rank r of A as well as $N \in \mathsf{K}[x]^{(n-r)\times n}$ of rank $n-r$ such that $NA = 0$.

As already seen before Theorem 2.5, a solution in the restrictive (*e.g.* generic) case when all minimal vectors have degrees in $O(d)$ is provided by a solution to $\mathsf{PartialNullSpace}_{n,d}$. In the general case the row degrees in a nullspace basis of A may be unbalanced, they range between 0 and nd [17, Theorem 3.3]. Previously known methods, whose cost is essentially driven by the highest Kronecker index, do not seem to allow the target complexity estimate $O\tilde{}(n^\omega d)$ (see for instance [17, Section 2]).

Our approach in [17] gives a randomized Las Vegas algorithm with complexity bound $O\tilde{}(n^\omega d)$ for a general matrix A. Randomization is first used for reducing the general nullspace problem to the full column rank case. This consists in evaluating the rank r of A at a random point $x = x_0$, then in compressing A to a full column rank matrix $\hat{A}$. This strategy may possibly underestimate the rank and give $\hat{A}$ with $r_0 = \operatorname{rank}\hat{A} < \operatorname{rank} A = r$. However the certification of the rank may be accomplished as follows [17, Section 7.2]. We compute $n - r_0$ linearly independent nullspace vectors for $\hat{A}$, and test by multiplication whether these vectors are in the nullspace of A. A positive answer implies that $r \le r_0$, therefore is an adequate certificate.

We also derive a particular strategy for the case $n \gg r$. Consequently, for a simplified explanation here, we now assume that A has full column rank n and dimensions $(n+m)\times n$ with $m = O(n)$.

The algorithm works in $\log n$ steps. At step i, $1 \le i \le \log n$, we compute a set of about $m/2^i$ nullspace vectors of degrees less that $\delta = 2^i d$. These vectors are obtained from solutions to $\mathsf{PartialNullSpace}_{m/2^i,2^i d}$ for nullspace vectors of bounded degree $\delta = 2^i d$, and involving matrices of decreasing dimensions $n + m/2^i$. Hence we essentially have a reduction to $\mathsf{PartialNullSpace}^*_{m,d}$. The compromise (2) of Theorem 2.5 is satisfied since $(2^i d)(m/2^i) = md = O(nd)$. We may point out that the proof of Theorem 2.5 for the cost of the partial nullspace itself relies on solutions to $\mathsf{MatFracExp}_{m,\delta}$, and $\mathsf{MatFracRec}_{m,\delta}$. Nullspace vectors are computed using a matrix fraction expansion /reconstruction scheme. The appropriate instances for $\mathsf{PartialNullSpace}_{m/2^i,2^i d}$, $1 \le i \le \log n$, are built as submatrices of the input matrix A. Our choices for these submatrices ensure the linear independency of the successive computed sets of nullspace vectors. The algorithm hence outputs a union

of a logarithmic number of sets of linearly independent nullspace vectors. Each set, corresponding to an instance of $\mathsf{PartialNullSpace}_{m/2^i,2^id}$, is a family of minimal vectors for a submatrix of A. The minimality is not preserved in general with respect to A, however we prove that small degree vectors are obtained [17, Proposition 7.1].

This reduction of $\mathsf{NullSpace}_{n,d}$ to $\mathsf{PartialNullSpace}^*_{n,d}$ and to $\mathsf{MatMul}^*_{n,d}$ for additional matrix multiplications establishes that a solution matrix N such that $NA = 0$ can be computed in $O^{\sim}(n^\omega d)$ operations in K by a randomized Las Vegas (certified) algorithm.

4.5 Factorization ($\mathsf{Factor}_{n,d}$)

Given a right factorization BA^{-1} of degree d of $H \in \mathsf{K}(x)^{n\times n}$, the problem is to compute polynomial matrices U and V such that $V^{-1}U = H$.

Corollary 3.7, together with the expansion of $H = BA^{-1}$, provides a solution to $\mathsf{FracMatRec}_{n,d}$ if H admits factorizations of degree d on both sides. The solution of the general case, we mean for an arbitrary left side factorization, induces several difficulties for dealing with unbalanced row degrees. These difficulties are bypassed using the techniques of Section 4.4.

By considering the polynomial matrix $[-A^T \quad B^T]$ and solving $\mathsf{NullSpace}_{\mathsf{2n,d}}$ we get U and V such that

$$[U \quad V] \begin{bmatrix} -A \\ B \end{bmatrix} = 0.$$

Arguments similar to those used in the proof of Theorem 3.6 lead to the fact that V is non-singular. Hence a solution $V^{-1}U$ to the factorization problem is computed in $O^{\sim}(n^\omega d)$ operations in K. Note that since a solution to $\mathsf{NullSpace}_{\mathsf{2n,d}}$ may not be minimal, the factorization $V^{-1}U$ may not be coprime.

Acknowledgement

We thank the referee for her/his helpful comments.

References

[1] B. Beckermann and G. Labahn. A uniform approach for the fast computation of matrix-type Padé approximants. *SIAM J. Matrix Anal. Appl.*, 15(3):804–823, July 1994.

[2] B. Beckermann, G. Labahn, and G. Villard. Normal forms for general polynomial matrices. *Journal of Symbolic Computation*. To appear.

[3] A. Bostan and É. Schost. Polynomial evaluation and interpolation on special sets of points. *Journal of Complexity*, 21(4):420–446, 2005.

[4] B. Bürgisser, C. Clausen, and M.A. Shokrollahi. *Algebraic Complexity Theory*, volume 315 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, 1997.

[5] D.G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28(7):693–701, 1991.

[6] G.D. Forney. Minimal bases of rational vector spaces, with applications to multivariable linear systems. *SIAM J. Control*, 13:493–520, 1975.

[7] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, second edition, 2003.

[8] P. Giorgi, C.-P. Jeannerod, and G. Villard. On the complexity of polynomial matrix computations. In *Proc. International Symposium on Symbolic and Algebraic Computation, Philadelphia, Pennsylvania, USA*, pages 135–142. ACM Press, August 2003.

[9] C.-P. Jeannerod and G. Villard. Essentially optimal computation of the inverse of generic polynomial matrices. *Journal of Complexity*, 21(1):72–86, 2005.

[10] T. Kailath. *Linear Systems*. Prentice-Hall, 1980.

[11] E. Kaltofen. On computing determinants without divisions. In *International Symposium on Symbolic and Algebraic Computation, Berkeley, California USA*, pages 342–349. ACM Press, July 1992.

[12] E. Kaltofen and G. Villard. On the complexity of computing determinants. *Computational Complexity*, 13:91–130, 2004.

[13] D.E. Knuth. The analysis of algorithms. In *Proc. International Congress of Mathematicians, Nice, France*, volume 3, pages 269–274, 1970.

[14] C. Oară and A. Varga. Minimal degree coprime factorization of rational matrices. *SIAM J. Matrix Anal. Appl.*, 21:245–278, 1999.

[15] A. Schönhage. Schnelle Berechnung von Kettenbruchentwicklungen. *Acta Informatica*, 1:139–144, 1971.

[16] A. Storjohann. High-order lifting and integrality certification. *Journal of Symbolic Computation*, 36(3-4):613–648, 2003. Special issue International Symposium on Symbolic and Algebraic Computation (ISSAC'2002). Guest editors: M. Giusti & L. M. Pardo.

[17] A. Storjohann and G. Villard. Computing the rank and a small nullspace basis of a polynomial matrix. In *Proc. International Symposium on Symbolic and Algebraic Computation, Beijing, China*, pages 309–316. ACM Press, July 2005.