Computing the rank and a small nullspace basis of a polynomial matrix

Arne Storjohann, Gilles Villard

Computing the rank and a small nullspace basis of a polynomial matrix

Arne Storjohann (U. Waterloo)
and Gilles Villard

January 2005

CENTRE NATIONAL
DE LA RECHERCHE
SCIENTIFIQUE

INRIA

Computing the rank and a small nullspace basis of a polynomial matrix

Arne Storjohann (U. Waterloo)
and Gilles Villard

January 2005

Abstract

We reduce the problem of computing the rank and a nullspace basis of a univariate polynomial matrix to polynomial matrix multiplication. For an input $n \times n$ matrix of degree d over a field K we give a rank and nullspace algorithm using about the same number of operations as for multiplying two matrices of dimension n and degree d. If the latter multiplication is done in $\mathsf{MM}(n, d) = O\tilde{\ }(n^\omega d)$ operations, with ω the exponent of matrix multiplication over K, then the algorithm uses $O\tilde{\ }(\mathsf{MM}(n, d))$ operations in K. For $m \times n$ matrices of rank r and degree d, the cost expression is $O\tilde{\ }(nmr^{\omega-2}d)$. The soft-O notation $O\tilde{\ }$ indicates some missing logarithmic factors. The method is randomized with Las Vegas certification. We achieve our results in part through a combination of matrix Hensel high-order lifting and matrix minimal fraction reconstruction, and through the computation of minimal or small degree vectors in the nullspace seen as a $\mathsf{K}[x]$-module.

Keywords: Linear algebra, polynomial matrix, matrix multiplication, nullspace basis, minimal polynomial basis, randomized algorithm.

Résumé

Nous réduisons le problème du calcul du rang et du noyau d'une matrice polynomiale d'une variable, à la multiplication de deux matrices polynomiales. Pour une matrice $n \times n$ de degré d en entrée sur un corps K, nous donnons un algorithme de calcul du rang et du noyau dont le coût est grosso modo celui du produit de deux matrices de dimension n et de degré d. Si cette multiplication s'effectue en $\mathsf{MM}(n, d) = O\tilde{\ }(n^\omega d)$ opérations, où ω est l'exposant du produit de matrices sur K, alors l'algorithme nécessite $O\tilde{\ }(\mathsf{MM}(n, d))$ opérations dans K. Pour une matrice $m \times n$ de rang r et degré d, l'expression générale est $O\tilde{\ }(nmr^{\omega-2}d)$. La notation « soft-O » indique des facteurs logarithmiques éludés. La méthode est probabiliste certifiée sur le mode Las Vegas. Nous obtenons ces résultats en combinant une remontée de Hensel aux grands ordres avec une reconstruction de fraction matricielle, ainsi que grâce à un calcul de bases minimales (ou presque) dans le noyau vu comme $\mathsf{K}[x]$-module.

Mots-clés: Algèbre linéaire, matrice polynomiale, produit de matrices, noyau d'une matrice, base polynomial minimale, algorithme probabiliste.

COMPUTING THE RANK AND A SMALL NULLSPACE BASIS

OF A POLYNOMIAL MATRIX

Arne Storjohann[*] and Gilles Villard[†]

Abstract

We reduce the problem of computing the rank and a nullspace basis of a univariate polynomial matrix to polynomial matrix multiplication. For an input $n \times n$ matrix of degree d over a field K we give a rank and nullspace algorithm using about the same number of operations as for multiplying two matrices of dimension n and degree d. If the latter multiplication is done in $\mathsf{MM}(n, d) = O^{\tilde{}}(n^\omega d)$ operations, with ω the exponent of matrix multiplication over K, then the algorithm uses $O^{\tilde{}}(\mathsf{MM}(n, d))$ operations in K. For $m \times n$ matrices of rank r and degree d, the cost expression is $O^{\tilde{}}(nmr^{\omega-2}d)$. The soft-O notation $O^{\tilde{}}$ indicates some missing logarithmic factors. The method is randomized with Las Vegas certification. We achieve our results in part through a combination of matrix Hensel high-order lifting and matrix minimal fraction reconstruction, and through the computation of minimal or small degree vectors in the nullspace seen as a $\mathsf{K}[x]$-module.

1 Introduction

Two $n \times n$ univariate polynomial matrices over a field K, whose entries have degree d at most, can be multiplied in $\mathsf{MM}(n, d) = O^{\tilde{}}(n^\omega d)$ operations in K [7, 9] where ω is the exponent of matrix multiplication over K [8, Chapter 15]. For $M \in \mathsf{K}[x]^{n \times n}$ of degree d we propose an algorithm that uses about the same number of operations for computing the rank r of M, and $n - r$ linearly independent vectors N_i in $\mathsf{K}[x]^n$ such that $N_i M = 0$, $1 \le i \le n - r$. The cost of the algorithm is $O^{\tilde{}}(\mathsf{MM}(n, d)) = O^{\tilde{}}(n^\omega d)$ operations in K. If M is $m \times n$ of rank r, a more precise and rank-sensitive expression of the cost is $O^{\tilde{}}(nmr^{\omega-2}d)$ (see Theorem 7.4). The soft-O notation $O^{\tilde{}}$ indicates missing logarithmic factors $\alpha(\log n)^\beta (\log d)^\gamma$ for three positive real constants α, β, γ. We mention previous works on the subject in Section 2. Our main idea is to combine matrix lifting techniques [31, 32], minimal bases computation and matrix fraction reconstruction [1, 15, 16], together with a degree / dimension compromise for keeping the cost of the computation as low as possible. Within the target complexity, lifting used alone only allows to obtain few vectors of large degrees, while minimal bases used alone only leads to an incomplete set of vectors of small degrees.

Our study extends the knowledge of the interaction between matrix multiplication and other basic linear algebra problems on matrices over $\mathsf{K}[x]$. Indeed, the interaction is quite well known for linear algebra over a field. For instance we refer to the survey [8, Chapter 16] for a list of problems on matrices in $\mathsf{K}^{n \times n}$ that can be solved in $O(n^\omega)$ or $O^{\tilde{}}(n^\omega)$ operations

---

[*]School of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1 Canada
http://www.scg.uwaterloo.ca/~astorjoh
[†]CNRS, LIP, École Normale Supérieure de Lyon 46, Allée d'Italie, 69364 Lyon Cedex 07, France
http://perso.ens-lyon.fr/gilles.villard

in K. Only recent results give an analogous view (although incomplete) of the situation for polynomial matrices. It is known that the following problems can be solved with $O\tilde{}(\mathsf{MM}(n, d))$ operations: *linear system solution, determinant, order d approximants, Smith normal form,* and, for a non-singular matrix, *column reduction* [15, 31, 32]. It is possible to compute the *inverse* of a generic matrix in essentially optimal time $O\tilde{}(n^3 d)$ [16]. We may also consider the problem of computing the Frobenius normal form, thus in particular the characteristic polynomial, of a square matrix. It does not seem to be known how to calculate the form in time $O\tilde{}(\mathsf{MM}(n, d))$. The best known estimate $O\tilde{}(n^{2.7} d)$ is given in [21] (see also [18]) with $\omega = 2.376$ [11].

Hence, we augment the above list of problems solved in $O\tilde{}(\mathsf{MM}(n, d))$ with the *certified computation of the rank* and a *nullspace basis*. This improvement is made possible by combining in a new way the key ideas of [15, 16, 31]. For the rank, the target complexity $O\tilde{}(\mathsf{MM}(n, d))$ was only attainable by a Monte Carlo (non-certified) approach consisting in computing the rank of $M(x_0)$ for x_0 a random value in K (see Lemma 7.3). In obtaining a certified value of the rank and a nullspace basis within the target complexity, a difficulty is related to the output size. For $M \in \mathsf{K}[x]^{2n \times n}$ of degree d and rank n, Gaussian elimination (fraction free or using evaluation / interpolation) leads to a basis of n vectors of degrees nd in $\mathsf{K}[x]^{2n}$ in the worst-case, hence to an output size in $\Theta(n^3 d)$. A complexity in $O\tilde{}(n^\omega d)$ must therefore rely on a different strategy.

We propose a sort of elimination scheme based on *minimal polynomial bases*. A minimal basis of the nullspace as $\mathsf{K}[x]$-module is a basis with lowest possible degrees (all necessary definitions are given in Section 3). For $M \in \mathsf{K}[x]^{2n \times n}$ as above, the total size of a minimal basis of the nullspace is in $O(n^2 d)$ (see Theorem 3.3). However, it is not known how to reduce the problem of computing such a basis to that of polynomial matrix multiplication. In the same context, minimal bases have been already used for computing the inverse of a polynomial matrix in [16], but only the generic case has been solved. Indeed, for a generic $M \in \mathsf{K}[x]^{2n \times n}$, the degrees in a minimal basis of the nullspace are all equal to the input degree d, and somehow, a basis is easy to compute in $O\tilde{}(\mathsf{MM}(n, d))$ operations [16, Section 4]. In the general case, the vector degrees in a minimal basis may be unbalanced, they range between 0 to nd. Known methods whose cost is essentially driven by the highest degree do not seem to allow our objective.

Our solution presented in Section 7 is to slightly relax the problem, and to compute a small degree—rather than minimal—nullspace basis in a logarithmic number of steps. We rely on the fact that even in the unbalanced degree case, the sum of the degrees remains bounded by nd (Theorem 3.3). Intuitively, at step k for $1 \leq k \leq \log_2 n$, we compute about $n/2^k$ vectors of degrees less than $2^k d$. Algorithm Nullspace(M) in Section 7 (whole nullspace) calls at most $\log_2 n$ times Algorithm Nullspace minimal vectors(M, δ) of Section 6 (nullspace vectors of bounded degree δ) with increasing degree thresholds δ. To keep the cost as low as possible, the degree increase requires to reduce the dimensions of involved matrices in the same proportion. We refer to an analogous degree / dimension compromise in [32, Section 17] for computing the Smith normal form, and in [16, Section 2] for inversion.

For a general view of the process, including successive compressions of the problem into smaller problems for reducing dimensions, consider

$$M = \begin{bmatrix} A \\ B \end{bmatrix} \in \mathsf{K}[x]^{m \times n} \tag{1}$$

with A square and non-singular. The rows of the matrix $[BA^{-1} \quad -I_{m-n}]$ give a basis of the

nullspace of M. However, as noticed previously a direct calculation of BA^{-1} would be too expensive. Now, note that if $[BA^{-1} \quad -I_{m-n}] = S^{-1}N$, for S and N two appropriate polynomial matrices, then the rows of $S[BA^{-1} \quad -I_{m-n}] = N$ are also in the nullspace. A key observation, see Section 4, is that considering a polynomial matrix N instead of $[BA^{-1} \quad -I_{m-n}]$ takes advantage of minimal bases properties and allows us to manipulate smaller degrees.

In algorithm Nullspace we proceed the following way. We deal with a small number of submatrices of the initial input for reducing the problem to

$$M = \begin{bmatrix} A \\ B \end{bmatrix} \in \mathsf{K}[x]^{(n+p) \times n}, \ 1 \leq p \leq n, \tag{2}$$

and introduce appropriate "compressing" matrices $P \in \mathsf{K}[x]^{n \times p}$ (successive choices of p are guided by the compromise with the degree). We start with a matrix lifting / fraction reconstruction phase. We compute an expansion of $H = BA^{-1}$ in $\mathsf{K}[[x]]^{p \times n}$ using [31, 32] to sufficiently high order, and "compress" it to $H_p = BA^{-1}P \in \mathsf{K}[[x]]^{p \times p}$. A reconstruction phase [1, 15] (see also the comments about coprime factorization in Section 2) then gives

$$S^{-1}N_p = BA^{-1}P. \tag{3}$$

We prove that "good" choices of P imply that S—denominator matrix for $BA^{-1}P$—is also a denominator matrix for BA^{-1} (Proposition 4.2) and that vectors in the nullspace of M can be recovered (Proposition 5.4). Indeed, the computation of $S[H \quad -I_{m-n}] \bmod x^{\delta+1}$ gives row vectors in the nullspace of degrees bounded by δ (Proposition 6.4). For a candidate Monte Carlo value r_o for the rank, in $\log_2 n$ steps of compression / uncompression (and choices of δ and p) combined with matrix lifting / matrix fraction reconstruction, we are able to compute candidate vectors for a nullspace basis. A final multiplication certifies that the rank is correct (i.e., $r_o = r$) and that a nullspace has actually been computed.

Although for each degree threshold δ we compute a minimal polynomial basis, the compression strategy unfortunately does not lead to a minimal polynomial basis for the whole nullspace. However, we prove especially in Proposition 7.1 that vectors with reasonably small degrees are obtained.

Our algorithms are randomized of Las Vegas kind— always correct, probably fast. Randomization is essentially linked to the compression stages where the matrices P are chosen at random of degree d in $\mathsf{K}[x]^{n \times p}$. We also use random matrices Q over K for linear independence preconditioning [10], or random evaluation points x_0 in K. Our results are proven for symbolic points x_o and matrices P and Q. By evaluation [12, 35, 29], the same results hold with high probability for random x_0, P and Q if K has enough elements, see Remark 7.6. The cost estimates might increase by poly-logarithmic factors in the case of small fields (with the introduction of an algebraic extension). We skip the details here, and refer for instance to the techniques used in [10, 19, 20] and to the references therein.

We study the cost of the algorithms by bounding the number of field operations in K on an algebraic random access machine. In [15] and [32], *ad hoc* cost functions have been defined for matrix polynomial problems that can be reduced recursively to matrix polynomial multiplication:

$$\mathsf{MM}'(n, d) = \sum_{i=0}^{\log_2 d} 2^i \mathsf{MM}(n, 2^{-i}d)$$

and

$$\overline{\mathsf{MM}}(n,d) = \sum_{i=0}^{\log_2 n} 4^i \mathsf{MM}(2^{-i}n, d) + n^2 (\log n)\mathsf{B}(d)$$

where $\mathsf{B}(d)$ is the cost for solving the extended gcd problem for two polynomial in $\mathsf{K}[x]$ of degree bounded by d. If $\mathsf{M}(d)$ is the number of operations in K required for multiplying two polynomials in $\mathsf{K}[x]$ of degree d, the Knuth [23] / Schönhage [28] half-gcd algorithm allows $\mathsf{B}(d) = O(\mathsf{M}(d)\log d)$. For the scalar polynomial multiplication we take $\mathsf{M}(d) = O(d\log d\log\log d)$ [9]. The reader may refer to Chapters 8 and 11 in [14] for more details and references about polynomial multiplication and gcd computation.

For simplifying the cost results in this paper we consider either that

$$\mathsf{MM}(n,d) = O(n^\omega \mathsf{M}(d)) \tag{4}$$

using the algorithm of [9], or, when the field K has at least $2d+1$ elements [6, 7],

$$\mathsf{MM}(n,d) = O(n^\omega d + n^2 \mathsf{M}(d)). \tag{5}$$

Hence from (4) and (5) we assume that

$$\mathsf{MM}'(n,d) = O(\mathsf{MM}(n,d)\log d), \quad \overline{\mathsf{MM}}(n,d) = O((\mathsf{MM}(n,d) + n^2\mathsf{B}(d))\log n). \tag{6}$$

Note that if $\omega > 2$ then $\overline{\mathsf{MM}}(n,d) = O(\mathsf{MM}(n,d) + n^2\mathsf{B}(d)\log n)$. If the assumption (6) is not made then some of our cost results that use $\mathsf{MM}(n,d)$ are not valid. However, we state our algorithms in terms of polynomial matrix multiplication; precise complexity estimates in terms of the *ad hoc* cost functions could be derived with some extra care.

2   Previous works

The rank and a basis for the nullspace of a matrix $M \in \mathsf{K}[x]^{m \times n}$ of degree d and rank r may be computed by fraction free Gaussian elimination in $O\tilde{\ }(nmr^{\omega-1}d)$ operations in K [30, Chapter 2]. The same asymptotic estimate may also be obtained using evaluation / interpolation techniques such as Chinese remaindering [14, Section 5.5].

Therefore, compared to these classical approaches, we improve the cost by a factor n in the worst-case ($2n \times n$ full column-rank matrix).

An elimination strategy specific to polynomial matrices is given in [25] that improves— asymptotically in the dimensions—on $O\tilde{\ }(nmr^{\omega-1}d)$, and computes the rank by a deterministic algorithm in $O(nmrd^2)$ operations in K, but how to incorporate matrix multiplication, and generalize the approach to computing the nullspace, is not known.

An alternative to the "matrix over the polynomials" approach above is to linearize the problem. A first type of linearization is to consider a degree one matrix of larger dimension with the same structural invariants (see the definition of the Kronecker indices in Section 3) [4]. A degree one matrix is a matrix pencil and an important literature exists on the topic. A minimal nullspace basis of a pencil may be computed through the calculation of the Kronecker canonical form. To our knowledge, the best known complexity for computing the Kronecker form of an $m \times n$ pencil is $O(m^2 n)$ [3, 24, 26]. Taking into account the dimension increase due to the linearization we may evaluate that computing a minimal basis of M would cost $O((md)^2(nd)) = O(m^2 n d^3)$. This approach is superior to ours concerning the quality of the

output basis which is minimal. However, it is unclear how it can lead to the reduction to polynomial matrix multiplication that we establish.

A second alternative and different linearization of the problem is to associate to M a generalized Sylvester matrix (i.e., a block-Toeplitz matrix [5]) or another type of resultant. This has been heavily used for control theory problems and in linear algebra. A polynomial vector of degree δ in the nullspace of M may be obtained from the nullspace of a block-Toeplitz of dimension about $n\delta$. This leads to costs too high by a factor of n when the degrees in a minimal nullspace basis are unbalanced. We are not aware of an approach based on successive compression here that would allow to save a factor n and to introduce polynomial matrix multiplication.

These two types of linearization correspond to two main approaches—based on state-space realizations or on resultants— for the problem of *coprime matrix fraction description* or *coprime factorization* [17, Chapter 6]. We see from (3) that we will use a solution to the latter problem a logarithmic number of times on the compressed matrices. If all matrices involved are of degree d, then we use the σ-basis algorithm of [1], and the corresponding reduction to polynomial matrix multiplication of [15]. A solution of the coprime factorization in case of unbalanced degree is, in a way similar to the block-Toeplitz approach, is faced with the question of saving a factor n in the cost. Known algorithms seem to have a cost driven only by the highest degree in the factorization, rather than by the sum of the involved degrees as we propose.

Our work is a derivation of an elimination scheme using minimal bases directly on polynomial matrices. Our compression / uncompression strategy can be compared to the techniques used for the staircase algorithm of [3, 26] for preserving a special structure. We somehow generalize the latter to the case of polynomial matrices for reducing the matrix description problem with input BA^{-1} to the polynomial matrix multiplication.

3   Preliminaries for polynomial matrices

We give here some definitions and results about minimal bases [13] and matrix fraction descriptions that will be used in the rest of the paper. For a comprehensive treatment we refer to [17, Chapter 6]. For a matrix $M \in \mathsf{K}[x]^{m \times n}$ of rank r and degree d, we call (left) nullspace the $\mathsf{K}(x)$-vector space of vectors $v \in \mathsf{K}(x)^m$ such that $vM = 0$. We will compute a basis of that space. The basis will be given by $m - r$ linearly independent polynomial vectors, and is related to the notion of minimal basis of the nullspace seen as a $\mathsf{K}[x]$-module.

Definition 3.1 *A basis* $N_1, \ldots, N_{m-r} \in \mathsf{K}[x]^m$ *with degrees* $\delta_1 \leq \ldots \leq \delta_{m-r}$ *of the nullspace of M seen as a $\mathsf{K}[x]$-module is called a* minimal basis *if any other nullspace basis with degrees* $\delta'_1 \leq \ldots \leq \delta'_{m-r}$ *satisfies* $\delta'_i \geq \delta_i$ *for* $1 \leq i \leq m - r$.

In the rest of the text, "basis" will usually refer to the vector space while "minimal basis" will refer to the module. The degrees $\delta_1, \ldots, \delta_{m-r}$ are structural invariants of the nullspace. They are called the minimal indices of the nullspace basis. The minimal indices of a nullspace basis of M are called the (left) *Kronecker indices* of M. A polynomial matrix $M \in \mathsf{K}[x]^{m \times n}$ is called *row-reduced* if its leading row coefficient matrix has full rank. It is called *irreducible* if its rank is full for all (finite) values of x (i.e., I_m is contained in the set of $\mathsf{K}[x]$-linear combinations of columns of M). These two definitions are used for characterizing minimal bases; we refer to [17, Theorem 6.5-10] for the proof of the following.

Theorem 3.2 *The rows of $N \in \mathsf{K}[x]^{m-r}$, such that $NM = 0$, form a minimal basis of the nullspace of M if and only if N is row-reduced and irreducible.*

A key point for keeping the cost of the computation low is the degree transfer between M and a minimal nullspace basis N. The McMillan degree of M of rank r is the maximum of the degrees of the determinants of $r \times r$ submatrices of M [17, Exercise 6.5-9].

Theorem 3.3 *The Kronecker indices and the McMillan degree of M satisfy*

$$\sum_{i=1}^{m-r} \delta_i \leq \text{McMillan-deg}\, M \tag{7}$$

with equality if M is irreducible.

Proof. Let I and J be row and column index sets such that the McMillan degree of the submatrix $M_{I,J}$ of M is equal to the one of M. By considering $M_{\cdot,J}$ we reduce ourselves to the full column-rank case. Define $I_c = \{1, \ldots, m\} \setminus I$, and the corresponding submatrices $A = M_{I,J}$ and $B = M_{I_c,J}$ of M. By unimodular column reduction we may assume that $M_{\cdot,J}$ is column reduced, since A carries the McMillan degree, BA^{-1} is proper. A minimal basis N gives the corresponding matrices $C = N_{\cdot,I}$ and $D = N_{\cdot,I_c}$ such that $CA + DB = 0$. The matrix D cannot be singular otherwise there would exist a vector $u \neq 0$ such that $uD = 0$ and $uC \neq 0$ (the latter since N is non-singular). This would give a non-zero vector u such that $uA = 0$ which is not possible. Hence, $D^{-1}C = BA^{-1}$. If M is irreducible, then since N is irreducible by definition, both latter fractions are irreducible. By [17, Theorem 6.5-1], $\deg \det D = \text{McMillan-deg}\, M$. Therefore, using the fact that $D^{-1}C$ is proper we know that $\deg \det D = \sum_{i=1}^{m-r} \delta_i$, and the Theorem is established. When M is not irreducible the same reasoning applies with $\deg \det D \leq \text{McMillan-deg}\, M$. $\qquad\square$

The reader may also refer to [2, Theorem 5.1] for bound (7). As dicussed in the introduction, Gaussian elimination is far too pessimistic when it results in a nullspace basis with degree sum in $\Theta(n^3 d)$. Theorem 3.3 shows that there exist minimal bases with degree sum in $O(n^2 d)$ whose computation should be cheaper.

We will use minimal bases in relation with left or right matrix fraction descriptions. A left fraction description $S^{-1}N$ is irreducible (or coprime) if any non-singular polynomial matrix and left common divisor U of S and N (i.e., $S = US'$ and $N = UN'$ for polynomial matrices S' and N') is unimodular. An analogous definition holds on the right.

Lemma 3.4 *The rows of $N = [N_p \; S]$, such that $NM = 0$, with S non-singular form a basis for the nullspace as a $\mathsf{K}[x]$-module if and only if $S^{-1}N_p$ is irreducible.*

Proof. We have that N is a basis if and only if it is irreducible, which in turn is equivalent to the fact that S and N_p are coprime [17, Lemma 6.3-6]. $\qquad\square$

For a rational matrix $\mathcal{H}$ we may define the $\mathsf{K}[x]$-module $\mathcal{P}_{\mathcal{H}}$ of polynomial vectors u such that $u\mathcal{H}$ is polynomial. We will use the following.

Lemma 3.5 *$S^{-1}N = \mathcal{H}$ is a coprime matrix description of $\mathcal{H}$ if and only if the rows of S form a basis of $\mathcal{P}_{\mathcal{H}}$.*

Proof. Consider T non-singular whose rows are in $\mathcal{P}_{\mathcal{H}}$. Then for a polynomial matrix M we have $T\mathcal{H} = TS^{-1}N = M$, hence $S^{-1}N = T^{-1}M$. Since $S^{-1}N$ is coprime, T is a left multiple of S [17, Lemma 6.5-5]. Conversely, if the rows of S form a basis of $\mathcal{P}_{\mathcal{H}}$, then $S^{-1}N$ is coprime. Otherwise, S would be a multiple of S_c for $S_c^{-1}N_c$ coprime, which would contradict the basis property. □

In Section 6 we will focus on computing only vectors of degrees bounded by a given δ in a nullspace minimal basis. We define their number $\kappa = \max\{1 \le i \le m - r \text{ s.t. } \delta_i \le \delta\}$ (the Kronecker indices are arranged in increasing order). Corresponding vectors are called κ *first minimal vectors* in the nullspace.

Remark 3.6 We will also manipulate the module generated by κ such vectors. As in Theorem 3.2, a corresponding submatrix $\tilde{N}$ with κ rows of N must be irreducible. As in Lemma 3.4, if $\tilde{N} = [\tilde{N}_p \quad \tilde{S}]$ then $\tilde{N}_p$ and $\tilde{S}$ have no left and non-singular common divisor other than unimodular. Since a minimal basis N of the nullspace is row-reduced, by the predictable-degree property [17, Theorem 6.3-13], any vector of degree less than δ must be in the sub-module generated by κ minimal vectors.

4   Matrix fraction descriptions for the nullspace

Let us consider a matrix $M = [A^T \quad B^T]^T \in \mathsf{K}[x]^{(n+p)\times n}$ of degree d as in (2) with A square $n \times n$ and invertible. Our study here and in next section focuses on the case $p \le n$ which is the heart of the method, and where all difficulties arise. The results here remain true but are trivial for $p > n$ (see Remark 6.6).

The rows of $\mathcal{H} = [H \quad -I_p] = [BA^{-1} \quad -I_p]$ form a nullspace basis of M. Hence, for N a minimal nullspace basis, there exists a transformation S in $\mathsf{K}(x)^{p\times p}$ such that $S\mathcal{H} = N$. With the special shape of $\mathcal{H}$ we deduce that S is a polynomial matrix in $\mathsf{K}[x]^{p\times p}$ whose columns are the last p columns of N. This leads to the following left matrix fraction description of $\mathcal{H}$:

$$\mathcal{H} = [H \quad -I_p] = [BA^{-1} \quad -I_p] = S^{-1}N. \tag{8}$$

The left fraction description $S^{-1}N$ must be irreducible otherwise there would exist another description $\mathcal{H} = (S')^{-1}N'$ with $N' \in \mathsf{K}[x]^{p\times(n+p)}$ having row degrees lexicographically smaller than the row degrees of N. Since $N'M = 0$ this would contradict the fact that N is minimal.

For reducing the cost of our approach we will introduce a (random) column compression H_p of H given by

$$H_p = HP = BA^{-1}P \in \mathsf{K}[x]^{p\times p} \tag{9}$$

with $P \in \mathsf{K}[x]^{n\times p}$.

In order to be appropriate for computing the nullspace of M, H_p must keep certain invariants of BA^{-1}. We establish in the rest of the section— see Proposition 4.2—that there exists a P such that, on the left, the description $H_p = S^{-1}(NP)$ *remains irreducible*. With the same P we show the existence, on the right, of a description whose *denominator matrix has relatively small degree*. The existence of such a P will ensure that the properties remains true for a random compression.

For proving above irreducibility and small degree properties after compression, we first linearize the problem by introducing a realization of the rational matrix H_p.

Lemma 4.1 *Let A be non-singular of degree less than d and determinantal degree $\nu \neq 0$ in $\mathsf{K}[x]^{n \times n}$. Let B be in $\mathsf{K}[x]^{p \times n}$. There exists a surjective function $\sigma : \mathsf{K}[x]^{n \times p} \to \mathsf{K}^{\nu \times p}$, and two matrices $X \in \mathsf{K}^{p \times \nu}$ and $A_o \in \mathsf{K}^{\nu \times \nu}$, such that for any P in $\mathsf{K}[x]^{n \times p}$*

$$H_p(x) = B(x)A(x)^{-1}P(x) = Q(x) + X(x - A_o)^{-1}\sigma(P), \tag{10}$$

with $Q \in \mathsf{K}[x]^{p \times p}$. If P is selected uniformly at random of degree at most $d - 1$, then $\sigma(P)$ is uniform random in $\mathsf{K}^{\nu \times p}$. Additionally, a matrix $S \in \mathsf{K}[x]^{p \times p}$ is the denominator of a left coprime description of BA^{-1} if and only if S is the denominator of a left coprime description of $X(x - A_o)^{-1}$.

Proof. We first establish (10) for B the identity matrix of dimension n and for A in column Popov form [27] (see also [17, §6.7.2]): A is column-reduced, i.e. its leading column coefficient matrix has full rank; in each row of A a unique entry has maximum degree and is monic. Let $d_1, d_2, \ldots, d_n$ be the column degrees of A, since A is in Popov form, $\nu = \sum_{i=1}^{n} d_i$. We first assume that the d_i's are greater than one. We follow the lines of the realization constructions in [17, §6.4]. Consider $D = \mathrm{diag}(x^{d_1}, x^{d_2}, \ldots, x^{d_n})$ and $\Psi = \mathrm{diag}([1 \ x \ \ldots x^{d_i-1}], 1 \leq i \leq n) \in \mathsf{K}^{n \times \nu}$. Since A is in column Popov form we have $A = D + \Psi A_L$ where $A_L \in \mathsf{K}^{\nu \times n}$ is given by the low degree coefficients of the entries of A. We also define $X = \mathrm{diag}([0, \ldots, 0, 1] \in \mathsf{K}^{1 \times d_i}, 1 \leq i \leq n) \in \mathsf{K}^{n \times \nu}$ and $D_o = \mathrm{diag}(C_{x^{d_1}}, C_{x^{d_2}}, \ldots, C_{x^{d_n}}) \in \mathsf{K}^{\nu \times \nu}$ whose diagonal blocks are matrices companion to the diagonal entries of D. One can directly check that $\Psi(x - D_o) = DX$. Taking $A_o = D_o - A_L X$ we get $\Psi(x - A_o) = \Psi(x - D_o + A_L X) = DX + \Psi A_L X$, hence $\Psi(x - A_o) = AX$, or, in other words,

$$A^{-1}\Psi = X(x - A_o)^{-1}. \tag{11}$$

If the row degrees of P are strictly lower than the d_i's then P may be decomposed into $P(x) = \Psi P_o$. This leads to $A^{-1}P = X(x - A_o)^{-1}P_o$ and we take $\sigma(P) = P_o$. If P has larger degrees, dividing P by A uniquely defines two polynomial matrices Q and R such that $R = P - AQ$ and such that the row degrees of R are less than the d_i's (see [17, Division Theorem 6.3-15]). Writing $R = \Psi R_o$ we get $A^{-1}P = A^{-1}(AQ + R) = Q + A^{-1}R = Q + X(x - A_o)^{-1}R_o$ and we take $\sigma(P) = R_o$.

Now, if some column degrees are zero, say exactly k of the d_i's, then for row and column permutations U_l and U_r and since A is in Popov form, we may write

$$U_l A U_r = \begin{bmatrix} \bar{A} & A_{12} \\ 0 & I \end{bmatrix}$$

where $\bar{A} \in \mathsf{K}[x]^{(n-k) \times (n-k)}$ has column degrees greater than one and A_{12} is a constant matrix in $\mathsf{K}^{(n-k) \times k}$. Applying (11) to $\bar{A}$ we get matrices $\bar{\Psi}, \bar{X}$ and $\bar{A}_o$ such that $\bar{A}^{-1}\bar{\Psi} = \bar{X}(x - \bar{A}_o)^{-1}$. Hence, if Ψ and X are constructed by augmenting $\bar{\Psi}$ and $\bar{X}$ with k zero rows, we get

$$(A^{-1}U_l^{-1})\Psi = U_r \begin{bmatrix} \bar{A}^{-1} & -\bar{A}^{-1}A_{12} \\ 0 & I \end{bmatrix} \Psi = U_r X(x - \bar{A}_o)^{-1}. \tag{12}$$

Then σ may be defined as previously, if $R = \Psi R_o$ is the remainder of the division of $U_l P$ by $U_l A U_r$, then $\sigma(P) = R_o$.

If R_o, with $R = U_l^{-1}\Psi R_o$, is the image of a matrix P of degree less than $d - 1$, we have $P = R + AQ$. For another matrix $R_0' \in \mathsf{K}^{\nu \times p}$, with $R' = \Psi R_o'$, this defines a unique matrix

$P' = R' + AQ$ of degree less than $d - 1$ such that $\sigma(P') = R'_o$. Hence any two matrices in $\mathsf{K}^{\nu \times p}$ have the same number of inverse images of degree less that $d - 1$ by σ. Together with the fact that the restriction of σ to the matrices of degree less than $d - 1$ is surjective, this shows that a uniform random choice of P of degree less than $d - 1$ leads to a uniform random choice $\sigma(P)$ in $\mathsf{K}^{\nu \times p}$.

For general matrices $A \in \mathsf{K}[x]^{n \times n}$ and $B \in \mathsf{K}[x]^{p \times n}$, let V be unimodular such that $\tilde{A} = AV$ is in Popov form. From the above we know that

$$\tilde{A}(x)^{-1} P(x) = Q(x) + X(x - A_o)^{-1} \sigma(P).$$

Taking $\tilde{X} = B(x)V(x)X - Q_B(x)(x - A_o)$ the remainder of the division of $B(x)V(x)X$ by $(x - A_o)$ this leads to

$$B(x)V(x)\tilde{A}(x)^{-1} P(x) = B(x)V(x)Q(x) + Q_B(x) + \tilde{X}(x - A_o)^{-1} \sigma(P)$$

which is

$$B(x)A(x)^{-1} P(x) = \tilde{Q}(x) + \tilde{X}(x - A_o)^{-1} \sigma(P)$$

where $\tilde{Q}$ is a matrix polynomial and X is a constant matrix as the remainder of a division by a matrix of degree one. This establishes (10) with an appropriate σ.

It remains to show the property on denominator matrices S. We use Lemma 3.5 and prove that SBA^{-1} and $SX(x - A_o)^{-1}$ are polynomials for the same denominator matrices S. In (12) the matrix $[I_{n-k}\ 0]^T$ is a submatrix of Ψ. Therefore $\tilde{A}^{-1}$ and $\tilde{A}^{-1}\Psi$ are polynomials for the same polynomial matrices SBV (B and V are fixed), and SBA^{-1} and $SBA^{-1}\Psi$ are polynomials for the same S. Finally notice that $S\tilde{X}(x - A_o)^{-1}$ is polynomial for the same set of matrices S since $\tilde{X}(x - A_o)^{-1}$ is the fractional part of $BA^{-1}\Psi$.   □

We now state the required properties for P, and prove them on the realization (10).

Proposition 4.2 *Let $A \in \mathsf{K}[x]^{n \times n}$ be non-singular of degree less than d and determinantal degree ν, and let $B \in \mathsf{K}[x]^{p \times n}$. Assume that $S \in \mathsf{K}[x]^{p \times p}$ is any denominator of a left irreducible fraction description of BA^{-1}. Then there exists a matrix P of degree less than $d - 1$ in $\mathsf{K}[x]^{n \times p}$ such that*

$$H_p = BA^{-1}P = CT^{-1} \tag{13a}$$

$$= S^{-1}N_p \in \mathsf{K}[x]^{p \times p} \tag{13b}$$

where CT^{-1} is a right irreducible description with $T \in \mathsf{K}[x]^{p \times p}$ of degree less than $\lceil \nu/p \rceil \leq (n/p)d + 1$, and where $S^{-1}N_p$ is a left irreducible description.

Proof. For $\nu = 0$ (BA^{-1} is a polynomial), the results hold with $T = S = I$. In the general case Lemma 4.1 gives

$$H_p(x) = B(x)A(x)^{-1}P(x) = Q(x) + X(x - A_o)^{-1}\sigma(P).$$

For studying denominators of irreducible descriptions of H_p one can forget its polynomial part, hence we now focus on the fraction $X(x - A_o)^{-1}\sigma(P)$. Lemma 4.1 also gives that there exists a left irreducible fraction description of $X(x - A_o)^{-1}$ with denominator S:

$$S(x)^{-1}N'(x) = X(x - A_o)^{-1}. \tag{14}$$

Through the application σ, choosing an adequate polynomial $P \in \mathsf{K}[x]^{n \times p}$ for H_p reduces to choosing an adequate constant $Y \in \mathsf{K}^{n \times p}$ for $X(x - A_o)^{-1}$.

We now use the formalism of minimum generating polynomials of matrix sequences introduced in [33, 34] and [21, §2]. By Lemma 2.8 in [21], finding a matrix Y such that $X(x - A_o)^{-1}Y = C'(x)T(x)^{-1}$ with T as expected, reduces to finding an appropriate Y with T the a right minimum generator of the sequence $\{XA_o^iY\}_{i \geq 0}$. From (14) we have that S is a left minimum generator of $\{XA_o\}_{i \geq 0}$. Therefore one may use the construction of [33, Corollary 6.4], together with [21, Theorem 2.12]. This provides a Y and a right minimum generator T of $\{XA_oY\}_{i \geq 0}$ with determinantal degree equal to the determinantal degree μ of S, and with degree bounded by $\lceil \mu/p \rceil \leq \lceil \nu/p \rceil$. A matrix P of degree $d - 1$ such that $\sigma(P) = Y$ is an appropriate choice for concluding the proof of (13a). Indeed, $C = QT + C'$ and T gives an appropriate right irreducible description. The corresponding left description $S^{-1}N_p$ is coprime by [17, Lemma 6.5-6], which establishes (13b). $\square$

Proposition 4.2 shows that if P has symbolic entries, then a right coprime description of $H_p = BA^{-1}P = CT^{-1}$ can be found with a denominator matrix of degree less than d, and with the same left denominators as for BA^{-1}.

Remark 4.3 Proposition 4.2 establishes the existence of appropriate descriptions $S^{-1}N_p$ and CT^{-1} for a symbolic P. As a consequence of Lemma 4.1, [33, Corollary 6.4] or [21, Section 2], and by evaluation [12, 35, 29], the same denominator properties will hold for a random matrix P.

5   From compressed minimal bases to minimal bases

As seen in Introduction, we will compute a small basis for the nullspace of the input matrix as a set of successive minimal bases of matrices like in (2). The latter minimal bases are computed in two main steps. We first compute the expansion of $H_p = BA^{-1}P$ and reconstruct a corresponding fraction (3) with denominator S. Then, if P is such that H_p satisfies (13b), we know that $S[BA^{-1} \; -I_p]$ is a polynomial matrix N, which by construction satisfies $NM = 0$.

In the spirit of the scalar polynomial case and of [1] for the matrix case, the reconstruction may be done via Padé approximation, and through the computation of particular bases of the nullspace of $[-I_p \; H_p^T]^T$. Indeed we have the equivalence between $S^{-1}N_p = H_p$ and $[N_p \; S] \cdot [-I_p \; H_p^T]^T = 0$. Hence the purpose of this section is to identify the bases of the nullspace of $[H_p^T \; -I_p]^T$ that actually lead to *minimal* bases N for M.

Through a conditioning of M let us first specify the location of the leading degree terms in the latter bases (see Theorem 3.3).

Lemma 5.1 *For M as in (1) there exist a matrix $Q \in \mathsf{K}^{(n+p) \times (n+p)}$ such that the McMillan degree of the top $n \times n$ submatrix of QM is equal to the McMillan degree of QM (and of M). This implies that if N is a minimal basis of the nullspace of QM, then $S = N_{\cdot,p+1..n+p}$ is row-reduced with row degrees the Kronecker indices $\delta_1, \ldots, \delta_p$.*

Proof. If I is a set of row indices such that $M_{I,\cdot}$ has determinantal degree $\sum_{i=1}^p \delta_i$, and let $Q \in \mathsf{K}^{(n+p) \times (n+p)}$ be a row permutation π such that $\pi(I) = \{1, \ldots, n\}$. Then the top n rows of QM give the McMillan degree. From [2, Theorem 5.1 (b)] the dominant degrees in a minimal nullspace basis of QM are in the columns $\{1, \ldots, n+p\} \setminus \{1, \ldots, n\} = \{n+1, \ldots, n+p\}$. $\square$

Remark 5.2 The property given by the multiplication by Q in Lemma 5.1 will hold for a random Q over K (compare to Remark 4.3).

In next sections, nullspace vectors v^T for M are easily obtained from nullspace vectors w^T for QM, indeed $v^T = w^T Q$ satisfies $v^T M = w^T Q M = 0$. This conditioning of M—and implicitly of N—will alllow us to compute S, and then deduce N, from a shifted minimal basis for the nullspace of $[-I_p \quad H_p^T]^T$. Shifted bases are defined as usual minimal bases by changing the notion of degree. For $\bar{t}$ a fixed multi-index in $\mathbb{Z}^m$, the $\bar{t}$-degree of a vector v in $\mathsf{K}[x]^m$ is

$$\bar{t}\text{-deg }v = \max_{1 \le i \le m} \{\deg v_i - \bar{t}_i\}. \tag{15}$$

Definition 5.3 A basis of a $\mathsf{K}[x]$-submodule of $\mathsf{K}[x]^m$, given by the rows of a matrix N, is called $\bar{t}$-minimal if N is row-reduced with respect to the $\bar{t}$-degree. Equivalently, $N \cdot x^{-\bar{t}}$ is row-reduced with respect to the usual degree (see [2, Definition 3.1]).

For $\bar{t} = [0, \dots, 0]$ the definition corresponds to the usual definition of minimal bases. The value $\bar{t} = [(d-1)_p, 0_p]$ below, where $(d-1)_p$ and 0_p respectively denote the values $d-1$ and 0 repeated p times, is chosen from the degree $d-1$ of the compression matrix P of Proposition 4.2. This value forces the row reduction in the last columns of the bases.

Proposition 5.4 Let $M \in \mathsf{K}[x]^{(n+p) \times n}$ be of full rank such that the matrix S, formed by the last p columns of a minimal basis N for its nullspace, is row-reduced with row degrees the Kronecker indices $\delta_1, \dots, \delta_p$. Assume that $P \in \mathsf{K}[x]^{n \times p}$ satisfies (13b). Let $\bar{t} = [(d-1)_p, 0_p] \in \mathbb{N}^{2p}$. Then $[N_p \quad S]$ is a $\bar{t}$-minimal basis for the nullspace of $[-I_p \quad H_p^T]^T$ if and only if $N = S[BA^{-1} \quad -I_p] = [\bar{N} \quad S]$ is a minimal basis for the nullspace of M.

Proof. We first prove that the condition is sufficient. If $[\bar{N} \quad S]$ is a minimal basis for the nullspace of M, the description $S^{-1}\bar{N}$ of BA^{-1} is irreducible (Lemma 3.4). The rows of $[N_p \quad S] = [\bar{N}P \quad S]$ are in the nullspace of $[-I_p \quad H_p^T]^T$. They form a basis of the latter nullspace since otherwise $S^{-1}N_p$ would not be irreducible which would contradict (13b). We assume that the rows of the bases are arranged by increasing degrees. The ith row of $\bar{N}P$ has degree less than $\delta_i + d - 1$, hence its $\bar{t}$-degree is less than δ_i, which in turn is less that the $\bar{t}$-degree of the ith row of S. Which shows that $[N_p \quad S]$ is row-reduced with respect to the $\bar{t}$-degree since S is row-reduced by assumption on M and by Theorem 3.2). The $\bar{t}$-minimality follows.

Conversely, if $[N_p \quad S]$ is a $\bar{t}$-minimal basis for the nullspace of $[-I_p \quad H_p^T]^T$, then by (13b) $N = S[BA^{-1} \quad I] = [\bar{N} \quad S]$ is a polynomial matrix, and $N_p = \bar{N}P$. Since $[N_p \quad S]$ is a basis, $S^{-1}N_p$ is irreducible (Lemma 3.4), hence also by (13b), $S^{-1}\bar{N}$ is irreducible and $[\bar{N} \quad S]$ is a basis for the nullspace of M. It remains to show that $[\bar{N} \quad S]$ is row-reduced. There exists a unimodular $p \times p$ matrix U such that $[\bar{N} \quad S] = U[\bar{L} \quad R]$ where $[\bar{L} \quad R]$ is a row-reduced basis for the nullspace of M, hence where R is row-reduced by assymption on M. By the predictable-degree property [17, Theorem 6.3-13], the degree of the ith row of S is $\deg S_i = \max_{j=1,\dots,p}\{\delta_j + \deg U_{ij}\}$. The degree of the ith row of $\bar{N}$ is $\deg \bar{N}_i \le \max_{j=1,\dots,p}\{\delta_j + \deg U_{ij}\}$. The $\bar{t}$-degree of the ith row of N_p may then be bounded as follows,

$$\bar{t}\text{-deg }((N_p)_i) = \deg((\bar{N}P)_i) - (d-1) \le \max_{j=1,\dots,p} \{\delta_j + \deg U_{ij}\} \le \deg S_i = \bar{t}\text{-deg }S_i.$$

Since $[N_p \ \ S]$ is row-reduced with respect to the $\bar{t}$-degree, this implies that S itself is row-reduced. By assumption on M the degrees of S are dominating in $[\bar{N} \ \ S]$, hence the latter matrix also is row-reduced and is a minimal basis for the nullspace of M. $\square$

For compressing matrices P which satisfy (13b), Proposition 5.4 establishes strong links between the nullspace of $[-I_p \ \ H_p^T]^T$ and the one of M. In particular, $[-I_p \ \ H_p^T]^T$ and M have the same Kronecker indices. For any given δ, there is a one-to-one correspondence between the vectors of $\bar{t}$-degree δ in the nullspace of $[-I_p \ \ H_p^T]^T$, and those of degree δ in the nullspace of M. This is seen from the "S" common part of the bases.

6   Computing nullspace minimal vectors

We still consider a full column-rank matrix M be of degree d as in (1). Let δ be a fixed integer and $\kappa(=\kappa(\delta))$ be the number of vectors of degree less than δ in a minimal basis N of the $\mathsf{K}[x]$-nullspace of M. In this section we study the cost for computing κ such vectors.

Algorithm Nullspace minimal vectors (M,δ)

> *Input:* $M \in \mathsf{K}[x]^{(n+p)\times n}$ of degree d, a degree threshold δ,
> $\quad\quad\quad\quad$ M has full column-rank.
> *Output:* $\kappa = \max\{1 \leq i \leq p$ s.t. $\delta_i \leq \delta\}$,
> $\quad\quad\quad\quad$ independent vectors $N_i \in \mathsf{K}[x]^{n+p}$ of degree δ_i, $1 \leq i \leq \kappa$, in the nullspace of M.

> (a) $M := QM$ for a random $Q \in \mathsf{K}^{(n+p)\times(n+p)}$;
> (b) $M := M(x + x_0)$ for x_0 random in K;
> $\quad\quad$ $A := M_{1..n,1..n}$, **if** $\det A(0) = 0$ **then** **fail**;   /* rank M is probably less than n */
> $\quad\quad$ $B := M_{n+1..n+p,1..n}$;
> $\quad\quad$ $\eta := \delta + d + \lceil nd/p\rceil$;
> (c) $H :=$ expansion of BA^{-1} mod x^η;
> (d) $H_p := HP$ for P random in $\mathsf{K}[x]^{n\times p}$ of degree less than $d-1$;
> $\quad\quad$ $\bar{t} = [(d-1)_p, 0_p] = [d-1, \ldots d-1, 0, \ldots, 0] \in \mathbb{N}^{2p}$;
> (e) $L := [\mathcal{N}_p \ \ \mathcal{S}] :=$ a σ-basis with respect to $\bar{t}$ for $[-I_p \ \ H_p^T]^T$ of order η;
> (f) $\kappa :=$ the number of rows of $[\mathcal{N}_p \ \ \mathcal{S}]$ of $\bar{t}$-degree at most δ;
> $\quad\quad$ select the corresponding κ rows S_i of $\mathcal{S}$ by increasing degrees, $1 \leq i \leq \kappa$;
> (g) $N_i := S_i[H \ \ -I_p]$ mod $x^{\delta+1}$, $1 \leq i \leq \kappa$;
> $\quad\quad$ $N_i(x) := N_i(x - x_0)Q$, $1 \leq i \leq \kappa$;
> $\quad\quad$ $\lambda := \#\{N_i$ s.t. $N_iM = 0\}$
> (h)   **if** $\lambda \neq \kappa$ **then** **fail**;   /* certification of κ */
> $\quad\quad$ $N^{(\delta)} :=$ the $\kappa \times (n+p)$ matrix whose rows are the N_i's;
> (i)   **if** $N^{(\delta)}$ is not row-reduced **then** **fail**;   /* certification of the minimality */
> $\quad\quad$ **else** **return** κ and N_i, $1 \leq i \leq \kappa$. $\square$

Algorithm Nullspace minimal vectors above starts with lifting on a compressed matrix (Proposition 4.2). Then it partially (subject to the degree threshold) computes a denominator matrix S through a partial $\bar{t}$-minimal basis computation. Using Proposition 5.4 the target nullspace vectors are finally obtained.

We prove the algorithm and its cost in the rest of the section. Step (a) is the conditioning seen in Section 5 to ensure the degree dominance of the last p columns of N. Together with the

randomized compression of Step (d) studied in Proposition 4.2 this will allow the computation of S at Step (e). Step (b) is a randomized choice for working with a matrix A non-singular at $x = 0$. The latter condition is required for computing at Step (c) the expansion of BA^{-1} by lifting [31, 32]. Step (e) partly reconstructs a description $S^{-1}N_p$ from a truncated expansion of H_p. The computation is explained in Lemma 6.3 below, and the selection of small degree rows at Step (f) is justified. Our approach for the reconstruction is very close to the column reduction of [15, §3]. A degree less than δ in S corresponds to a $\bar{t}$-degree (see (15)) less than δ in $[N_p \ \ S]$ (the compression using P increases the degree in N_p by $d-1$), and to a degree less than δ in N. Step (g) applies Proposition 5.4 for partly reconstructing the nullspace of M, and Steps (h) and (i) certify the outputs.

We now partly reconstruct H_p (i.e., partly compute a $\bar{t}$-minimal basis at Step (e), and the denominator matrix S at Step (f)) from a minimal "nullspace basis expansion"—or σ-basis [1]. This is done by generalizing [15, §3] or [2, §4.2], especially for the partial computation aspects.

Definition 6.1 *Let G be in $\mathsf{K}[[x]]^{q \times p}$. Let $\bar{t}$ be a fixed multi-index in $\mathbb{Z}^q$. A σ-basis of (matrix-)order d with respect to $\bar{t}$ for G is a matrix polynomial L in $\mathsf{K}[x]^{q \times q}$ such that:*

I) $L(x)G(x) \equiv 0 \mod x^d$;

II) *every $v \in \mathsf{K}[x]^q$ such that $v(x)G(x) = O(x^d)$ admits a unique decomposition $v^T = \sum_{i=1}^{q} \alpha_i L_i$ where, for $1 \leq i \leq q$, L_i is the ith row of L, and α_i is a scalar polynomial in $\mathsf{K}[x]$ such that $\deg \alpha_i + \bar{t}$-$\deg L_i \leq \bar{t}$-$\deg v$.*

The reader may notice that we have slightly adapted the notion of order of the original Definition 3.2 of [1] for a fully matrix point of view. We also use the notion of shifted degree (see [2]) equivalently to the notion of defect used in [1, Definition 3.1]. The following shows that a σ-basis to sufficiently high order contains a minimal basis.

Lemma 6.2 *Let us assume that a minimal nullspace basis of G has κ vectors of $\bar{t}$-degree at most δ, and consider a σ-basis L with respect to $\bar{t}$. For an approximation order greater than $\delta + 1$, at least κ rows in L have $\bar{t}$-degree at most δ.*

Proof. See the proof of [16, Proposition 5]. We consider the κ rows of degree less than δ in a minimal nullspace basis of G. We order them by increasing degrees $\delta_1, \delta_2, \ldots, \delta_\kappa$. The first row v_1 has degree δ_1 therefore by II) of Definition 6.1, v_1 can be written as

$$v_1 = \sum_{i=1}^{q} \alpha_i L_i, \text{ with } \deg \alpha_i + \bar{t}\text{-}\deg L_i \leq \bar{t}\text{-}\deg v_1 = \delta_1.$$

We deduce that one row of L has $\bar{t}$-degree at most δ_1. Now if L has $i-1$ rows of degrees $\delta_1, \ldots, \delta_{i-1}$, with v_i of $\bar{t}$-degree δ_i, then the same reasoning as for v_1 shows that L has a row of degree less than δ_i, linearly independent with respect to the first $i-1$ chosen ones. The proof is concluded with $i = \kappa$. □

Next lemma identify the situation when a σ-basis will give the exact information we need. We assume that we are in the situation of Proposition 5.4, in particular S in the minimal bases has row degrees $\delta_1, \ldots, \delta_p$, the Kronecker indices of M and of $[-I_p \ \ H_p^T]^T$. We fix a value δ and define $\kappa = \max\{1 \leq i \leq p \text{ s.t. } \delta_i \leq \delta\}$, and $\bar{t} = [(d-1)_p, 0_p] \in \mathbb{N}^{2p}$.

Lemma 6.3 *Let us assume we are in the situation of Proposition 5.4. Let L be a σ-basis for $[-I_p \;\; H_p^T]^T$, with respect to $\bar{t}$, and of order of approximation at least $\eta = \delta + d + \lceil nd/p \rceil$. Then exactly κ rows of L have $\bar{t}$-degree at most δ, are in the nullspace of $[-I_p \;\; H_p^T]^T$, and have $\bar{t}$-degrees $\delta_1, \ldots, \delta_\kappa$.*

Proof. We generalize the proof of [15, Lemma 3.7] to the partial computation of the basis, and to the shifted case. We first verify that at most κ rows of L have $\bar{t}$-degree less than δ, then we prove their existence. Note that the rows of L are linearly independent [1].

Let $L_i = [\bar{L}_i \;\; S_i] \in \mathsf{K}[x]^{2p}$ be a row of L of $\bar{t}$-degree at most δ. From I) in Definition 6.1,

$$S_i(x)H_p(x) \equiv \bar{L}_i(x) \mod x^\eta,$$

and from the assumption (13a) on P,

$$S_i(x)C(x) \equiv \bar{L}_i(x)T(x) \mod x^\eta. \tag{16}$$

We now look at the degrees in both sides of latter identity. We have $\deg S_i = \bar{t}\text{-}\deg S_i \leq \delta$. By assumption on M, the degree of BA^{-1} is at most zero, hence the degree of H_p is at most $d - 1$. The latter is also true for CT^{-1} in (13a), which implies that $\deg C \leq \deg T + d - 1$. Using the degree bound on T in Proposition 4.2, the left side term of (16) thus have degree at most $\eta - 1$. In addition, $\deg \bar{L}_i = \bar{t}\text{-}\deg(\bar{L}_i) + (d-1) \leq \delta + d - 1$. Hence both sides in (16) have degree at most $\eta - 1$ and we deduce that

$$S_i(x)C(x) = \bar{L}_i(x)T(x). \tag{17}$$

It follows that $L_i = [\bar{L}_i \;\; S_i]$ is in nullspace of $[-I_p \;\; H_p^T]^T$. Using the equivalence with the nullspace of M in Proposition 5.4, one may associate to $L_i = [\bar{L}_i \;\; S_i]$ a row vector $N_i = [\bar{N}_i \;\; S_i]$, with $\bar{N}_i P = \bar{L}_i$, of degree less than δ in the nullspace of M (the "S" part is row-degree dominant). Since the rows L_i are linearly independent, the rows N_i of degree less than δ, corresponding to the L_i's of $\bar{t}$-degree less than δ, are linearly independent. At most κ such rows can exist.

We now show that κ rows of $\bar{t}$-degree at most δ exist in L. We consider the κ rows of degrees $\delta_1, \ldots, \delta_\kappa$ in a minimal basis $N = [\bar{N} \;\; S]$ of the nullspace of M. They give κ rows of $\bar{t}$-degree at most δ in the nullspace of $[-I_p \;\; H_p^T]^T$. Using Lemma 6.2 they lead to κ rows of $\bar{t}$-degree at most δ in L, which are in the nullspace by (17), hence their $\bar{t}$-degrees are $\delta_1, \ldots, \delta_\kappa$ by minimality. Note that the linear independency in N is preserved for the nullspace of $[-I_p \;\; H_p^T]^T$ since the column-reduced part S is in common. $\qquad\square$

From there, we can establish the correctness of the algorithm.

Proposition 6.4 *Let $M \in \mathsf{K}[x]^{(n+p)\times n}$ be of full column-rank with Kronecker indices $\delta_1, \ldots, \delta_p$. Algorithm* Nullspace minimal vectors *with inputs M and $\delta \in \mathbb{N}$ returns $\kappa = \max\{1 \leq i \leq p \text{ s.t. } \delta_i \leq \delta\}$, and κ first minimal vectors of the nullspace of M. The algorithm is randomized, it either fails or returns correct values (Las Vegas fashion).*

Proof. We first verify that if the random choices of x_0, Q and P work as expected then the result is correct. We will then prove that if the algorithm does not return fail then we are in the previous case. Note that the random shift x_0 does not modify the problem. Indeed, $\operatorname{rank} M(x) = \operatorname{rank} M(x + x_0)$, and since a matrix whose rows form a minimal basis is irreducible, the Kronecker indices are invariant under a shift.

Using Lemma 5.1, the role of Q is twofold: the top $n \times n$ submatrix of M becomes non-singular, and the dominant degrees in the nullspace are in the last columns. If $\det A(0) \neq 0$ then the rest of the computation is valid, in particular the expansion of BA^{-1} at Step (c). The basis L of order η as required can be computed from the expansion of BA^{-1} to the order η [1, 15]. If the choices of Q and P are successful then Lemma 6.3 ensures that the value of κ is correct; the corresponding rows of L are in the nullspace of $[-I_p \quad H_p^T]^T$. The nullspace correspondence of Proposition 5.4 then shows that $S_i[BA^{-1} \quad -I_p]$ is a polynomial row of degree less than δ, hence the computation of N_i can be done modulo $x^{\delta+1}$.

We now study the certification of the outputs. If $\det A(0) \neq 0$ then we know that M has full column-rank. The algorithm may then potentially fail with respect to the output value κ, there could actually be less or more minimal vectors of degrees at most δ. It may also fail with respect to the minimality of the output vectors. In any case, the computation of λ ensures that the returned N_i's are in the nullspace.

To avoid confusion we now denote by κ_o the output value and keep κ for the correct (unknown) value. Let us first see that $\kappa_o \geq \kappa$. Indeed, to the κ rows $N_i = [\bar{N}_i \quad S_i]$ of degree less than δ is in the nullspace of M, one may associate κ rows $[\bar{N}_i P \quad S_i]$ of $\bar{t}$-degree less than δ in the nullspace of $[-I_p \quad H_p^T]^T$. Since $\det A \neq 0$, we know that the S_i's are linearly independent. Hence we have κ linearly independent rows of $\bar{t}$-degree less than δ in the nullspace of $[-I_p \quad H_p^T]^T$. Then by Lemma 6.2, there must be $\kappa_o \geq \kappa$ rows of $\bar{t}$-degree less than δ in L. If $\lambda = \kappa_o$ then we have found κ_o linearly independent rows (from the S_i's) of degree less than δ (the degree is forced by construction at Step (g)) in the nullspace of M (test at Step (h)), hence $\kappa_o > \kappa$ cannot happen, and $\kappa_o = \kappa$. The returned value κ is always correct. In the latter case the returned vectors are linearly independent in the nullspace and satisfy the degree constraint.

We finally show that the returned vectors must be minimal. The corresponding κ rows, say $L_i = [\bar{L}_i \quad S_i]$ for $1 \leq i \leq \kappa$, in the nullspace of $[-I_p \quad H_p^T]^T$, must be minimal. Otherwise, by II) of Definition 6.1, a row of smaller $\bar{t}$-degree would have been selected in L. In particular, the matrix formed by the $\bar{L}_i$'s and the one formed by the S_i's are left relatively prime (no common left divisor other than unimodular). The κ computed rows $N_i = [\bar{N}_i \quad S_i]$ satisfy $\bar{N}_i P = \bar{L}_i$, hence the matrix formed by the $\bar{N}_i$'s and the one formed by the S_i's are also left relatively prime. Let $N_o^{(\delta)}$ be the $\kappa \times (n+p)$ matrix whose rows are the computed N_i's, and let $N^{(\delta)}$ be a $\kappa \times (n+p)$ matrix whose rows are κ first minimal vectors for the nullspace of M. Then the primality implies that there exist a unimodular U such that $N_o^{(\delta)} = UN^{(\delta)}$ (see Remark 3.6). Therefore the rows of $N_o^{(\delta)}$ have minimal degrees if and only if $N_o^{(\delta)}$ is row-reduced. The check is made at Step (i). □

From the arguments used in the proof of Proposition 6.4 we see that the algorithm may fail because the computed value κ is too large. This will essentially happen for bad choices of P, when the nullspace of the compressed matrix (see (13b)), and the approximating σ-basis (see (13a)), does not reflect the nullspace of M correctly. Then, even for correct values of κ, the minimality may not be ensured without the test at Step (i). A bad choice of Q, depending on P, may lead to a row reduction in the non-dominant part of the basis (see Lemma 5.1), and to a loss of minimality (see Proposition 5.4)*. A correctly computed value of κ may lead to a smaller value λ after the truncation (g) of a non-minimal vector.

---

*An improvement would be to combine both conditionings into a unique one with three different effects, left and right fractions for H_p, and location of the dominant degrees.

We also note that the minimality condition could be relaxed in the algorithm. Avoiding the last certificate would lead to the Las Vegas computation of κ independent vectors (possibly non-minimal) in the nullspace.

Corollary 6.5 *Let $M \in \mathsf{K}[x]^{(n+p) \times n}$ be of full column-rank and degree d with $1 \leq p \leq 2n$, and let $d \leq \delta \leq nd$. Minimal independent vectors in the nullspace of M, of degrees the Kronecker indices less than δ, can be computed by a randomized Las Vegas (certified) algorithm in*

$$O(\lceil p\delta/nd \rceil \, \mathsf{MM}(n, d) \log n + (n/p)\mathsf{MM}(p, \delta) + \mathsf{MM}(p, \delta) \log \delta + n^2 \mathsf{B}(d) \log n)$$

operations in K*. The cost is*

$$O(\mathsf{MM}(n, d) \log(nd) + n^2 \mathsf{B}(d) \log n + n\mathsf{M}(nd)) \tag{18}$$

when $p\delta/(nd) = O(1)$.

Proof. We use either (4) or (5), and the corresponding simplifications (6) for studying the cost of Algorithm Nullspace minimal vectors.

Steps (a) and (b) uses $O(\mathsf{MM}(n, d) + n^2\mathsf{M}(d))$ operations. From [32, Proposition 15], the cost for computing the expansion H is $O(\log(\eta/d)\lceil p\eta/nd \rceil \mathsf{MM}(n, d) + \overline{\mathsf{MM}}(n, d))$. This gives $O((\lceil (p\delta)/(nd) \rceil \, \mathsf{MM}(n, d) + n^2\mathsf{B}(d)) \log n)$ for $\eta = O(\delta + nd/p)$. Step (d) is a polynomial matrix multiplication that can be done in $O(\mathsf{MM}(n, d))$ operations. For the computation of the σ-basis at Step (e) we use the algorithm of [15, §2] based on polynomial matrix multiplication. The corresponding cost from [15, Theorem 2.4] is $O(\mathsf{MM}'(p, \eta) + \eta\mathsf{MM}(p))$, hence $O(\mathsf{MM}(p, \delta) \log \delta)$, or $O(\mathsf{MM}(n, d) \log d)$ if $p\delta/nd = O(1)$. Step (g) is a polynomial matrix multiplication modulo $x^{\delta+1}$ that can be computed in $(n/p)\mathsf{MM}(p, \delta)$ operations, this is $O(\mathsf{MM}(n, d) \log d + n\mathsf{M}(nd))$ when $p\delta/nd = O(1)$. The shift of the N_i's is done in at most $O(\sum_{i=1}^{p} n\mathsf{M}(\delta_i))$ operations, which is less than $O((n/p)\mathsf{MM}(p, \delta))$, or than $O(n\mathsf{M}(nd))$ for $p\delta/nd = O(1)$. The subsequent multiplication by Q has lower cost. Then we compute $N_i M$ for $1 \leq i \leq \kappa$, where N_i has degree δ_i, and $\sum_{i=1}^{p} \delta_i \leq nd$. Doing this computation directly as the product of a $\kappa \times (n + p)$ matrix with possibly large degrees, by an $(n + p) \times n$ matrix of degree d would be too expensive. Instead, we split the large degree entries of the N_i's and form an $n \times (n+p)$ matrix $\tilde{N}$ of degree d, and recover the products $N_i M$ from the multiplication $\tilde{N}M$. The corresponding cost is $O(\mathsf{MM}(n, d))$. The final check (i) is done in $O(n^\omega + n^2d)$ operations. $\square$

We see from (18) that computing vectors in the nullspace at essentially the cost of multiplying two polynomial matrices relies on the *compromise between p and δ*. The algorithm is a combination of *matrix lifting* and *matrix fraction reconstruction*. Many vectors of small degrees are computed using lifting to a limited order and large matrix reconstruction. Conversely, few vectors of large degrees are computed from a high-order lifting and reconstruction with matrices of small dimensions.

Remark 6.6 Note that the random compression P is introduced for $p < n$. Still, the algorithm is proven for $p \geq n$. In the latter case however, for the sake of simplicity, one may work directly with $H_p = H$ at Step (d).

7   Small degree nullspace basis computation

Corollary 6.5 which uses for (18) a compromise between p and δ, does not directly allow a low-cost computation of large degree vectors in a nullspace of large dimension. For the latter situation, and for computing a whole set of linearly independent vectors in the nullspace of a matrix M in $\mathsf{K}[x]^{(n+q)\times n}$, we need to successively restrict ourselves to smaller nullspace dimensions (while increasing the degree). Here we take the notation $m = n+q$ for M as in (1). We keep the notation p for submatrices (2), and successive compressions, as in Sections 4-6 .

7.1   Full column-rank and $n < m \leq 2n$ case

Let $M \in \mathsf{K}[x]^{(n+q)\times n}$ with $1 \leq q \leq n$ be of degree d and rank n. The way we restrict ourselves to smaller nullspaces is derived from the following observation. Let C be in $\mathsf{K}^{(n+p)\times(n+q)}$ with $1 \leq p \leq q$. If $CM \in \mathsf{K}[x]^{(n+p)\times n}$ also has full column-rank, then let $\delta_1, \ldots, \delta_p$ be its Kronecker indices, and with the degree threshold $\delta = 2nd/p$ take $\kappa(\delta) = \max\{1 \leq i \leq p \text{ s.t. } \delta_i \leq \delta\}$. Since $\sum_1^p \delta_i \leq nd$, at most $nd/\delta = p/2$, hence $\lfloor p/2 \rfloor$, vectors in a minimal basis of the nullspace of CM may have degrees more than δ, therefore $\kappa(\delta) \geq \lceil p/2 \rceil$. From at least $p/2$ minimal vectors $D_1, \ldots, D_\kappa \in \mathsf{K}[x]^{n+p}$ of degrees at most $2nd/p$ in the nullspace of CM, we obtain κ corresponding vectors $N_i = D_i C \in \mathsf{K}[x]^{n+q}$ in the nullspace of M.

Algorithm $\mathsf{Nullspace}_{2n}(M)$

> *Input:* $M \in \mathsf{K}[x]^{(n+q)\times n}$ of degree d,
> $\qquad\qquad$ M has full column-rank and $1 \leq q \leq n$.
> *Output:* q "small" linearly independent polynomial vectors in the nullspace of M.
>
> $M := QM$ for a random $Q \in \mathsf{K}^{(n+q)\times(n+q)}$;
> **if** $\det M_{1..n,1..n}(x_0) = 0$ for x_0 random in K **then fail**;
> $I = \{\}$;
> $p := q$;
> **while** $\#I < q$
> $\qquad$ (a) $\{i_1, \ldots, i_p\} := \{n+1, \ldots, n+q\} \setminus I$;
> $\qquad$ (b) $\delta := 2nd/p$;
> $\qquad$ (c)   construct $C \in \mathsf{K}^{(n+p)\times(n+q)}$ with $C_{i,i} := 1$, $1 \leq i \leq n$, $C_{n+j,i_j} := 1$, $1 \leq j \leq p$,
> $\qquad\qquad$ and $C_{i,j} := 0$ otherwise;
> $\qquad$ (d) $\bar{M} := CM \in \mathsf{K}[x]^{(n+p)\times n}$;
> $\qquad$ (e) $\{\kappa, \{D_i, 1 \leq i \leq \kappa\}\} := \mathsf{Nullspace\ minimal\ vectors}(\bar{M}, \delta)$;
> $\qquad\qquad$ $N_i^{(\delta)} = D_i C$, $1 \leq i \leq \kappa$;
> $\qquad$ (f) $N^{(\delta)} :=$ the $\kappa \times (n+q)$ matrix whose rows are the $N_i^{(\delta)}$'s;
> $\qquad$ (g) $J := \kappa$ column indices greater than $n+1$ such that $N_{1..\kappa,J}$ is non-singular;
> $\qquad$ (h) $I := I \cup J$, $p := p - \kappa$;
> $\qquad$ (i) $N := [N^T \ (N^{(\delta)})^T]^T$;   /* update the nullspace */
> $N := NQ$;
> **return** N_i, $1 \leq i \leq q$. $\square$

Algorithm $\mathsf{Nullspace}_{2n}$ is proven in Proposition 7.1 below. Let us first give the general idea. For computing the whole nullspace, the algorithm generates a sequence of decreasing

dimensions p at Step (h). Using the observation made previously, each time the algorithm passes through the "while loop" the dimension is divided by at least two, hence at most $O(\log_2 q)$ stages are necessary. This corresponds to $O(\log_2 q)$ calls to Nullspace minimal vectors with input CM. Each time the dimension is decreased, the degree threshold is increased in the same proportion at Step (b), we preserve the invariant

$$p\delta/(nd) = 2. \tag{19}$$

The latter identity will be used for applying the cost estimate (18) of Corollary 6.5.

The proof of Proposition 7.1 will check that q vectors in the nullspace are actually computed. In addition, the algorithm has to ensure their linear independency. The latter is done on the fly, and will first rely on the initial conditioning with Q for working with a top $n \times n$ non-singular submatrix. The vectors for updating the nullspace are computed at Step (e) and Step (f) in the nullspace of $M_{\bar{I},1..n}$, with $\bar{I} = \{1, 2, \ldots, n, i_1, i_2, \ldots, i_p\}$. This is done through the construction of the compression matrix C at Step (c) which selects the corresponding rows of M. The choice of the indices $\{i_1, i_2, \ldots, i_p\}$ at Step (a), complements the index choices at Step (g) that are kept in I at Step (h) for previous stages, and will provide the linear independency by construction. Another perhaps simpler strategy for ensuring independency could be based on randomization.

Our approach is "greedy", all vectors of degree under the threshold δ are kept. It is unclear how using a formal "divide and conquer" would make a difference.

Proposition 7.1 *Let $M \in \mathsf{K}[x]^{(n+q)\times n}$ with $1 \leq q \leq n$ be of full column-rank. Algorithm* Nullspace$_{2n}$ *computes q linearly independent polynomial vectors in the nullspace of M. If M has degree d then the sum of the degrees of the output vectors is less than $nd\lceil \log_2 q \rceil$. The algorithm is randomized, it either fails or returns correct values (Las Vegas fashion).*

Proof. The initial multiplication by Q and the corresponding failure test ensure that the top $n \times n$ matrix of M is invertible when the algorithm enters the "while loop" (if the algorithms fails then M probably has rank less than n). At Step (f), κ vectors in the nullspace of M are computed, indeed, $D_i\bar{M} = D_iCM = 0$ directly gives $N_i^{(\delta)}M = D_iCM = 0$. The number of elements of I is increased by κ at Step (h), hence is equal to the current total number of computed vectors. Since $\kappa \leq q - \#I$, if the algorithm terminates then exactly q nullspace vectors are obtained. In addition we have already seen that κ is at least $\lceil p/2 \rceil$, therefore, if we denote by p_{new} the new value of p at Step (h), we have $p_{\text{new}} \leq \lfloor p/2 \rfloor$, which means that the algorithm terminates after having passed through the "while loop" at most $\lceil \log_2 q \rceil$ times.

Algorithm Nullspace minimal vectors returns κ linearly independent row vectors $D_i \in \mathsf{K}[x]^{n+p}$ at Step (e). Let D be the $\kappa \times (n+p)$ matrix whose rows are the D_i's. We respectively denote the kth column of $N^{(\delta)}$ and D, by $N_{\cdot,k}^{(\delta)}$ and $D_{\cdot,k}$. The construction of C leads to:

$$N_{\cdot,i_j}^{(\delta)} = D_{\cdot,n+j}, \text{ if } 1 \leq j \leq p, \tag{20a}$$

$$N_{\cdot,k}^{(\delta)} = 0, \text{ otherwise.} \tag{20b}$$

Since the top $n \times n$ matrix of M, and consequently the one of $\bar{M}$, is non-singular, κ linearly independent columns may be found among the last p columns of D. Therefore, from (20a), κ linearly independent columns J may be found among the columns $i_1, \ldots, i_p$ of $N^{(\delta)}$. This shows that Step (g) is valid. In addition, at subsequent stages, from Step (a) and (20b), the

non-zero columns involved between $n+1$ and q will be outside J, the corresponding nullspace vectors will thus be linearly independent from $N_i^{(\delta)}$, $1 \le i \le \kappa$. At each stage the $N_i^{(\delta)}$'s are linearly independent, and are independent from those computed subsequently, hence we have proven that the algorithm returns q linearly independent nullspace vectors.

Each of the times the algorithm passes through the "while loop", the sum of the degrees of the computed vectors is bounded by the sum nd of the Kronecker indices. Indeed, these vectors are minimal for the nullspace of the submatrix $\bar{M}$. Hence the sum of the degrees in output is less than $nd\lceil \log_2 q \rceil$.   □

The computed vectors D_i's are minimal in the nullspace of CM but the minimality is not preserved in general for the vectors N_i's in the nullspace of M. The output basis for the nullspace as $\mathsf{K}(x)$-vector space may not be a basis for the $\mathsf{K}[x]$-module. However, Proposition 7.1 shows that if the sum of the Kronecker indices is nd (the maximum possible), then the sum of the computed degrees is only within $\lceil \log_2 q \rceil$ times the optimum. We notice also that the vectors computed at the first stage are minimal vectors by Proposition 6.4, hence the algorithm reaches the optimum for a generic matrix M (the whole nullspace is computed with $p = q$). It would be interesting to study the loss of minimality compared to the Kronecker indices in the general case.

We also remark that the algorithm could be slightly modified for computing a row-reduced nullspace matrix N. The intermediate bases matrices $D \in \mathsf{K}[x]^{\kappa \times (n+p)}$ whose rows are the D_i's are row-reduced by Proposition 6.4. By Lemma 5.1 the dominant degrees are in the last p columns. The column index selection of Step (g) may be specialized for choosing indices corresponding to dominant degrees. From there, the proof of Proposition 7.1 for establishing that the computed vectors are independent may be extended to the fact that the output matrix N is row-reduced. This could be certified at the end of the Algorithm $\mathsf{Nullspace_{2n}}$ as done at Step (i) of Algorithm $\mathsf{Nullspace\ minimal\ vectors}$.

Corollary 7.2 *Let $M \in \mathsf{K}[x]^{(n+q) \times n}$ be of full column-rank and degree d with $1 \le q \le n$, q polynomial vectors whose degree sum is less than $nd\lceil \log_2 q \rceil$ can be computed in*

$$O((\mathsf{MM}(n,d)\log(nd) + n^2\mathsf{B}(d)\log n + n\mathsf{M}(nd))\log q) \qquad (21)$$

operations in K by a randomized Las Vegas (certified) algorithm.

Proof. We study the cost of Algorithm $\mathsf{Nullspace_{2n}}$. The conditioning with the matrix Q and the failure test use at most $O(\mathsf{MM}(n,d))$ operations. We claim that the dominating cost is the body of the loop is the call to Algorithm $\mathsf{Nullspace\ minimal\ vectors}$. Since $O(\log q)$ calls are sufficient, and since $p\delta/(nd) = 2$, (21) is a consequence of (18) in Corollary 6.5. Step (d) is the extraction of a submatrix. The computations $N_i^{(\delta)} = D_i C \in \mathsf{K}[x]^{n+q}$, for $1 \le i \le \kappa$, can be done in $O(n^2d)$ since the degree sum of the $N_i^{(\delta)}$'s is less than nd. The choice of κ column indices at Step (g) can be made in $O(n^\omega + n^2d)$ operations.   □

7.2 General case

We now work with a general matrix $M \in \mathsf{K}[x]^{m \times n}$ of degree d. We compute the rank r of M and $m - r$ linearly independent and "small" polynomial vectors in the nullspace. Our strategy first uses Monte Carlo techniques for computing a value $r_o \le r$, equal to r with high probability.

Lemma 7.3 *Let M be in $\mathsf{K}[x]^{m \times n}$ of degree d. A matrix $\tilde{M} \in \mathsf{K}[x]^{m \times r_o}$ of degree d and full column-rank with $r_o \leq r$, such that with high probability $r_o = r$ and its nullspace is equal to the nullspace of M, can be computed in $O(nm\mathsf{MM}(r, d)/r^2)$ operations in K by a randomized Monte Carlo (non-certified) algorithm.*

Proof. The matrix M can be evaluated at a random value x_0 in K in $O(mnd)$ operations. With high probability the rank is preserved. Then the rank $r_o \leq r$ after evaluation can be computed over K in $O(nmr^{\omega-2})$ operations (see [22] and [30, Chapter 3]). We compute $\tilde{M} = MR$ for R a random $n \times r_o$ matrix over K in $O(nm\mathsf{MM}(r, d)/r^2)$. □

Lemma 7.3 reduces the problem to the full column-rank case. We then apply the results of previous sections for computing $m - r_o$ candidate independent vectors in the nullspace of $\tilde{M}$. We finally test by multiplication whether the $m - r_o$ vectors are actually in the nullspace of M. A positive answer implies that $r \leq r_o$, therefore certifies that $r = r_o$, and that a correct nullspace representation has been constructed.

The case $m \leq 2r_o$ has been treated in Section 7.1. It remains to handle in particular the situation $m \gg r_o$. The sum of the Kronecker indices is at most $r_o d$, hence at most r_o vectors may have degrees greater than d. For $m > 2r_o$, we apply the technique of successive row indices selection of Section 7.1 for computing $m - 2r_o$ independent vectors of degrees less than d, and will terminate by computing r_o vectors of possibly higher degrees using the case $m = 2r_o$.

Algorithm Nullspace(M)

> *Input:* $M \in \mathsf{K}[x]^{m \times n}$ of degree d.
> *Output:* $r = \operatorname{rank} M$,
> $m - r$ "small" linearly polynomial vectors in the nullspace of M.

> (a)   compute r_o and $\tilde{M} = MR \in \mathsf{K}[x]^{m \times r_o}$ using Lemma 7.3;
>   **if** $m = r_o$ **then return** m and $\{\}$;
> $q := \lceil (m - 2r_o)/r_o \rceil$;
> (b)   randomly ensure that the top $r_o \times r_o$ submatrix of $\tilde{M}$ is non-singular or **fail**;
> (c) $\{N_i, 1 \leq i \leq m - 2r_o\} := $ Nullspace minimal vectors($\tilde{M}^{(k)}, d$), $1 \leq k \leq q$;
> (d) $\{N_i', 1 \leq i \leq \min\{m, 2r_o\} - r_o\} := $ Nullspace$_{2n}$($\tilde{M}^{(q+1)}$);
> N in $\mathsf{K}[x]^{(m-r_o) \times m}$ the matrix whose rows are the N_i's and the N_i''s;
> (e)   **if** $NM \neq 0$ **then fail**;
>   else **return** r_o and N_i, $1 \leq i \leq m - r_o$.   □

For the first $m - 2r_o$ vectors of degrees less than d we work in $q = \lceil (m - 2r_o)/r_o \rceil$ stages, and successively consider submatrices $\tilde{M}^{(1)}, \ldots, \tilde{M}^{(q)} \in \mathsf{K}[x]^{\iota \times r_o}$ of $\tilde{M}$, with $2r_o < \iota \leq 3r_o$. More precisely, $\tilde{M}^{(k)} \in \mathsf{K}[x]^{3r_o \times r_o}$ for $1 \leq k \leq q - 1$, and $\tilde{M}^{(q)} \in \mathsf{K}[x]^{(m-(q-1)r_o) \times r_o}$. Like in Algorithm Nullspace$_{2n}$ we always ensure by randomization that the top $r_o \times r_o$ submatrix is non-singular. Each of the $\tilde{M}^{(k)}$'s has at least $\iota - 2r_o$ nullspace vectors of degree at most d. Therefore, in at most q calls to Algorithm Nullspace minimal vectors (see also Remark 6.6) on the $\tilde{M}^{(k)}$'s with $\delta = d$ we compute $(q-1)r_o + (m - (q-1)r_o - 2r_o) = m - 2r_o$ nullspace vectors of degrees less than d. This is exactly in q calls if exactly $\iota - 2r_o$ nullspace vectors have degree less than d at each call, or if exactly $\iota - 2r_o$ vectors are kept. Otherwise, a greedy strategy as in previous section may need less calls. Without giving the details here,

we remark that *ad hoc* successive index choices for constructing the submatrices $\tilde{M}^{(k)}$'s will lead to $m - 2r_o$ linearly independent vectors (see Proposition 7.1 and its proof). Once this is done, we are led to a remaining $\min\{m, 2r_o\} \times r_o$ matrix $\tilde{M}^{(q+1)}$ whose nullspace can be computed by Algorithm Nullspace$_{2n}$. If $m \leq 2r_o$ then $\tilde{M}^{(q+1)}$ is simply the input matrix M. Again, we ensure independency by *ad hoc* row index choices.

   We do not further detail the proof of the algorithm which relies on similar techniques than those used for the proof of Proposition 7.1. The $m - r_o$ computed vectors at Step (c) and Step (d) are in the nullspaces of full rank submatrices with r_o columns of $\tilde{M}$, hence are in the nullspace of $\tilde{M}$. The check (e) ensures that they are in the nullspace of M.

Theorem 7.4 *Let $M \in \mathsf{K}[x]^{m \times n}$ be of degree d. The rank r of M and $m - r$ linearly independent polynomial vectors in the nullspace of M can be computed in*

$$O(nm\mathsf{MM}(r,d)/r^2 + (m/r + \log r)(\mathsf{MM}(r,d)\log(rd) + r^2\mathsf{B}(d)\log r + r\mathsf{M}(rd))) \qquad (22)$$

hence $O\tilde{\ }(nmr^{\omega-2}d)$ operations in K by a randomized Las Vegas (certified) algorithm. The degree sum of the computed nullspace vectors is less than $rd\lceil \log_2 r \rceil + (m - 2r)d$.

Proof. The cost for computing $\tilde{M}$ using Lemma 7.3 is bounded by $O(nm\mathsf{MM}(r,d)/r^2)$. The top $r_o \times r_o$ matrix is made non-singular by pre-multiplication by a random constant matrix $Q \in \mathsf{K}^{m \times m}$ (see Algorithm Nullspace$_{2n}$) in $O(\mathsf{MM}(n,d))$. Since only the first r_o rows of M need to be modified, the first r_o rows of Q are randomly chosen in K, and the last $m - r_o$ are fixed to $[0 \ \ I_{m-r_o}]^T$. The cost of the multiplication by Q is $O((m/r)(\mathsf{MM}(r,d)))$. At Step (c) we run Algorithm Nullspace minimal vectors $q = O(m/r)$ times on matrices of dimensions $O(r)$. Each call has cost (18) with $n = r$. Then at Step (d) one call to Algorithm Nullspace$_{2n}$ has cost (21) with m and n in $O(r)$. The two latter costs give the factor of $O(m/r + \log r)$ in (22). The final check at Step (e) is done in $q + 1$ multiplications using the special form of the intermediate results of Step (c) and Step (d). For one output of Nullspace minimal vectors at Step (c), the check is done in $O(n/r)\mathsf{MM}(r,d)$ operations, therefore q calls lead to a check in $O((nm)\mathsf{MM}(r,d)/r^2)$. As done in Corollary 6.5 for computing λ, the check involving the output of Algorithm Nullspace$_{2n}$ is done by splitting the large degrees in the N_i''s, and by forming an $(\min\{m, 2r_o\} - r_o) \times m$ matrix of degree d, the multiplication by M is done in $O((nm)\mathsf{MM}(r,d)/r^2)$ operations.

   The degree bound follows from the fact that the minimal vectors computations of Step (c) lead to $m - 2r$ vectors of degrees at most d. Proposition 7.1 gives the term $rd\lceil \log_2 r \rceil$ for the degree sum bound for Step (d) outputs. $\square$

   For $m \leq 2r$ we have already commented after Proposition 7.1 the quality of the degree sum bound $rd\lceil \log_2 r \rceil$. For $m \gg r$, since the sum of the Kronecker indices is no more than rd, we see that the bound we propose in Theorem 7.4 is within a factor asymptotically m/r from the optimal. A more accurate "tri-parameter" analysis—with respect to n, m and r—remains to be done. It may first require slight modifications of the σ-basis algorithm of [1, 15] that we use for computing minimal vectors, and a corresponding cost analysis especially with respect to r when $m \gg r$.

   We conclude with a simplified expression of the cost for $n = m$ and using $r \leq n$. The polynomial matrix multiplication has cost given by (4) or (5), and we take $\mathsf{M}(d) = O(d \log d \log \log d)$ [9].

Corollary 7.5 *The rank r of $M \in \mathsf{K}[x]^{n \times n}$ of degree d, and $m - r$ linearly independent polynomial vectors in the nullspace of M can be computed in*

$$O(\mathsf{MM}(n, d)(\log^2 n + \log n \log d) + n^2 \mathsf{B}(d) \log^2 n \log \log n)$$

hence $\tilde{O}(n^\omega d)$ operations in K by a randomized Las Vegas (certified) algorithm.

Remark 7.6 We did not detail the probability analysis. Random values in K occur for: the choice of P concerning the denominator matrix S and the right fraction degree bound in Proposition 4.2; the choice of Q in Lemma 5.1 for the degree dominance of the last columns in bases, and as linear independence conditioning in the different algorithms; the point x_0 in Algorithms Minimal nullspace vectors and Nullspace$_{2n}$; the random conditioning of M into $\tilde{M}$ in Lemma 7.3. Our algorithms are deterministic if random values are replaced by symbolic variables. For a given input matrix M, the algorithm succeeds if the random values do not form a zero of a fixed polynomial over K in the latter variables. This happens only with small probability if the random values are chosen from a subset of K of appropriate cardinality [12, 35, 29] (see also our comments in Introduction).

Concluding remarks

We compute a $\mathsf{K}(x)$-nullspace basis of an input matrix over $\mathsf{K}[x]$ as the union of few minimal $\mathsf{K}[x]$-basis of submatrices of the input matrix. It remains to compute a minimal basis with an analogous complexity estimate. A possible direction of work here is to ensure the irreducibility of the output basis either on the fly or *a posteriori*.

Subsequent work may also concern the applicability of our compression / uncompression scheme to other problems such as questions about matrix approximants or block structured matrices.

Computing a nullspace basis is added to the recent list of problems that can be solved in about the same number of operations as for multiplying two matrix polynomials. We hope that this will help in making progress for the characteristic polynomial [18, 21], and for (non-generic) matrix inversion [16].

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. Research report LIP 2002-1, Laboratoire LIP, ENS Lyon, France, 2002. To appear in *Journal of Symbolic Computation*.

[3] T. Beelen and P.M. Van Dooren. An improved algorithm for the computation of Kronecker's canonical form of a singular pencil. *Linear Algebra and its Applications*, 105:9–65, 1988.

[4] T. Beelen and P.M. Van Dooren. A pencil approach for embedding a polynomial matrix into a unimodular matrix. *SIAM J. Matrix Anal. Appl.*, 9(1):77–89, Jan. 1988.

[5] R.R. Bitmead, S.Y. Kung, B.D.O. Anderson, and T. Kailath. Greatest common divisors via generalized Sylvester and Bezout matrices. *IEEE Trans. Automat. Control.*, 23(6):1043–1047, 1978.

[6] A. Bostan. *Algorithmique efficace pour des opérations de base en calcul formel.* PhD thesis, École Polytechnique, Palaiseau, France, December 2003.

[7] A. Bostan and E. Schost. Polynomial evaluation and interpolation on special sets of points. Preprint 2004-02, Laboratoire STIX, École Polytechnique, Palaiseau, France, January 2004.

[8] P. Bürgisser, M. Clausen, and M.A. Shokrollahi. *Algebraic Complexity Theory.* Volume 315, Grundlehren der mathematischen Wissenschaften. Springer-Verlag, 1997.

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

[10] L. Chen, W. Eberly, E. Kaltofen, B.D. Saunders, W.J. Turner, and G. Villard. Efficient matrix preconditioners for black box linear algebra. *Linear Algebra and its Applications*, 343-344:119–146, 2002.

[11] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. of Symbolic Computations*, 9(3):251–280, 1990.

[12] R. A. DeMillo and R. J. Lipton. A probabilistic remark on algebraic program testing. *Information Process. Letters*, 7(4):193–195, 1978.

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

[14] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra.* Cambridge University Press, 1999.

[15] 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.

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

[17] T. Kailath. *Linear systems.* Prentice Hall, 1980.

[18] 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.

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

[20] E. Kaltofen and B.D. Saunders. On Wiedemann's method of solving sparse linear systems. In *Proc. AAECC-9*, LNCS 539, Springer Verlag, pages 29–38, 1991.

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

[22] W. Keller-Gehrig. Fast algorithms for the characteristic polynomial. *Theoretical Computer Science*, 36:309–317, 1985.

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

[24] P. Misra, P. Van Dooren, and A. Varga. Computation of structural invariants of generalized state-space systems. *Automatica*, 30:1921–1936, 1994.

[25] T. Mulders and A. Storjohann. On lattice reduction for polynomial matrices. *Journal of Symbolic Computation*, 35(4):377–401, 2003.

[26] C. Oară and P. Van Dooren. An improved algorithm for the computation of structural invariants of a system pencil and related geometric aspects. *Systems and Control Letters*, 30:38–48, 1997.

[27] V.M. Popov. Some properties of control systems with irreducible matrix transfer functions. In *Lecture Notes in Mathematics*, volume 144, pages 169–180. Springer Verlag, Berlin, 1970.

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

[29] J.T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27:701–717, 1980.

[30] A. Storjohann. *Algorithms for Matrix Canonical Forms*. PhD thesis, Institut für Wissenschaftliches Rechnen, ETH-Zentrum, Zurich, Switzerland, November 2000.

[31] A. Storjohann. High-Order Lifting (Extended Abstract). In *Proc. International Symposium on Symbolic and Algebraic Computation, Lille, France*, pages 246–254. ACM Press, July 2002.

[32] 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.

[33] G. Villard. A study of Coppersmith's block Wiedemann algorithm using matrix polynomials, Feb. 1997. RR 975-I-M IMAG Grenoble, France.

[34] G. Villard. Further analysis of Coppersmith's block Wiedemann algorithm for the solution of sparse linear systems. In *International Symposium on Symbolic and Algebraic Computation, Maui, Hawaii, USA*, pages 32–39. ACM Press, July 1997.

[35] R.E. Zippel. Probabilistic algorithms for sparse polynomials. In *Proc. EUROSAM*, volume 72 of *Lect. Notes in Comput. Sci.*, pages 216–226. Springer Verlag, 1979.