

Fast computation of shifted Popov forms of polynomial matrices via systems of modular polynomial equations

Vincent Neiger

► **To cite this version:**

Vincent Neiger. Fast computation of shifted Popov forms of polynomial matrices via systems of modular polynomial equations. 41st International Symposium on Symbolic and Algebraic Computation, Jul 2016, Waterloo, ON, Canada. Proceedings of the 41st International Symposium on Symbolic and Algebraic Computation. <10.1145/2930889.2930936>. <hal-01266014v2>

HAL Id: hal-01266014

<https://hal.inria.fr/hal-01266014v2>

Submitted on 12 May 2016

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

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

Fast Computation of Shifted Popov Forms of Polynomial Matrices via Systems of Modular Polynomial Equations

Vincent Neiger
ENS de Lyon
Laboratoire LIP, CNRS, Inria, UCBL, U. Lyon
vincent.neiger@ens-lyon.fr

ABSTRACT

We give a Las Vegas algorithm which computes the shifted Popov form of an $m \times m$ nonsingular polynomial matrix of degree d in expected $\tilde{O}(m^\omega d)$ field operations, where ω is the exponent of matrix multiplication and $\tilde{O}(\cdot)$ indicates that logarithmic factors are omitted. This is the first algorithm in $\tilde{O}(m^\omega d)$ for shifted row reduction with arbitrary shifts.

Using partial linearization, we reduce the problem to the case $d \leq \lceil \sigma/m \rceil$ where σ is the generic determinant bound, with σ/m bounded from above by both the average row degree and the average column degree of the matrix. The cost above becomes $\tilde{O}(m^\omega \lceil \sigma/m \rceil)$, improving upon the cost of the fastest previously known algorithm for row reduction, which is deterministic.

Our algorithm first builds a system of modular equations whose solution set is the row space of the input matrix, and then finds the basis in shifted Popov form of this set. We give a deterministic algorithm for this second step supporting arbitrary moduli in $\tilde{O}(m^{\omega-1}\sigma)$ field operations, where m is the number of unknowns and σ is the sum of the degrees of the moduli. This extends previous results with the same cost bound in the specific cases of order basis computation and M-Padé approximation, in which the moduli are products of known linear factors.

Keywords

Shifted Popov form; polynomial matrices; row reduction; Hermite form; system of modular equations.

1. INTRODUCTION

In this paper, we consider two problems of linear algebra over the ring $\mathbb{K}[X]$ of univariate polynomials, for some field \mathbb{K} : computing the shifted Popov form of a matrix, and solving systems of modular equations.

1.1 Shifted Popov form

A polynomial matrix \mathbf{P} is row reduced [22, Section 6.3.2] if its rows have some type of minimal degree (we give precise

definitions below). Besides, if \mathbf{P} satisfies an additional normalization property, then it is said to be in Popov form [22, Section 6.7.2]. Given a matrix \mathbf{A} , the efficient computation of a (row) reduced form of \mathbf{A} and of the Popov form of \mathbf{A} has received a lot of attention recently [14, 28, 16].

In many applications one rather considers the degrees of the rows of \mathbf{P} shifted by some integers which specify degree weights on the columns of \mathbf{P} , for example in list-decoding algorithms [2, 7], robust Private Information Retrieval [12], and more generally in polynomial versions of the Copper-Smith method [9, 10]. A well-known specific shifted Popov form is the Hermite form; there has been recent progress on its fast computation [17, 15, 35]. The case of an arbitrary shift has been studied in [6].

For a shift $\mathbf{s} = (s_1, \dots, s_n) \in \mathbb{Z}^n$, the \mathbf{s} -degree of $\mathbf{p} = [p_1, \dots, p_n] \in \mathbb{K}[X]^{1 \times n}$ is $\max_{1 \leq j \leq n} (\deg(p_j) + s_j)$; the \mathbf{s} -row degree of $\mathbf{P} \in \mathbb{K}[X]^{m \times n}$ is $\text{rdeg}_{\mathbf{s}}(\mathbf{P}) = (d_1, \dots, d_m)$ with d_i the \mathbf{s} -degree of the i -th row of \mathbf{P} . Then, the \mathbf{s} -leading matrix of $\mathbf{P} = [p_{i,j}]_{i,j}$ is the matrix $\text{lm}_{\mathbf{s}}(\mathbf{P}) \in \mathbb{K}^{m \times n}$ whose entry (i, j) is the coefficient of degree $d_i - s_j$ of $p_{i,j}$.

Now, we assume that $m \leq n$ and \mathbf{P} has full rank. Then, \mathbf{P} is said to be \mathbf{s} -reduced [22, 6] if $\text{lm}_{\mathbf{s}}(\mathbf{P})$ has full rank. For a full rank $\mathbf{A} \in \mathbb{K}[X]^{m \times n}$, an \mathbf{s} -reduced form of \mathbf{A} is an \mathbf{s} -reduced matrix \mathbf{P} whose row space is the same as that of \mathbf{A} ; by row space we mean the $\mathbb{K}[X]$ -module generated by the rows of the matrix. Equivalently, \mathbf{P} is left-unimodularly equivalent to \mathbf{A} and the tuple $\text{rdeg}_{\mathbf{s}}(\mathbf{P})$ sorted in nondecreasing order is lexicographically minimal among the \mathbf{s} -row degrees of all matrices left-unimodularly equivalent to \mathbf{A} .

Specific \mathbf{s} -reduced matrices are those in \mathbf{s} -Popov form [22, 5, 6], as defined below. One interesting property is that the \mathbf{s} -Popov form is canonical: there is a unique \mathbf{s} -reduced form of \mathbf{A} which is in \mathbf{s} -Popov form, called the \mathbf{s} -Popov form of \mathbf{A} .

DEFINITION 1.1 (PIVOT). Let $\mathbf{p} = [p_j]_j \in \mathbb{K}[X]^{1 \times n}$ be nonzero and let $\mathbf{s} \in \mathbb{Z}^n$. The \mathbf{s} -pivot index of \mathbf{p} is the largest index j such that $\text{rdeg}_{\mathbf{s}}(\mathbf{p}) = \deg(p_j) + s_j$. Then we call p_j and $\deg(p_j)$ the \mathbf{s} -pivot entry and the \mathbf{s} -pivot degree of \mathbf{p} .

We remark that adding a constant to the entries of \mathbf{s} does not change the notion of \mathbf{s} -pivot. For example, we will sometimes assume $\min(\mathbf{s}) = 0$ without loss of generality.

DEFINITION 1.2 (SHIFTED POPOV FORM). Let $m \leq n$, let $\mathbf{P} \in \mathbb{K}[X]^{m \times n}$ be full rank, and let $\mathbf{s} \in \mathbb{Z}^n$. Then, \mathbf{P} is said to be in \mathbf{s} -Popov form if the \mathbf{s} -pivot indices of its rows are strictly increasing, the corresponding \mathbf{s} -pivot entries are monic, and in each column of \mathbf{P} which contains a pivot the nonpivot entries have degree less than the pivot entry.

In this case, the \mathbf{s} -pivot degree of \mathbf{P} is $\boldsymbol{\delta} = (\delta_1, \dots, \delta_m) \in \mathbb{N}^m$, with δ_i the \mathbf{s} -pivot degree of the i -th row of \mathbf{P} .

Here, although we will encounter Popov forms of rectangular matrices in intermediate nullspace computations, our main focus is on computing shifted Popov forms of *square nonsingular matrices*. For the general case, studied in [6], a fast solution would require further developments. A square matrix in \mathbf{s} -Popov form has its \mathbf{s} -pivot entries on the diagonal, and its \mathbf{s} -pivot degree is the tuple of degrees of its diagonal entries and coincides with its column degree.

PROBLEM 1 (SHIFTED POPOV NORMAL FORM).

Input: *the base field \mathbb{K} , a nonsingular matrix $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$, a shift $\mathbf{s} \in \mathbb{Z}^m$.*

Output: *the \mathbf{s} -Popov form of \mathbf{A} .*

Two well-known specific cases are the Popov form [27, 22] for the *uniform* shift $\mathbf{s} = \mathbf{0}$, and the Hermite form [19, 22] for the shift $\mathbf{h} = (0, \delta, 2\delta, \dots, (m-1)\delta) \in \mathbb{N}^m$ with $\delta = m \deg(\mathbf{A})$ [6, Lemma 2.6]. For a broader perspective on shifted reduced forms, we refer the reader to [6].

For such problems involving $m \times m$ matrices of degree d , one often wishes to obtain a cost bound similar to that of polynomial matrix multiplication in the same dimensions: $\tilde{\mathcal{O}}(m^\omega d)$ operations in \mathbb{K} . Here, ω is so that we can multiply $m \times m$ matrices over a commutative ring in $\mathcal{O}(m^\omega)$ operations in that ring, the best known bound being $\omega < 2.38$ [11, 25]. For example, one can compute $\mathbf{0}$ -reduced [14, 16], $\mathbf{0}$ -Popov [28], and Hermite [15, 35] forms of $m \times m$ nonsingular matrices of degree d in $\tilde{\mathcal{O}}(m^\omega d)$ field operations.

Nevertheless, d may be significantly larger than the average degree of the entries of the matrix, in which case the cost $\tilde{\mathcal{O}}(m^\omega d)$ seems unsatisfactory. Recently, for the computation of order bases [30, 34], nullspace bases [36], interpolation bases [20, 21], and matrix inversion [37], fast algorithms do take into account some types of average degrees of the matrices rather than their degree. Here, in particular, we achieve a similar improvement for the computation of shifted Popov forms of a matrix.

Given $\mathbf{A} = [a_{i,j}]_{ij} \in \mathbb{K}[X]^{m \times m}$, we denote by $\sigma(\mathbf{A})$ the *generic bound for $\deg(\det(\mathbf{A}))$* [16, Section 6], that is,

$$\sigma(\mathbf{A}) = \max_{\pi \in S_m} \sum_{1 \leq i \leq m} \overline{\deg}(a_{i, \pi_i}) \quad (1)$$

where S_m is the set of permutations of $\{1, \dots, m\}$, and $\overline{\deg}(p)$ is defined over $\mathbb{K}[X]$ as $\overline{\deg}(0) = 0$ and $\overline{\deg}(p) = \deg(p)$ for $p \neq 0$. We have $\deg(\det(\mathbf{A})) \leq \sigma(\mathbf{A}) \leq m \deg(\mathbf{A})$, and $\sigma(\mathbf{A}) \leq \min(|\text{rdeg}(\mathbf{A})|, |\text{cdeg}(\mathbf{A})|)$ with $|\text{rdeg}(\mathbf{A})|$ and $|\text{cdeg}(\mathbf{A})|$ the sums of the row and column degrees of \mathbf{A} . We note that $\sigma(\mathbf{A})$ can be substantially smaller than $|\text{rdeg}(\mathbf{A})|$ and $|\text{cdeg}(\mathbf{A})|$, for example if \mathbf{A} has one row and one column of uniformly large degree and other entries of low degree.

THEOREM 1.3. *There is a Las Vegas randomized algorithm which solves Problem 1 in expected $\tilde{\mathcal{O}}(m^\omega [\sigma(\mathbf{A})/m]) \subseteq \tilde{\mathcal{O}}(m^\omega \deg(\mathbf{A}))$ field operations.*

The ceiling function indicates that the cost is $\tilde{\mathcal{O}}(m^\omega)$ when $\sigma(\mathbf{A})$ is small compared to m , in which case \mathbf{A} has mostly constant entries. Here we are mainly interested in the case $m \in \mathcal{O}(\sigma(\mathbf{A}))$: the cost bound may be written $\tilde{\mathcal{O}}(m^{\omega-1} \sigma(\mathbf{A}))$ and is both in $\tilde{\mathcal{O}}(m^{\omega-1} |\text{rdeg}(\mathbf{A})|)$ and $\tilde{\mathcal{O}}(m^{\omega-1} |\text{cdeg}(\mathbf{A})|)$.

Previous work on fast algorithms related to Problem 1 is summarized in Table 1. The fastest known algorithm for the

Ref.	Problem	Cost bound
[18]	Hermite form	$\tilde{\mathcal{O}}(m^4 d)$
[31]	Hermite form	$\tilde{\mathcal{O}}(m^{\omega+1} d)$
[33]	Popov & Hermite forms	$\tilde{\mathcal{O}}(m^{\omega+1} d + (md)^\omega)$
[1, 2]	weak Popov form	$\tilde{\mathcal{O}}(m^{\omega+1} d)$
[26]	Popov & Hermite forms	$\mathcal{O}(m^3 d^2)$
[14]	$\mathbf{0}$ -reduction	$\tilde{\mathcal{O}}(m^\omega d)$ *
[28]	Popov form of $\mathbf{0}$ -reduced	$\tilde{\mathcal{O}}(m^\omega d)$
[17]	Hermite form	$\tilde{\mathcal{O}}(m^\omega d)$ *
[16]	$\mathbf{0}$ -reduction	$\tilde{\mathcal{O}}(m^\omega d)$
[35]	Hermite form	$\tilde{\mathcal{O}}(m^\omega d)$
[16]+[28]	\mathbf{s} -Popov form for any \mathbf{s}	$\tilde{\mathcal{O}}(m^\omega (d + \mu))$
Here	\mathbf{s} -Popov form for any \mathbf{s}	$\tilde{\mathcal{O}}(m^\omega [\sigma(\mathbf{A})/m])$ *

Table 1: Fast algorithms for shifted reduction problems ($d = \deg(\mathbf{A})$; * = probabilistic; $\mu = \max(\mathbf{s}) - \min(\mathbf{s})$).

$\mathbf{0}$ -Popov form is deterministic and has cost $\tilde{\mathcal{O}}(m^\omega d)$ with $d = \deg(\mathbf{A})$; it first computes a $\mathbf{0}$ -reduced form of \mathbf{A} [16], and then its $\mathbf{0}$ -Popov form via normalization [28]. Obtaining the Hermite form in $\tilde{\mathcal{O}}(m^\omega d)$ was first achieved by a probabilistic algorithm in [15], and then deterministically in [35].

For an arbitrary \mathbf{s} , the algorithm in [6] is fraction-free and uses a number of operations that is, depending on \mathbf{s} , at least quintic in m and quadratic in $\deg(\mathbf{A})$.

When \mathbf{s} is not uniform there is a folklore solution based on the fact that \mathbf{Q} is in \mathbf{s} -Popov form if and only if \mathbf{QD} is in $\mathbf{0}$ -Popov form, with $\mathbf{D} = \text{diag}(X^{s_1}, \dots, X^{s_m})$ and assuming $\mathbf{s} \geq 0$. Then, this solution computes the $\mathbf{0}$ -Popov form \mathbf{P} of \mathbf{AD} using [16, 28] and returns \mathbf{PD}^{-1} . This approach uses $\tilde{\mathcal{O}}(m^\omega (d + \mu))$ operations where $\mu = \max(\mathbf{s}) - \min(\mathbf{s})$, which is not satisfactory when μ is large. For example, its cost for computing the Hermite form is $\tilde{\mathcal{O}}(m^{\omega+2} d)$. This is the worst case since one can assume without loss of generality that $\mu \in \mathcal{O}(m \deg(\det(\mathbf{A}))) \subseteq \mathcal{O}(m^2 d)$ [21, Appendix A].

Here we obtain, to the best of our knowledge, the best known cost bound $\tilde{\mathcal{O}}(m^\omega [\sigma(\mathbf{A})/m]) \subseteq \tilde{\mathcal{O}}(m^\omega d)$ for an arbitrary shift \mathbf{s} . This removes the dependency in μ , which means in some cases a speedup by a factor m^2 . Besides, this is also an improvement for both specific cases $\mathbf{s} = \mathbf{0}$ and $\mathbf{s} = \mathbf{h}$ when \mathbf{A} has unbalanced degrees.

One of the main difficulties in row reduction algorithms is to control the size of the manipulated matrices, that is, the number of coefficients from \mathbb{K} needed for their dense representation. A major issue when dealing with arbitrary shifts is that the size of an \mathbf{s} -reduced form of \mathbf{A} may be beyond our target cost. This is a further motivation for focusing on the computation of the \mathbf{s} -Popov form of \mathbf{A} : by definition, the sum of its column degrees is $\deg(\det(\mathbf{A}))$, and therefore its size is at most $m^2 + m \deg(\det(\mathbf{A}))$, independently of \mathbf{s} .

Consider for example $\mathbf{A} = \begin{bmatrix} \mathbf{B}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_2 \end{bmatrix}$ for any $\mathbf{0}$ -reduced \mathbf{B}_1 and \mathbf{B}_2 in $\mathbb{K}[X]^{m \times m}$. Then, taking $\mathbf{s} = (0, \dots, 0, d, \dots, d)$ with $d > 0$, $\begin{bmatrix} \mathbf{B}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_2 \end{bmatrix}$ is an \mathbf{s} -reduced form of \mathbf{A} for any $\mathbf{C} \in \mathbb{K}[X]^{m \times m}$ with $\deg(\mathbf{C}) \leq d$; for some \mathbf{C} it has size $\Theta(m^2 d)$, with d arbitrary large independently of $\deg(\mathbf{A})$.

Furthermore, the size of the unimodular transformation leading from \mathbf{A} to \mathbf{P} may be beyond the target cost, which is why fast algorithms for $\mathbf{0}$ -reduction and Hermite form do not directly perform unimodular transformations on \mathbf{A} to reduce the degrees of its entries. Instead, they proceed in two steps: first, they work on \mathbf{A} to find some equations which describe its row space, and then they find a basis of solutions to these equations in $\mathbf{0}$ -reduced form or Hermite form. We will follow a similar two-step strategy for an arbitrary shift.

It seems that some new ingredient is needed, since for both $\mathbf{s} = \mathbf{0}$ and $\mathbf{s} = \mathbf{h}$ the fastest algorithms use shift-specific properties at some point of the process: namely, the facts that a $\mathbf{0}$ -reduced form of \mathbf{A} has degree at most $\deg(\mathbf{A})$ and that the Hermite form of \mathbf{A} is triangular.

As in [17], we first compute the Smith form \mathbf{S} of \mathbf{A} and partial information on a right unimodular transformation \mathbf{V} ; this is where the probabilistic aspect comes from. This gives a description of the row space of \mathbf{A} as the set of row vectors $\mathbf{p} \in \mathbb{K}[X]^{1 \times m}$ such that $\mathbf{pV} = \mathbf{qS}$ for some $\mathbf{q} \in \mathbb{K}[X]^{1 \times m}$. Since \mathbf{S} is diagonal, this can be seen as a system of modular equations: the second step is the fast computation of a basis of solutions in \mathbf{s} -Popov form, which is our new ingredient.

1.2 Systems of modular equations

Hereafter, $\mathbb{K}[X]_{\neq 0}$ denotes the set of nonzero polynomials. We fix some moduli $\mathfrak{M} = (\mathbf{m}_1, \dots, \mathbf{m}_n) \in \mathbb{K}[X]_{\neq 0}^n$, and for $\mathbf{A}, \mathbf{B} \in \mathbb{K}[X]^{m \times n}$ we write $\mathbf{A} = \mathbf{B} \bmod \mathfrak{M}$ if there exists $\mathbf{Q} \in \mathbb{K}[X]^{m \times n}$ such that $\mathbf{A} = \mathbf{B} + \mathbf{Q} \text{diag}(\mathfrak{M})$. Given $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$ specifying the equations, we call *solution for* $(\mathfrak{M}, \mathbf{F})$ any $\mathbf{p} \in \mathbb{K}[X]^{1 \times m}$ such that $\mathbf{pF} = 0 \bmod \mathfrak{M}$.

The set of all such \mathbf{p} is a $\mathbb{K}[X]$ -submodule of $\mathbb{K}[X]^{1 \times m}$ which contains $\text{lcm}(\mathbf{m}_1, \dots, \mathbf{m}_n) \mathbb{K}[X]^{1 \times m}$, and is thus free of rank m [24, p. 146]. Then, we represent any basis of this module as the rows of a matrix $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$, called a *solution basis for* $(\mathfrak{M}, \mathbf{F})$. Here, for example for the application to Problem 1, we are interested in such bases that are \mathbf{s} -reduced, in which case \mathbf{P} is said to be an *s-minimal solution basis for* $(\mathfrak{M}, \mathbf{F})$. The unique such basis which is in \mathbf{s} -Popov form is called the *s-Popov solution basis for* $(\mathfrak{M}, \mathbf{F})$.

PROBLEM 2 (MINIMAL SOLUTION BASIS).

Input: the base field \mathbb{K} , moduli $\mathfrak{M} = (\mathbf{m}_1, \dots, \mathbf{m}_n) \in \mathbb{K}[X]_{\neq 0}^n$, a matrix $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$ such that $\deg(\mathbf{F}_{*,j}) < \deg(\mathbf{m}_j)$, a shift $\mathbf{s} \in \mathbb{Z}^m$.

Output: an \mathbf{s} -minimal solution basis for $(\mathfrak{M}, \mathbf{F})$.

Well-known specific cases of this problem are *Hermite-Padé approximation* with a single equation modulo some power of X , and *M-Padé approximation* [3, 32] with moduli that are products of known linear factors. Moreover, an *s-order basis for* \mathbf{F} and $(\sigma_1, \dots, \sigma_n)$ [34] is an \mathbf{s} -minimal solution basis for $(\mathfrak{M}, \mathbf{F})$ with $\mathfrak{M} = (X^{\sigma_1}, \dots, X^{\sigma_n})$.

An overview of fast algorithms for Problem 2 is given in Table 2. For M-Padé approximation, and thus in particular for order basis computation, there is an algorithm to compute the \mathbf{s} -Popov solution basis using $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ operations, with $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n)$ [21]. Here, for $n \in \mathcal{O}(m)$, we extend this result to arbitrary moduli.

THEOREM 1.4. *Assuming $n \in \mathcal{O}(m)$, there is a deterministic algorithm which solves Problem 2 using $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ field operations, with $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n)$, and returns the \mathbf{s} -Popov solution basis for $(\mathfrak{M}, \mathbf{F})$.*

We note that Problem 2 is a minimal interpolation basis problem [5, 20] when the so-called *multiplication matrix* \mathbf{M} is block diagonal with companion blocks. Indeed, \mathbf{p} is a solution for $(\mathfrak{M}, \mathbf{F})$ if and only if \mathbf{p} is an *interpolant for* (\mathbf{E}, \mathbf{M}) [20, Definition 1.1], where $\mathbf{E} \in \mathbb{K}^{m \times \sigma}$ is the concatenation of the coefficient vectors of the columns of \mathbf{F} and $\mathbf{M} \in \mathbb{K}^{\sigma \times \sigma}$ is $\text{diag}(\mathbf{M}_1, \dots, \mathbf{M}_n)$ with \mathbf{M}_j the companion

matrix associated with \mathbf{m}_j . In this context, the multiplication $\mathbf{p} \cdot \mathbf{E}$ defined by \mathbf{M} as in [5, 20] precisely corresponds to $\mathbf{pF} \bmod \mathfrak{M}$.

In particular, Theorem 1.4 follows from [20, Theorem 1.4] when $\sigma \in \mathcal{O}(m)$. If some of the moduli have small degree, we use this result for base cases of our recursive algorithm.

Ref.	Cost bound	Moduli	Particularities
[3, 32]	$\mathcal{O}(m^2\sigma^2)$	split	
[4]	$\mathcal{O}(m\sigma^2)$	$\mathbf{m}_j = X^{\sigma/n}$	partial basis
[4]	$\tilde{\mathcal{O}}(m^\omega\sigma)$	$\mathbf{m}_j = X^{\sigma/n}$	
[14]	$\tilde{\mathcal{O}}(m^\omega\sigma/n)$	$\mathbf{m}_j = X^{\sigma/n}$	
[30]	$\tilde{\mathcal{O}}(m^\omega \lceil \sigma/m \rceil)$	$\mathbf{m}_j = X^{\sigma/n}$	partial basis, $ \mathbf{s} \leq \sigma$
[34]	$\tilde{\mathcal{O}}(m^\omega \lceil \sigma/m \rceil)$	$\mathbf{m}_j = X^{\sigma/n}$	$ \mathbf{s} \leq \sigma$
[8]	$\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$, probabilistic	any	returns a single small degree solution
[20]	$\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$	split	$ \mathbf{s} \leq \sigma$
[20]	$\tilde{\mathcal{O}}(m\sigma^{\omega-1})$	any	\mathbf{s} -Popov, $\sigma \in \mathcal{O}(m)$
[21]	$\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$	split	\mathbf{s} -Popov
Here	$\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$	any	\mathbf{s} -Popov

Table 2: Fast algorithms for Problem 2 ($n \in \mathcal{O}(m)$); *partial basis* = returns small degree rows of an \mathbf{s} -minimal solution basis; *split* = product of known linear factors).

In the case of M-Padé approximation, knowing the moduli as products of linear factors leads to rewriting the problem as a minimal interpolation basis computation with \mathbf{M} in Jordan form [5, 20]. Since \mathbf{M} is upper triangular, one can then rely on recurrence relations to solve the problem iteratively [3, 32, 4, 5]. The fast algorithms in [4, 14, 34, 20, 21], beyond the techniques used to achieve efficiency, are essentially divide-and-conquer versions of this iterative solution and are thus based on the same recurrence relations.

However, for arbitrary moduli the matrix \mathbf{M} is not triangular and there is no such recurrence in general. Then, a natural idea is to relate solution bases to nullspace bases: Problem 2 asks to find \mathbf{P} such that there is some quotient \mathbf{Q} with $[\mathbf{P}|\mathbf{Q}]\mathbf{N} = \mathbf{0}$ for $\mathbf{N} = [\mathbf{F}^\top | -\text{diag}(\mathfrak{M})]^\top$. More precisely, $[\mathbf{P}|\mathbf{Q}]$ can be obtained as a \mathbf{u} -minimal nullspace basis of \mathbf{N} for the shift $\mathbf{u} = (\mathbf{s} - \min(\mathbf{s}), \mathbf{0}) \in \mathbb{N}^{m+n}$.

Using recent ingredients from [17, 21] outlined in the next paragraphs, the main remaining difficulty is to deal with this nullspace problem when $n = 1$. Here, we give a $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ algorithm to solve it using its specific properties: \mathbf{N} is the column $[\mathbf{F}^\top | \mathbf{m}_1]^\top$ with $\deg(\mathbf{F}) < \deg(\mathbf{m}_1) = \sigma$, and the last entry of \mathbf{u} is $\min(\mathbf{u})$. First, when $\max(\mathbf{u}) \in \mathcal{O}(\sigma)$ we show that $[\mathbf{P}|\mathbf{Q}]$ can be efficiently obtained as a submatrix of the \mathbf{u} -Popov order basis for \mathbf{N} and order $\mathcal{O}(\sigma)$. Then, when $\max(\mathbf{u})$ is large compared to σ and assuming \mathbf{u} is sorted non-decreasingly, \mathbf{P} has a lower block triangular shape. We show how this shape can be revealed, along with the \mathbf{s} -pivot degree of \mathbf{P} , using a divide-and-conquer approach which splits \mathbf{u} into two shifts of amplitude about $\max(\mathbf{u})/2$.

Then, for $n \geq 1$ we use a divide-and-conquer approach on n which is classical in such contexts: two solution bases $\mathbf{P}^{(1)}$ and $\mathbf{P}^{(2)}$ are computed recursively in shifted Popov form and are multiplied together to obtain the \mathbf{s} -minimal solution basis $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$ for $(\mathfrak{M}, \mathbf{F})$. However this product is usually not in \mathbf{s} -Popov form and may have size beyond our target cost. Thus, as in [21], instead of computing $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$, we use $\mathbf{P}^{(2)}$ and $\mathbf{P}^{(1)}$ to deduce the \mathbf{s} -pivot degree of \mathbf{P} .

In both recursions above, we focus on finding the \mathbf{s} -pivot degree of \mathbf{P} . Using ideas and results from [17, 21], we show that this knowledge about the degrees in \mathbf{P} allows us to complete the computation of \mathbf{P} within the target cost.

2. FAST COMPUTATION OF THE SHIFTED POPOV SOLUTION BASIS

Hereafter, we call *s-minimal degree* of $(\mathfrak{M}, \mathbf{F})$ the *s-pivot degree* δ of the *s-Popov* solution basis for $(\mathfrak{M}, \mathbf{F})$; δ coincides with the column degree of this basis. A central result for the cost analysis is that $|\delta| = \delta_1 + \dots + \delta_m$ is at most $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n)$. This is classical for M-Padé approximation [32, Theorem 4.1] and holds for minimal interpolation bases in general (see for example [20, Lemma 7.17]).

2.1 Solution bases from nullspace bases and fast algorithm for known minimal degree

This subsection summarizes and slightly extends results from [17, Section 3]. We first show that the *s-Popov* solution basis for $(\mathfrak{M}, \mathbf{F})$ is the principal $m \times m$ submatrix of the *u-Popov* nullspace basis of $[\mathbf{F}^\top | \text{diag}(\mathfrak{M})]^\top$ for some $\mathbf{u} \in \mathbb{Z}^{m+n}$.

LEMMA 2.1. *Let $\mathfrak{M} = (\mathbf{m}_1, \dots, \mathbf{m}_n) \in \mathbb{K}[X]_{\neq 0}^n$, $\mathbf{s} \in \mathbb{Z}^m$, $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$ with $\deg(\mathbf{F}_{*,j}) < \deg(\mathbf{m}_j)$, $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$, and $\mathbf{w} \in \mathbb{Z}^n$ be such that $\max(\mathbf{w}) \leq \min(\mathbf{s})$. Then, \mathbf{P} is the *s-Popov* solution basis for $(\mathfrak{M}, \mathbf{F})$ if and only if $[\mathbf{P}|\mathbf{Q}]$ is the *u-Popov* nullspace basis of $[\mathbf{F}^\top | \text{diag}(\mathfrak{M})]^\top$ for some $\mathbf{Q} \in \mathbb{K}[X]^{m \times n}$ and $\mathbf{u} = (\mathbf{s}, \mathbf{w}) \in \mathbb{Z}^{m+n}$. In this case, $\deg(\mathbf{Q}) < \deg(\mathbf{P})$ and $[\mathbf{P}|\mathbf{Q}]$ has *s-pivot* index $(1, 2, \dots, m)$.*

PROOF. Let $\mathbf{N} = [\mathbf{F}^\top | \text{diag}(\mathfrak{M})]^\top$. It is easily verified that \mathbf{P} is a solution basis for $(\mathfrak{M}, \mathbf{F})$ if and only if there is some $\mathbf{Q} \in \mathbb{K}[X]^{m \times n}$ such that $[\mathbf{P}|\mathbf{Q}]$ is a nullspace basis of \mathbf{N} .

Now, having $\deg(\mathbf{F}_{*,j}) < \deg(\mathbf{m}_j)$ implies that any $[\mathbf{p}|\mathbf{q}] \in \mathbb{K}[X]^{1 \times (m+n)}$ in the nullspace of \mathbf{N} satisfies $\deg(\mathbf{q}) < \deg(\mathbf{p})$, and since $\max(\mathbf{w}) \leq \min(\mathbf{s})$ we get $\text{rdeg}_{\mathbf{w}}(\mathbf{q}) < \text{rdeg}_{\mathbf{s}}(\mathbf{p})$. In particular, for any matrix $[\mathbf{P}|\mathbf{Q}] \in \mathbb{K}[X]^{m \times (m+n)}$ such that $[\mathbf{P}|\mathbf{Q}]\mathbf{N} = 0$, we have $\text{lm}_{\mathbf{u}}([\mathbf{P}|\mathbf{Q}]) = [\text{lm}_{\mathbf{s}}(\mathbf{P}) | \mathbf{0}]$. This implies that \mathbf{P} is in *s-Popov* form if and only if $[\mathbf{P}|\mathbf{Q}]$ is in *u-Popov* form with *s-pivot* index $(1, \dots, m)$. \square

We now show that, when we have *a priori* knowledge about the *s-pivot* entries of a *s-Popov* nullspace basis, it can be computed efficiently via an *s-Popov* order basis.

LEMMA 2.2. *Let $\mathbf{s} \in \mathbb{Z}^{m+n}$ and let $\mathbf{N} \in \mathbb{K}[X]^{(m+n) \times n}$ be of full rank. Let $\mathbf{B} \in \mathbb{K}[X]^{m \times (m+n)}$ be the *s-Popov* nullspace basis for \mathbf{N} , (π_1, \dots, π_m) be its *s-pivot* index, $(\delta_1, \dots, \delta_m)$ be its *s-pivot* degree, and $\delta \geq \deg(\mathbf{B})$ be a degree bound. Then, let $\mathbf{u} = (u_1, \dots, u_{m+n}) \in \mathbb{Z}_{\leq 0}^{m+n}$ with*

$$u_j = \begin{cases} -\delta - 1 & \text{if } j \notin \{\pi_1, \dots, \pi_m\}, \\ -\delta_i & \text{if } j = \pi_i. \end{cases}$$

Writing $(\sigma_1, \dots, \sigma_n)$ for the column degree of \mathbf{N} , let $\tau_j = \sigma_j + \delta + 1$ for $1 \leq j \leq n$ and let \mathbf{A} be the *u-Popov* order basis for \mathbf{N} and (τ_1, \dots, τ_n) . Then, \mathbf{B} is the submatrix of \mathbf{A} formed by its rows at indices $\{\pi_1, \dots, \pi_m\}$.

PROOF. First, \mathbf{B} is in *u-Popov* form with $\text{rdeg}_{\mathbf{u}}(\mathbf{B}) = \mathbf{0}$. Define $\mathbf{C} \in \mathbb{K}[X]^{(m+n) \times (m+n)}$ whose *i*-th row is $\mathbf{B}_{j,*}$ if $i = \pi_j$ and $\mathbf{A}_{i,*}$ if $i \notin \{\pi_1, \dots, \pi_m\}$: we want to prove $\mathbf{C} = \mathbf{A}$.

Let $\mathbf{p} = [p_j]_j \in \mathbb{K}[X]^{1 \times (m+n)}$ be a row of \mathbf{A} , and assume $\text{rdeg}_{\mathbf{u}}(\mathbf{p}) < 0$. This means $\deg(p_j) < -u_j$ for all j , so that $\deg(\mathbf{p}) < \max(-\mathbf{u}) = \delta + 1$. Then, for all $1 \leq j \leq n$ we have $\deg(\mathbf{p}\mathbf{N}_{*,j}) < \sigma_j + \delta + 1 = \tau_j$, and from $\mathbf{p}\mathbf{N}_{*,j} = 0 \pmod{X^{\tau_j}}$ we obtain $\mathbf{p}\mathbf{N}_{*,j} = 0$, which is absurd by minimality of \mathbf{B} . As a result, $\text{rdeg}_{\mathbf{u}}(\mathbf{A}) \geq \mathbf{0} = \text{rdeg}_{\mathbf{u}}(\mathbf{B})$ componentwise.

Besides, $\mathbf{C}\mathbf{F} = \mathbf{0} \pmod{(X^{\tau_1}, \dots, X^{\tau_n})}$ and since \mathbf{C} has its *u-pivot* entries on the diagonal, it is *u-reduced*: by minimality of \mathbf{A} , we obtain $\text{rdeg}_{\mathbf{u}}(\mathbf{A}) = \text{rdeg}_{\mathbf{u}}(\mathbf{C})$. Then, it is easily verified that \mathbf{C} is in *u-Popov* form, hence $\mathbf{C} = \mathbf{A}$. \square

In particular, computing the *s-Popov* nullspace basis \mathbf{B} , when its *s-pivot* index, its *s-pivot* degree, and $\delta \geq \deg(\mathbf{B})$ are known, can be done in $\tilde{\mathcal{O}}(m^{\omega-1}(\sigma + n\delta))$ with $\sigma = \sigma_1 + \dots + \sigma_n$ using the order basis algorithm in [21].

As for Problem 2, with Lemma 2.1 this gives an algorithm for computing \mathbf{P} and the quotients $\mathbf{Q} = -\mathbf{P}\mathbf{F}/\text{diag}(\mathfrak{M})$ when we know *a priori* the *s-minimal* degree δ of $(\mathfrak{M}, \mathbf{F})$. Here, we would choose $\delta = \max(\delta) \geq \deg([\mathbf{P}|\mathbf{Q}])$: in some cases $\delta = \Theta(\sigma)$ and this has cost bound $\tilde{\mathcal{O}}(m^{\omega-1}(\sigma + n\sigma))$, which exceeds our target $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$. An issue is that \mathbf{Q} has size $\mathcal{O}(mn\sigma)$ when \mathbf{P} has columns of large degree; yet here we are not interested in \mathbf{Q} . This can be solved using partial linearization to expand the columns of large degree in \mathbf{P} into more columns of smaller degree as in the next result, which holds in general for interpolation bases [21, Lemma 4.2].

LEMMA 2.3. *Let $\mathfrak{M} \in \mathbb{K}[X]_{\neq 0}^n$ with entries having degrees $(\sigma_1, \dots, \sigma_n)$. Let $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$ and $\mathbf{s} \in \mathbb{Z}^m$. Furthermore, let $\delta = (\delta_1, \dots, \delta_m)$ denote the *s-minimal* degree of $(\mathfrak{M}, \mathbf{F})$.*

Writing $\sigma = \sigma_1 + \dots + \sigma_n$, let $\delta = \lceil \sigma/m \rceil \geq 1$, and for $i \in \{1, \dots, m\}$ write $\delta_i = (\alpha_i - 1)\delta + \beta_i$ with $\alpha_i \geq 1$ and $0 \leq \beta_i < \delta$, and let $\tilde{m} = \alpha_1 + \dots + \alpha_m$. Define $\tilde{\delta} \in \mathbb{N}^{\tilde{m}}$ as

$$\tilde{\delta} = (\underbrace{\delta, \dots, \delta}_{\alpha_1}, \beta_1, \dots, \underbrace{\delta, \dots, \delta}_{\alpha_m}, \beta_m) \quad (2)$$

and the expansion-compression matrix $\mathcal{E} \in \mathbb{K}[X]^{\tilde{m} \times m}$ as

$$\mathcal{E} = \begin{bmatrix} 1 \\ X^\delta \\ \vdots \\ X^{(\alpha_1-1)\delta} & \ddots & \\ & & 1 \\ & & X^\delta \\ & & \vdots \\ & & X^{(\alpha_m-1)\delta} \end{bmatrix}. \quad (3)$$

*Let $\mathbf{d} = -\tilde{\delta} \in \mathbb{Z}^{\tilde{m}}$ and $\mathbf{P} \in \mathbb{K}[X]^{\tilde{m} \times \tilde{m}}$ be the *d-Popov* solution basis for $(\mathfrak{M}, \mathcal{E}\mathbf{F} \pmod{\mathfrak{M}})$. Then, \mathbf{P} has *d-pivot* degree $\tilde{\delta}$ and the *s-Popov* solution basis for $(\mathfrak{M}, \mathbf{F})$ is the submatrix of $\mathbf{P}\mathcal{E}$ formed by its rows at indices $\{\alpha_1 + \dots + \alpha_i, 1 \leq i \leq m\}$.*

This leads to Algorithm 1, which solves Problem 2 efficiently when the *s-minimal* degree δ is known *a priori*.

ALGORITHM 1 (KNOWNDEGPOLMODSYS).

Input: polynomials $\mathfrak{M} = (\mathbf{m}_1, \dots, \mathbf{m}_n) \in \mathbb{K}[X]_{\neq 0}^n$, a matrix $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$ with $\deg(\mathbf{F}_{*,j}) < \deg(\mathbf{m}_j)$, a shift $\mathbf{s} \in \mathbb{Z}^m$, $\delta = (\delta_1, \dots, \delta_m)$ the *s-minimal* degree of $(\mathfrak{M}, \mathbf{F})$.

Output: the *s-Popov* solution basis for $(\mathfrak{M}, \mathbf{F})$.

1. $\delta \leftarrow \lceil (\deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n))/m \rceil$,
 $\alpha_i \leftarrow \lfloor \delta_i/\delta \rfloor + 1$ for $1 \leq i \leq m$, $\tilde{m} \leftarrow \alpha_1 + \dots + \alpha_m$,
 $\tilde{\delta}$ as in (2), \mathcal{E} as in (3), $\tilde{\mathbf{F}} \leftarrow \mathcal{E}\mathbf{F} \pmod{\mathfrak{M}}$
2. $\mathbf{u} \leftarrow (-\tilde{\delta}, -\delta - 1, \dots, -\delta - 1) \in \mathbb{Z}^{\tilde{m}+n}$
 $\tau \leftarrow (\deg(\mathbf{m}_j) + \delta + 1)_{1 \leq j \leq n}$
3. $\tilde{\mathbf{P}} \leftarrow$ the *u-Popov* order basis for $[\tilde{\mathbf{F}}^\top | \text{diag}(\mathfrak{M})]^\top$ and τ
 $\mathbf{P} \leftarrow$ the principal $\tilde{m} \times \tilde{m}$ submatrix of $\tilde{\mathbf{P}}$
4. Return the submatrix of $\mathbf{P}\mathcal{E}$ formed by the rows at indices $\alpha_1 + \dots + \alpha_i$ for $1 \leq i \leq m$

PROPOSITION 2.4. *Algorithm KNOWNDEGPOLMODSYS is correct. Writing $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n)$ and assuming $\sigma \geq m \geq n$, it uses $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ operations in \mathbb{K} .*

PROOF. By Lemmas 2.3 and 2.1, since $\min(-\tilde{\delta}) > -\delta - 1$ and $\mathbf{u} = (-\tilde{\delta}, -\delta - 1, \dots, -\delta - 1)$, the $-\tilde{\delta}$ -*Popov* solution

basis for $(\mathfrak{M}, \tilde{\mathbf{F}})$ is the principal $\tilde{m} \times \tilde{m}$ submatrix of the \mathbf{u} -Popov nullspace basis \mathbf{B} for $[\tilde{\mathbf{F}}^\top \text{diag}(\mathfrak{M})]^\top$, and \mathbf{B} has \mathbf{u} -pivot index $\{1, \dots, \tilde{m}\}$, \mathbf{u} -pivot degree $\tilde{\delta}$, and $\deg(\mathbf{B}) \leq \delta$. Then, by Lemma 2.2, \mathbf{B} is formed by the first \tilde{m} rows of $\tilde{\mathbf{P}}$ at Step 3, hence \mathbf{P} is the \mathbf{d} -Popov solution basis for $(\mathfrak{M}, \mathbf{F})$. The correctness then follows from Lemma 2.3.

Since $|\delta| \leq \sigma$, \mathcal{E} has $\tilde{m} \leq 2m$ rows and $\mathcal{E}\mathbf{F} \bmod \mathfrak{M}$ can be computed in $\tilde{\mathcal{O}}(m\sigma)$ operations using fast polynomial division [13]. The cost bound of Step 3 follows from [21, Theorem 1.4] since $\tau_1 + \dots + \tau_n = \sigma + n(1 + \lceil \sigma/m \rceil) \in \mathcal{O}(\sigma)$. \square

2.2 The case of one equation

We now present our main new ingredients, focusing on the case $n = 1$. First, we show that when the shift \mathbf{s} has a small *amplitude* $\text{amp}(\mathbf{s}) = \max(\mathbf{s}) - \min(\mathbf{s})$, one can solve Problem 2 via an order basis computation at small order.

LEMMA 2.5. *Let $\mathbf{m} \in \mathbb{K}[X]_{\neq 0}$, $\mathbf{s} \in \mathbb{Z}^m$, and $\mathbf{F} \in \mathbb{K}[X]^{m \times 1}$ with $\deg(\mathbf{F}) < \deg(\mathbf{m}) = \sigma$. Then, for any $\tau \geq \text{amp}(\mathbf{s}) + 2\sigma$, the \mathbf{s} -Popov solution basis for (\mathbf{m}, \mathbf{F}) is the principal $m \times m$ submatrix of the \mathbf{u} -Popov order basis for $[\mathbf{F}^\top | \mathbf{m}]^\top$ and τ , with $\mathbf{u} = (\mathbf{s}, \min(\mathbf{s})) \in \mathbb{Z}^{m+1}$.*

PROOF. Let $\mathbf{A} = \begin{bmatrix} \mathbf{p} & \mathbf{q} \\ \mathbf{p} & \mathbf{q} \end{bmatrix}$ denote the \mathbf{u} -Popov order basis for $[\mathbf{F}^\top | \mathbf{m}]^\top$ and τ , where $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$ and $q \in \mathbb{K}[X]$. Consider $\mathbf{B} = [\tilde{\mathbf{P}} | \tilde{\mathbf{q}}]$ the \mathbf{u} -Popov nullspace basis of $[\mathbf{F}^\top | \mathbf{m}]^\top$: thanks to Lemma 2.1, it is enough to prove that $\mathbf{B} = [\mathbf{P} | \mathbf{q}]$.

First, we have $\text{rdeg}(\mathbf{p}) \leq \deg(q)$ by choice of \mathbf{u} , so that $q\mathbf{m} \neq 0$ implies $\deg(\mathbf{p}\mathbf{F} + q\mathbf{m}) = \deg(q) + \sigma$. Since $\mathbf{p}\mathbf{F} + q\mathbf{m} = 0 \bmod X^\tau$, this gives $\deg(q) + \sigma \geq \tau$. This also shows that the \mathbf{u} -pivot entries of \mathbf{B} are located in $\tilde{\mathbf{P}}$.

Then, since the sum of the \mathbf{u} -pivot degrees of \mathbf{A} is at most τ , the sum of the \mathbf{s} -pivot degrees of \mathbf{P} is at most σ ; with $[\mathbf{P} | \mathbf{q}]$ in \mathbf{u} -Popov form, this gives $\deg(\mathbf{q}) < \sigma + \text{amp}(\mathbf{s}) \leq \tau - \sigma$. We obtain $\deg(\mathbf{P}\mathbf{F} + \mathbf{q}\mathbf{m}) < \tau$, so that $\mathbf{P}\mathbf{F} + \mathbf{q}\mathbf{m} = 0$. Thus, the minimality of \mathbf{B} and \mathbf{A} gives the conclusion. \square

When $\text{amp}(\mathbf{s}) \in \mathcal{O}(\sigma)$, this gives a fast solution to our problem. In what follows, we present a divide-and-conquer approach on $\text{amp}(\mathbf{s})$, with base case $\text{amp}(\mathbf{s}) \in \mathcal{O}(\sigma)$.

We first give an overview, assuming \mathbf{s} is non-decreasing. A key ingredient is that when $\text{amp}(\mathbf{s})$ is large compared to σ , then \mathbf{P} has a lower block triangular shape, since it is in \mathbf{s} -Popov form with sum of \mathbf{s} -pivot degrees $|\delta| \leq \sigma$. Typically, if $s_{i+1} - s_i \geq \sigma$ for some i then $\mathbf{P} = \begin{bmatrix} \mathbf{P}^{(1)} & \mathbf{0} \\ * & \mathbf{P}^{(2)} \end{bmatrix}$ with $\mathbf{P}^{(1)} \in \mathbb{K}[X]^{i \times i}$. Even though the block sizes are unknown in general, we show that they can be revealed efficiently along with δ by a divide-and-conquer algorithm, as follows.

First, we use a recursive call with the first j entries $\mathbf{s}^{(0)}$ of \mathbf{s} and $\mathbf{F}^{(0)}$ of \mathbf{F} , where j is such that $\text{amp}(\mathbf{s}^{(0)})$ is about half of $\text{amp}(\mathbf{s})$. This reveals the first $i \leq j$ entries $\delta^{(1)}$ of δ and the first i rows $[\mathbf{P}^{(1)} | \mathbf{0}]$ of \mathbf{P} , with $\mathbf{P}^{(1)} \in \mathbb{K}[X]^{i \times i}$. A central point is that $\text{amp}(\mathbf{s}^{(2)})$ is about half of $\text{amp}(\mathbf{s})$ as well, where $\mathbf{s}^{(2)}$ is the tail of \mathbf{s} starting at the entry $i + 1$.

Then, knowing the degrees $\delta^{(1)}$ allows us to set up an order basis computation that yields a *residual*, that is, a column $\mathbf{G} \in \mathbb{K}[X]^{(m-i) \times 1}$ and a modulus \mathfrak{n} such that we can continue the computation of \mathbf{P} using a second recursive call, which consists in computing the $\mathbf{s}^{(2)}$ -Popov solution basis for $(\mathfrak{n}, \mathbf{G})$. From these two calls we obtain δ , and then we recover \mathbf{P} using Algorithm 1.

Now we present the details. We fix $\mathbf{F} \in \mathbb{K}[X]^{m \times 1}$, $\mathbf{m} \in \mathbb{K}[X]_{\neq 0}$ with $\sigma = \deg(\mathbf{m}) > \deg(\mathbf{F})$, $\mathbf{s} \in \mathbb{Z}^m$, \mathbf{P} the \mathbf{s} -Popov

solution basis for (\mathbf{m}, \mathbf{F}) , and δ its \mathbf{s} -pivot degree. In what follows, $\boldsymbol{\pi}^{\mathbf{s}} = (\pi_1, \dots, \pi_m)$ is any permutation of $\{1, \dots, m\}$ such that $(s_{\pi_1}, \dots, s_{\pi_m})$ is non-decreasing.

Then, for $\mathbf{t} = (t_1, \dots, t_m) \in \mathbb{Z}^m$ we write $\mathbf{t}_{[i:j]}$ for the subtuple of \mathbf{t} formed by its entries at indices $\{\pi_i, \dots, \pi_j\}$, and for a matrix $\mathbf{M} \in \mathbb{K}[X]^{m \times m}$ we write $\mathbf{M}_{[i:j, k:l]}$ for the submatrix of \mathbf{M} formed by its rows at indices $\{\pi_i, \pi_{i+1}, \dots, \pi_j\}$ and columns at indices $\{\pi_k, \pi_{k+1}, \dots, \pi_l\}$. The main ideas in this subsection can be understood by focusing on the case of a non-decreasing \mathbf{s} , taking $\pi_i = i$ for all i : then we have $\mathbf{t}_{[i:j]} = (t_i, t_{i+1}, \dots, t_j)$ and $\mathbf{M}_{[i:j, k:l]} = (\mathbf{M}_{u,v})_{i \leq u \leq j, k \leq v \leq l}$.

We now introduce the notion of splitting index, which will help us to locate zero blocks in \mathbf{P} .

DEFINITION 2.6 (SPLITTING INDEX). *Let $\mathbf{d} \in \mathbb{N}^m$, $\mathbf{t} \in \mathbb{Z}^m$ and $\boldsymbol{\pi}^{\mathbf{t}} = (\mu_i)_i$. Then, $i \in \{1, \dots, m-1\}$ is a splitting index for (\mathbf{d}, \mathbf{t}) if $d_{\mu_j} + t_{\mu_j} - t_{\mu_{j+1}} < 0$ for all $j \in \{1, \dots, i\}$.*

In particular, if i is a splitting index for (δ, \mathbf{s}) , then we have $[\mathbf{P}_{[i:i]} | \mathbf{P}_{[i:i+1:i]}] = [\mathbf{P}_{[i:i]} | \mathbf{0}]$. Our algorithm first looks for such a splitting index, and then uses $\mathbf{P}_{[i:i+1:i]} = \mathbf{0}$ to split the problem into two subproblems of dimensions i and $m-i$.

To find a splitting index, we rely on the following property: if (\mathbf{d}, \mathbf{t}) does not admit a splitting index, then $|\mathbf{d}| > \text{amp}(\mathbf{t})$. This allows us to partition \mathbf{s} into ℓ subtuples which all contain a splitting index, as follows.

Given $\alpha \in \mathbb{Z}_{>0}$ we let $\ell = 1 + \lceil \text{amp}(\mathbf{s})/\alpha \rceil$ and we consider the subtuples $\mathbf{s}_1, \dots, \mathbf{s}_\ell$ of \mathbf{s} where \mathbf{s}_k consists of the entries of \mathbf{s} in $\{\min(\mathbf{s}) + (k-1)\alpha, \dots, \min(\mathbf{s}) + k\alpha - 1\}$; this gives a subroutine $\text{PARTITION}(\mathbf{s}, \alpha) = (\mathbf{s}_1, \dots, \mathbf{s}_\ell)$. Now we take $\alpha \geq 2\sigma$ and we assume $s_{\pi_{i+1}} - s_{\pi_i} \leq \sigma$ for $1 \leq i < m$ without loss of generality [21, Appendix A]. Then, for $1 \leq k < \ell$, since $|\delta| \leq \sigma$ and $\text{amp}(\mathbf{t}) \geq \sigma$ with $\mathbf{t} = (\mathbf{s}_k, \min(\mathbf{s}_{k+1}))$, by the above remark \mathbf{s}_k contains a splitting index for (δ, \mathbf{s}) .

Still, we do not know in advance which entries of \mathbf{s}_k correspond to splitting indices for (δ, \mathbf{s}) . Thus we recursively compute the \mathbf{s} -Popov solution basis $\mathbf{P}^{(0)}$ for $\mathbf{s}_1, \dots, \mathbf{s}_{\ell/2}$, and we are now going to prove that this gives us a splitting index which divides the computation into two subproblems, the first of which has been already solved by computing $\mathbf{P}^{(0)}$.

LEMMA 2.7. *Let $j \in \{2, \dots, m\}$, $\mathbf{s}^{(0)} = \mathbf{s}_{[j:]}$, $\mathbf{P}^{(0)}$ be the $\mathbf{s}^{(0)}$ -Popov solution basis for $(\mathbf{m}, \mathbf{F}_{[j:]})$, and $\delta^{(0)}$ be its $\mathbf{s}^{(0)}$ -pivot degree. Suppose that there is a splitting index $i \leq j$ for $(\delta^{(0)}, \mathbf{s}^{(0)})$. Let $\mathbf{P}^{(1)} \in \mathbb{K}[X]^{i \times i}$ be the $\mathbf{s}^{(1)}$ -Popov solution basis for $(\mathbf{m}, \mathbf{F}_{[i:]})$ with $\mathbf{s}^{(1)} = \mathbf{s}_{[i:]}$, and let $\delta^{(1)}$ be its $\mathbf{s}^{(1)}$ -pivot degree. Then i is a splitting index for (δ, \mathbf{s}) and $\mathbf{P}_{[i:i]} = \mathbf{P}^{(1)} = \mathbf{P}^{(0)}$, hence $\delta_{[i:i]} = \delta^{(1)} = \delta^{(0)}$ (where $\mathbf{P}^{(0)}$ and $\delta^{(0)}$ are indexed by $\{\pi_1, \dots, \pi_j\}$ sorted increasingly).*

PROOF. Since i is a splitting index for $(\delta^{(0)}, \mathbf{s}^{(0)})$ we have $[\mathbf{P}_{[i:i]}^{(0)} | \mathbf{P}_{[i:i+1:i]}^{(0)}] = [\mathbf{Q} | \mathbf{0}]$ for some $\mathbf{Q} \in \mathbb{K}[X]^{i \times i}$. Now, for any $\mathbf{B} \in \mathbb{K}[X]^{m \times m}$ with $[\mathbf{B}_{[i:i, i]} | \mathbf{B}_{[i:i+1:i]}] = [\mathbf{P}^{(1)} | \mathbf{0}]$, $\mathbf{B}_{[i:i]}$ is in \mathbf{s} -Popov form with its rows being solutions for $(\mathfrak{M}, \mathbf{F})$. Then, by minimality of \mathbf{P} , $\mathbf{P}_{[i:i]}$ has \mathbf{s} -pivot degree at most $\delta^{(1)}$ componentwise, so that i is also a splitting index for (δ, \mathbf{s}) , and in particular $[\mathbf{P}_{[i:i]} | \mathbf{P}_{[i:i+1:i]}] = [\mathbf{R} | \mathbf{0}]$ for some $\mathbf{R} \in \mathbb{K}[X]^{i \times i}$. It remains to prove that $\mathbf{Q} = \mathbf{R} = \mathbf{P}^{(1)}$.

Since $\mathbf{R}\mathbf{F}_{[i]} = 0 \bmod \mathfrak{m}$ and $\mathbf{R} = \mathbf{P}_{[i:i]}$ is in $\mathbf{s}^{(1)}$ -Popov form, proving that all solutions $\mathbf{p} \in \mathbb{K}[X]^{1 \times i}$ for $(\mathfrak{m}, \mathbf{F}_{[i]})$ are in the row space of \mathbf{R} is enough to obtain $\mathbf{R} = \mathbf{P}^{(1)}$. Since $\mathbf{q} \in \mathbb{K}[X]^{1 \times m}$ defined by $[\mathbf{q}_{[i]} | \mathbf{q}_{[i+1:i]}] = [\mathbf{p} | \mathbf{0}]$ is a solution for (\mathbf{m}, \mathbf{F}) , $\mathbf{q} = \boldsymbol{\lambda}\mathbf{P}$ for some $\boldsymbol{\lambda} \in \mathbb{K}[X]^{1 \times m}$. Now \mathbf{P} is nonsingular, thus $\mathbf{P}_{[i:i+1:i]} = \mathbf{0}$ implies that $[\boldsymbol{\lambda}_{[i]} | \boldsymbol{\lambda}_{[i+1:i]}] =$

$[\boldsymbol{\mu}|\mathbf{0}]$ with $\boldsymbol{\mu} \in \mathbb{K}[X]^{1 \times i}$, hence $\mathbf{p} = \mathbf{q}_{[i]} = \boldsymbol{\lambda}_{[i]} \mathbf{P}_{[i,i]} + \boldsymbol{\lambda}_{[i+1:i]} \mathbf{P}_{[i+1:i,i]} = \boldsymbol{\mu} \mathbf{Q}$. Similar arguments give $\mathbf{Q} = \mathbf{P}^{(1)}$. \square

The next two lemmas show that knowing $\boldsymbol{\delta}^{(1)}$, which is $\boldsymbol{\delta}_{[i]}$, allows us to compute a so-called *residual* (\mathbf{n}, \mathbf{G}) from which we can complete the computation of $\boldsymbol{\delta}$ and \mathbf{P} .

LEMMA 2.8. *Let $\mathbf{s}^{(2)} = \mathbf{s}_{[i+1:]}$, $\mathbf{d} = -\boldsymbol{\delta}^{(1)} + \min(\mathbf{s}^{(2)}) - 2\sigma \in \mathbb{Z}^i$, $\mathbf{v} \in \mathbb{Z}^m$ be such that $[\mathbf{v}_{[i]}|\mathbf{v}_{[i+1:]}] = [\mathbf{d}|\mathbf{s}^{(2)}]$, and $\mathbf{u} = (\mathbf{v}, \min(\mathbf{d})) \in \mathbb{Z}^{m+1}$. Let $\begin{bmatrix} \mathbf{A} & \mathbf{q} \\ \mathbf{p} & q \end{bmatrix}$ be the \mathbf{u} -Popov order basis for $[\mathbf{F}^\top|\mathbf{m}]^\top$ and 2σ , where $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ and $q \in \mathbb{K}[X]$. Then we have $\deg(q) \geq \sigma$, $\mathbf{A}_{[i,i,i+1]} = \mathbf{0}$, $\mathbf{p}_{[i+1:]} = \mathbf{0}$, and $[\mathbf{A}_{[i,i,i]}|\mathbf{q}_{[i]}] = [\mathbf{P}^{(1)}|\mathbf{q}^{(1)}]$ with $\mathbf{q}^{(1)} = -\mathbf{P}^{(1)}\mathbf{F}_{[i]}/m$.*

PROOF. Since $\mathbf{u} = (\mathbf{v}, \min(\mathbf{v}))$ we have $\deg(\mathbf{p}) \leq \deg(q)$, and since $\deg(\mathbf{F}) < \deg(m)$ the degree of $\mathbf{p}\mathbf{F} + \mathbf{q}m$ is $\deg(q) + \sigma$; then $\mathbf{p}\mathbf{F} + \mathbf{q}m = 0 \pmod{X^{2\sigma}}$ implies $\deg(q) + \sigma \geq 2\sigma$. Now, since \mathbf{A} is in \mathbf{v} -Popov form and $\deg(\mathbf{A}) \leq 2\sigma - \deg(q) < 2\sigma$, from $\min(\mathbf{s}^{(2)}) \geq \max(\mathbf{d}) + 2\sigma$ we get $\mathbf{A}_{[i,i,i+1]} = \mathbf{0}$. Besides, $\mathbf{p}_{[i+1:]} = \mathbf{0}$ since either $\deg(q) < 2\sigma$ and then $\min(\mathbf{s}^{(2)}) > \min(\mathbf{d}) + \deg(q)$, or \mathbf{A} is the identity matrix and then $\mathbf{p} = \mathbf{0}$.

Furthermore, by Lemma 2.1 $[\mathbf{P}^{(1)}|\mathbf{q}^{(1)}]$ is the $(\mathbf{d}, \min(\mathbf{d}))$ -Popov nullspace basis for $[\mathbf{F}^\top|\mathbf{m}]^\top$, with $(\mathbf{d}, \min(\mathbf{d}))$ -pivot index $\{1, \dots, i\}$, $(\mathbf{d}, \min(\mathbf{d}))$ -pivot degree $\boldsymbol{\delta}^{(1)}$ and degree at most $\max(\boldsymbol{\delta}^{(1)})$. Then, as in the proof of Lemma 2.2, one can show that $[\mathbf{A}_{[i,i,i]}|\mathbf{q}_{[i]}] = [\mathbf{P}^{(1)}|\mathbf{q}^{(1)}]$. \square

Thus, up to row and column permutations this order basis is $\begin{bmatrix} \mathbf{P}^{(1)} & \mathbf{0} & \mathbf{q}^{(1)} \\ * & \mathbf{P}^{(2)} & * \\ * & * & q \end{bmatrix}$ with $\mathbf{P}^{(2)} = \mathbf{A}_{[i+1:i+1]} \in \mathbb{K}[X]^{(m-i) \times (m-i)}$ in $\mathbf{s}^{(2)}$ -Popov form; let $\boldsymbol{\delta}^{(2)}$ denote its $\mathbf{s}^{(2)}$ -pivot degree.

LEMMA 2.9. *Let $\mathbf{n} = X^{-2\sigma}(\mathbf{p}_{[i+1:]} \mathbf{F}_{[i+1:]} + \mathbf{q}m) \in \mathbb{K}[X]$ and $\mathbf{G} = X^{-2\sigma}(\mathbf{A}_{[i+1:i+1]} \mathbf{F} + \mathbf{q}_{[i+1:]} \mathbf{m}) \in \mathbb{K}[X]^{(m-i) \times 1}$. Then, $\deg(\mathbf{G}) < \deg(\mathbf{n}) \leq \sigma - |\boldsymbol{\delta}^{(1)}| - |\boldsymbol{\delta}^{(2)}|$. Let $\mathbf{P}^{(3)}$ be the \mathbf{t} -Popov solution basis for (\mathbf{n}, \mathbf{G}) with $\mathbf{t} = \text{rdeg}_{\mathbf{s}^{(2)}}(\mathbf{P}^{(2)})$ and $\boldsymbol{\delta}^{(3)}$ be its \mathbf{t} -pivot degree. Then, $(\boldsymbol{\delta}_{[i]}, \boldsymbol{\delta}_{[i+1:]}) = (\boldsymbol{\delta}^{(1)}, \boldsymbol{\delta}^{(2)} + \boldsymbol{\delta}^{(3)})$.*

PROOF. The sum $|\boldsymbol{\delta}^{(1)}| + |\boldsymbol{\delta}^{(2)}| + \deg(q)$ of the \mathbf{u} -pivot degrees of $\begin{bmatrix} \mathbf{A} & \mathbf{q} \\ \mathbf{p} & q \end{bmatrix}$ is at most the order 2σ . Thus, we have $\deg(\mathbf{n}) = \deg(q) - \sigma \leq \sigma - |\boldsymbol{\delta}^{(1)}| - |\boldsymbol{\delta}^{(2)}|$, $\deg(\mathbf{A}_{[i+1:i+1]}) < |\boldsymbol{\delta}^{(1)}| \leq \sigma$, $\deg(\mathbf{A}_{[i+1:i+1]}) \leq |\boldsymbol{\delta}^{(2)}| \leq \sigma$, and $\deg(\mathbf{q}_{[i+1:]}) < \deg(q)$. This implies $\deg(\mathbf{G}) < \deg(q) - \sigma = \deg(\mathbf{n})$.

Let $\mathbf{q}^{(3)} = -\mathbf{P}^{(3)}\mathbf{G}/\mathbf{n}$ and $t = \text{rdeg}_{\mathbf{u}}([\mathbf{p}|q]) = \deg(q) + \min(\mathbf{d}) \leq \min(\mathbf{s}^{(2)}) \leq \min(\mathbf{t})$. By Lemma 2.1, $[\mathbf{P}^{(3)}|\mathbf{q}^{(3)}]$ is the (\mathbf{t}, t) -Popov nullspace basis for $[\mathbf{G}^\top|\mathbf{n}]^\top$. Defining $\mathbf{B} \in \mathbb{K}[X]^{m \times m}$ and $\mathbf{c} \in \mathbb{K}[X]^{m \times 1}$ by $\begin{bmatrix} \mathbf{B}_{[i,i,i]} & \mathbf{B}_{[i,i,i+1]} & \mathbf{c}_{[i]} \\ \mathbf{B}_{[i+1:i+1,i]} & \mathbf{B}_{[i+1:i+1,i+1]} & \mathbf{c}_{[i+1:]} \end{bmatrix} =$

$\begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{P}^{(3)} & \mathbf{q}^{(3)} \end{bmatrix}$, then $[\mathbf{B} \ \mathbf{c}] \begin{bmatrix} \mathbf{A} & \mathbf{q} \\ \mathbf{p} & q \end{bmatrix}$ is a \mathbf{u} -minimal nullspace basis of $[\mathbf{F}^\top|\mathbf{m}]^\top$ [36, Theorem 3.9]. Thus Lemma 2.1 implies that $\tilde{\mathbf{P}} = [\mathbf{B} \ \mathbf{c}] \begin{bmatrix} \mathbf{A} \\ \mathbf{p} \end{bmatrix}$ is a \mathbf{v} -minimal solution basis for (\mathbf{m}, \mathbf{F}) .

It is easily checked that $\tilde{\mathbf{P}}$ is in \mathbf{v} -Popov form, so that the \mathbf{v} -Popov form of $\tilde{\mathbf{P}}$ is \mathbf{P} and its \mathbf{v} -pivot degree is $\boldsymbol{\delta}$. Besides $\begin{bmatrix} \mathbf{P}_{[i,i,i]} & \mathbf{P}_{[i,i,i+1]} \\ \mathbf{P}_{[i+1:i+1,i]} & \mathbf{P}_{[i+1:i+1,i+1]} \end{bmatrix} = \begin{bmatrix} \mathbf{P}^{(3)} & \mathbf{P}^{(1)} \\ \mathbf{A}_{2,1} + \mathbf{q}^{(3)} \mathbf{A}_{3,1} & \mathbf{P}^{(3)} \mathbf{P}^{(2)} \end{bmatrix}$, so that $(\boldsymbol{\delta}_{[i]}, \boldsymbol{\delta}_{[i+1:]}) = (\boldsymbol{\delta}^{(1)}, \boldsymbol{\delta}^{(2)} + \boldsymbol{\delta}^{(3)})$ [21, Section 3]. \square

This results in Algorithm 2. It takes as input α which dictates the amplitude of the subtuples that partition \mathbf{s} ; as mentioned above, the initial call can be made with $\alpha = 2\sigma$.

PROPOSITION 2.10. *Algorithm POLMODSYSONE is correct and uses $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ operations in \mathbb{K} .*

PROOF. The correctness follows from the results in this subsection. By [21, Theorem 1.4], each leaf of the recursion at Step 1.a in dimension m uses $\tilde{\mathcal{O}}(m^{\omega-1}\alpha)$ operations.

ALGORITHM 2 (POLMODSYSONE).

Input: a polynomial $\mathbf{m} \in \mathbb{K}[X]_{\neq 0}$ of degree σ , a column $\mathbf{F} \in \mathbb{K}[X]^{m \times 1}$ with $\deg(\mathbf{F}) < \deg(m)$, a shift $\mathbf{s} \in \mathbb{Z}^m$, a parameter $\alpha \in \mathbb{Z}_{>0}$ with $\alpha \geq 2\sigma$.

Output: the \mathbf{s} -Popov solution basis for (\mathbf{m}, \mathbf{F}) and the \mathbf{s} -minimal degree $\boldsymbol{\delta}$ of (\mathbf{m}, \mathbf{F}) .

1. If $\text{amp}(\mathbf{s}) \leq 2\alpha$:
 - a. $\mathbf{A} \leftarrow$ the $(\mathbf{s}, \min(\mathbf{s}))$ -Popov order basis for $[\mathbf{F}^\top|\mathbf{m}]^\top$ and $2\alpha + 2\sigma$; return the principal $m \times m$ submatrix of \mathbf{A} and the degrees of its diagonal entries
2. Else: /* $\ell = 1 + \lfloor \text{amp}(\mathbf{s})/\alpha \rfloor \geq 3$ */
 - a. $(\mathbf{s}_1, \dots, \mathbf{s}_\ell) \leftarrow \text{PARTITION}(\mathbf{s}, \alpha)$,
 $j \leftarrow$ sum of the lengths of $\mathbf{s}_1, \dots, \mathbf{s}_{\lfloor \ell/2 \rfloor}$, $\mathbf{s}^{(0)} \leftarrow \mathbf{s}_{[j]}$,
 $(\mathbf{P}^{(0)}, \boldsymbol{\delta}^{(0)}) \leftarrow \text{POLMODSYSONE}(\mathbf{m}, \mathbf{F}_{[j]}, \mathbf{s}^{(0)}, \alpha)$
 - b. $i \leftarrow$ the largest splitting index for $(\boldsymbol{\delta}^{(0)}, \mathbf{s}^{(0)})$, $\boldsymbol{\delta}^{(1)} \leftarrow \boldsymbol{\delta}_{[i]}^{(0)}$, $\mathbf{s}^{(2)} \leftarrow \mathbf{s}_{[i+1:]}$, $\mathbf{d} = -\boldsymbol{\delta}^{(1)} + \min(\mathbf{s}^{(2)}) - 2\sigma$, $\mathbf{v} \in \mathbb{Z}^m$ with $[\mathbf{v}_{[i]}|\mathbf{v}_{[i+1:]}] \leftarrow [\mathbf{d}|\mathbf{s}^{(2)}]$, $\mathbf{u} = (\mathbf{v}, \min(\mathbf{d}))$
 - c. $\begin{bmatrix} \mathbf{A} & \mathbf{q} \\ \mathbf{p} & q \end{bmatrix} \leftarrow$ \mathbf{u} -Popov order basis for $[\mathbf{F}^\top|\mathbf{m}]^\top$ and 2σ ,
 $\boldsymbol{\delta}^{(2)} \leftarrow$ the $\mathbf{s}^{(2)}$ -pivot degree of $\mathbf{A}_{[i+1:i+1]}$,
 $\mathbf{G} \leftarrow X^{-2\sigma}(\mathbf{A}_{[i+1:i+1]} \mathbf{F} + \mathbf{q}_{[i+1:]} \mathbf{m})$,
 $\mathbf{n} \leftarrow X^{-2\sigma}(\mathbf{p}_{[i+1:]} \mathbf{F}_{[i+1:]} + \mathbf{q}m)$.
 - d. $\mathbf{t} \leftarrow \mathbf{s}^{(2)} + \boldsymbol{\delta}^{(2)} = \text{rdeg}_{\mathbf{s}^{(2)}}(\mathbf{A}_{[i+1:i+1]})$,
 $(\mathbf{P}^{(3)}, \boldsymbol{\delta}^{(3)}) \leftarrow \text{POLMODSYSONE}(\mathbf{n}, \mathbf{G}, \mathbf{t}, \alpha)$
 - e. $\boldsymbol{\delta} \in \mathbb{N}^m$ with $(\boldsymbol{\delta}_{[i]}, \boldsymbol{\delta}_{[i+1:]}) \leftarrow (\boldsymbol{\delta}^{(1)}, \boldsymbol{\delta}^{(2)} + \boldsymbol{\delta}^{(3)})$
 $\mathbf{P} \leftarrow \text{KNOWNDEGPOLMODSYS}(\mathbf{m}, \mathbf{F}, \mathbf{s}, \boldsymbol{\delta})$
 - f. Return $(\mathbf{P}, \boldsymbol{\delta})$

Running the algorithm with initial input $\alpha = 2\sigma$, the recursive tree has depth $\mathcal{O}(\log(\ell)) = \mathcal{O}(\log(1 + \text{amp}(\mathbf{s})/2\sigma))$, with $\text{amp}(\mathbf{s})/2\sigma \in \mathcal{O}(m^2)$ [21, Appendix A]. All recursive calls are for a modulus of degree $\sigma < \alpha$. The order basis computation at Step 2.c uses $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ operations; the computation of \mathbf{G} and \mathbf{n} at Step 2.c can be done in time $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ using partial linearization as in Lemma 2.11 below; Step 2.e uses $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ operations by Proposition 2.4.

On a given level of the tree, the sum of the dimensions of the column vector in input of each sub-problem is in $\mathcal{O}(m)$. Since $a^{\omega-1} + b^{\omega-1} \leq (a+b)^{\omega-1}$ for all $a, b > 0$, each level of the tree uses a total of $\tilde{\mathcal{O}}(m^{\omega-1}\alpha)$ operations. \square

2.3 Fast divide-and-conquer algorithm

Now that we have an efficient algorithm for $n = 1$, our main algorithm uses a divide-and-conquer approach on n . Similarly to [21, Algorithm 1], from the two bases obtained recursively we first deduce the \mathbf{s} -minimal degree $\boldsymbol{\delta}$, and then we use this knowledge to compute \mathbf{P} with Algorithm 1. When $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n) \in \mathcal{O}(m)$, we rely on the algorithm LINEARIZATIONMIB in [20, Algorithm 9].

The computation of the so-called *residual* at Step 3.c can be done efficiently using partial linearization, as follows.

LEMMA 2.11. *Let $\mathfrak{M} = (\mathbf{m}_j)_j \in \mathbb{K}[X]_{\neq 0}^n$, $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$, $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$ with $m \geq n$ and $\deg(\mathbf{F}_{*,j}) < \sigma_j = \deg(\mathbf{m}_j)$, and let $\sigma \geq m$ such that $\sigma \geq \sigma_1 + \dots + \sigma_n$ and $|\text{cdeg}(\mathbf{P})| \leq \sigma$. Then $\mathbf{P}\mathbf{F} \pmod{\mathfrak{M}}$ can be computed in $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ operations.*

PROOF. Using notation from Lemma 2.3, we let $\tilde{\mathbf{P}} \in \mathbb{K}[X]^{m \times \tilde{m}}$ such that $\mathbf{P} = \tilde{\mathbf{P}}\mathcal{E}$ and $\deg(\tilde{\mathbf{P}}) < \lceil |\text{cdeg}(\mathbf{P})|/m \rceil$. As above, $\tilde{\mathbf{F}} = \mathcal{E}\mathbf{F} \pmod{\mathfrak{M}}$ can be computed in time $\tilde{\mathcal{O}}(m\sigma)$. Here we want to compute $\mathbf{P}\mathbf{F} \pmod{\mathfrak{M}} = \tilde{\mathbf{P}}\tilde{\mathbf{F}} \pmod{\mathfrak{M}}$.

ALGORITHM 3 (POLMODSYS).

Input: polynomials $\mathfrak{M} = (\mathbf{m}_1, \dots, \mathbf{m}_n) \in \mathbb{K}[X]_{\neq 0}^n$, a matrix $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$ with $\deg(\mathbf{F}_{*,j}) < \deg(\mathbf{m}_j)$, a shift $\mathbf{s} \in \mathbb{Z}^m$.

Output: the \mathbf{s} -Popov solution basis for $(\mathfrak{M}, \mathbf{F})$ and the \mathbf{s} -minimal degree δ of $(\mathfrak{M}, \mathbf{F})$.

1. If $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n) \leq m$:
 - a. Build $\mathbf{E} \in \mathbb{K}^{m \times \sigma}$ and $\mathbf{M} \in \mathbb{K}^{\sigma \times \sigma}$ as in Section 1.2
 - b. Return LINEARIZATIONMIB($\mathbf{E}, \mathbf{M}, \mathbf{s}, 2^{\lceil \log_2(\sigma) \rceil}$)
2. Else if $n = 1$: Return POLMODSYSONE($\mathbf{m}_1, \mathbf{F}, \mathbf{s}, 2\sigma$)
3. Else:
 - a. $\mathfrak{M}^{(1)}, \mathbf{F}^{(1)} \leftarrow (\mathbf{m}_1, \dots, \mathbf{m}_{\lfloor n/2 \rfloor}), \mathbf{F}_{*,1 \dots \lfloor n/2 \rfloor}$
 - $\mathfrak{M}^{(2)}, \mathbf{F}^{(2)} \leftarrow (\mathbf{m}_{\lfloor n/2 \rfloor + 1}, \dots, \mathbf{m}_n), \mathbf{F}_{*,\lfloor n/2 \rfloor + 1 \dots n}$
 - b. $\mathbf{P}^{(1)}, \delta^{(1)} \leftarrow \text{POLMODSYS}(\mathfrak{M}^{(1)}, \mathbf{F}^{(1)}, \mathbf{s})$
 - c. $\mathbf{R} \leftarrow \mathbf{P}^{(1)} \mathbf{F}^{(2)} \bmod \mathfrak{M}^{(2)}$
 - d. $\mathbf{P}^{(2)}, \delta^{(2)} \leftarrow \text{POLMODSYS}(\mathfrak{M}^{(2)}, \mathbf{R}, \text{rdeg}_{\mathbf{s}}(\mathbf{P}^{(1)}))$
 - e. $\mathbf{P} \leftarrow \text{KNOWNDEGPOLMODSYS}(\mathfrak{M}, \mathbf{F}, \mathbf{s}, \delta^{(1)} + \delta^{(2)})$
 - f. Return $(\mathbf{P}, \delta^{(1)} + \delta^{(2)})$

We have $\deg(\tilde{\mathbf{P}}) \leq \lceil \sigma/m \rceil \leq 2\sigma/m$. Since $|\text{cdeg}(\tilde{\mathbf{F}})| < \sigma$ and $n \leq m \leq \tilde{m} \leq 2m$, $\tilde{\mathbf{F}}$ can be partially linearized into $\mathcal{O}(m)$ columns of degree $\mathcal{O}(\sigma/m)$. Then, $\tilde{\mathbf{P}}\tilde{\mathbf{F}}$ is computed in $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ operations. The j -th column of $\tilde{\mathbf{P}}\tilde{\mathbf{F}}$ has $\tilde{m} \leq 2m$ rows and degree less than $\sigma_j + 2\sigma/m$: it can be reduced modulo \mathbf{m}_j in $\tilde{\mathcal{O}}(\sigma + m\sigma_j)$ operations [13, Chapter 9]; summing over $1 \leq j \leq n$ with $n \leq m$, this is in $\tilde{\mathcal{O}}(m\sigma)$. \square

PROOF OF THEOREM 1.4. The correctness and the cost $\tilde{\mathcal{O}}(m^{\omega-1}\sigma)$ for Steps 1 and 2 of Algorithm 3 follow from [20, Theorem 1.4] and Proposition 2.10. With the costs of Steps 3.c and 3.e given in Proposition 2.4 and Lemma 2.11, we obtain the announced cost bound.

Now, using notation in Step 3, suppose $\mathbf{P}^{(1)}$ and $\mathbf{P}^{(2)}$ are the \mathbf{s} - and $\text{rdeg}_{\mathbf{s}}(\mathbf{P}^{(1)})$ -Popov solution bases for $(\mathfrak{M}^{(1)}, \mathbf{F}^{(1)})$ and $(\mathfrak{M}^{(2)}, \mathbf{R})$. Then $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$ is a solution basis for $(\mathfrak{M}, \mathbf{F})$: if \mathbf{p} is a solution for $(\mathfrak{M}, \mathbf{F})$, it is one for $(\mathfrak{M}^{(1)}, \mathbf{F}^{(1)})$ and thus $\mathbf{p} = \lambda\mathbf{P}^{(1)}$ for some λ , and it is one for $(\mathfrak{M}^{(2)}, \mathbf{F}^{(2)})$ so that $\mathbf{p}\mathbf{F}^{(2)} = \lambda\mathbf{P}^{(1)}\mathbf{F}^{(2)} = \lambda\mathbf{R} = \mathbf{0} \bmod \mathfrak{M}^{(2)}$ and thus $\lambda = \mu\mathbf{P}^{(2)}$ for some μ ; then $\mathbf{p} = \mu\mathbf{P}^{(2)}\mathbf{P}^{(1)}$.

Then $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$ is an \mathbf{s} -minimal solution basis for $(\mathfrak{M}, \mathbf{F})$ and its \mathbf{s} -Popov form has \mathbf{s} -pivot degree $\delta^{(1)} + \delta^{(2)}$ [21, Section 3]. The correctness follows from Proposition 2.4. \square

3. FAST COMPUTATION OF THE SHIFTED POPOV FORM OF A MATRIX

3.1 Fast shifted Popov form algorithm

Our fast method for computing the \mathbf{s} -Popov form of a nonsingular $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ uses two steps, as follows.

1. Compute the Smith form of \mathbf{A} , giving the moduli \mathfrak{M} , and a corresponding right unimodular transformation, giving the equations \mathbf{F} , so that \mathbf{A} is a solution basis for $(\mathfrak{M}, \mathbf{F})$.
2. Find the \mathbf{s} -Popov solution basis for $(\mathfrak{M}, \mathbf{F})$.

We first show the correctness of this approach.

LEMMA 3.1. *Let $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ be nonsingular and $\mathbf{S} = \mathbf{UAV}$ be the Smith form of \mathbf{A} with \mathbf{U} and \mathbf{V} unimodular. Let $\mathfrak{M} \in \mathbb{K}[X]_{\neq 0}^m$ and $\mathbf{F} \in \mathbb{K}[X]^{m \times m}$ be such that $\mathbf{S} = \text{diag}(\mathfrak{M})$ and $\mathbf{F} = \mathbf{V} \bmod \mathfrak{M}$. Then \mathbf{A} is a solution basis for $(\mathfrak{M}, \mathbf{F})$.*

PROOF. Let $\mathbf{p} \in \mathbb{K}[X]^{1 \times m}$. If \mathbf{p} is in the row space of \mathbf{A} then \mathbf{p} is a solution for $(\mathfrak{M}, \mathbf{F})$ since $\mathbf{AV} = \mathbf{U}^{-1}\mathbf{S}$ with \mathbf{U}^{-1} over $\mathbb{K}[X]$. Now if $\mathbf{p}\mathbf{F} = \mathbf{0} \bmod \mathfrak{M}$, then $\mathbf{p}\mathbf{V} = \mathbf{q}\mathbf{S}$ for some \mathbf{q} and $\mathbf{p} = \mathbf{q}\mathbf{U}\mathbf{A}$ is in the row space of \mathbf{A} . \square

Concerning the cost of Step 1, such \mathfrak{M} and \mathbf{F} can be obtained in expected $\tilde{\mathcal{O}}(m^\omega \deg(\mathbf{A}))$ operations, by computing

- 1.a \mathbf{R} a row reduced form of \mathbf{A} [16, Theorem 18],
- 1.b $\text{diag}(\mathfrak{M})$ the Smith form of \mathbf{R} [29, Algorithm 12],
- 1.c (\ast, \mathbf{F}) a reduced Smith transform for \mathbf{R} [15, Figure 3.2];

as in [15, Figure 6.1], Steps 1.b and 1.c should be performed in conjunction with the preconditioning techniques detailed in [23]. One may take for \mathfrak{M} only the nontrivial Smith factors, and for \mathbf{F} only the nonzero columns of the transform.

The product of the moduli in \mathfrak{M} is $\det(\mathbf{A})$ so that the sum of their degrees is $\deg(\det(\mathbf{A}))$. Then, according to Theorem 1.4, Step 2 of the algorithm outlined above costs $\tilde{\mathcal{O}}(m^{\omega-1} \deg(\det(\mathbf{A})))$ operations. Thus this algorithm solves Problem 1 in expected $\tilde{\mathcal{O}}(m^\omega \deg(\mathbf{A}))$ field operations.

3.2 Reducing to almost uniform degrees

In this subsection, we use the partial linearization techniques from [16, Section 6] to prove the following result.

PROPOSITION 3.2. *Let $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ be nonsingular and let $\mathbf{s} \in \mathbb{Z}^m$. With no field operation, one can build a nonsingular $\tilde{\mathbf{A}} \in \mathbb{K}[X]^{\tilde{m} \times \tilde{m}}$ and a shift $\mathbf{u} \in \mathbb{Z}^{\tilde{m}}$ such that $\tilde{m} \leq 3m$, $\deg(\tilde{\mathbf{A}}) \leq \lceil \sigma(\mathbf{A})/m \rceil$, and the \mathbf{s} -Popov form of \mathbf{A} is the principal $m \times m$ submatrix of the \mathbf{u} -Popov form of $\tilde{\mathbf{A}}$.*

With the algorithm in the previous subsection, this implies Theorem 1.3. In the specific case of Hermite form computation, for which there is a deterministic algorithm with cost bound $\tilde{\mathcal{O}}(m^\omega \deg(\mathbf{A}))$ [35], one can verify that this leads to a deterministic algorithm using $\tilde{\mathcal{O}}(m^\omega \lceil \sigma(\mathbf{A})/m \rceil)$ operations. (However, for $\mathbf{s} = \mathbf{0}$ this does not give a $\tilde{\mathcal{O}}(m^\omega \lceil \sigma(\mathbf{A})/m \rceil)$ deterministic algorithm for the Popov form using [16, 28], since the corresponding \mathbf{u} is $(\mathbf{0}, t, \dots, t)$ with $t \geq \deg(\mathbf{A})$.)

DEFINITION 3.3 (COLUMN PARTIAL LINEARIZATION). *Let $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ and $\delta = (\delta_i)_i \in \mathbb{N}^m$. Then let $\delta = 1 + \lfloor (\delta_1 + \dots + \delta_m)/m \rfloor$, let $\alpha_i \geq 1$ and $0 \leq \beta_i < \delta$ be such that $\delta_i = (\alpha_i - 1)\delta + \beta_i$ for $1 \leq i \leq m$, let $\tilde{m} = \alpha_1 + \dots + \alpha_m$, and let $\mathcal{E} = [\mathbf{I} | \mathbf{E}^T]^T \in \mathbb{K}[X]^{\tilde{m} \times m}$ be the expansion-compression matrix with \mathbf{I} the identity matrix and*

$$\mathbf{E} = \begin{bmatrix} X^\delta & & & & & \\ & \ddots & & & & \\ & & X^{(\alpha_1-1)\delta} & & & \\ & & & \ddots & & \\ & & & & X^\delta & \\ & & & & & \ddots \\ & & & & & & X^{(\alpha_m-1)\delta} \end{bmatrix}. \quad (4)$$

The column partial linearization $\mathcal{L}_\delta^c(\mathbf{A}) \in \mathbb{K}[X]^{\tilde{m} \times \tilde{m}}$ of \mathbf{A} is defined as follows:

- the first m rows of $\mathcal{L}_\delta^c(\mathbf{A})$ form the unique matrix $\tilde{\mathbf{A}} \in \mathbb{K}[X]^{m \times \tilde{m}}$ such that $\mathbf{A} = \tilde{\mathbf{A}}\mathcal{E}$ and $\tilde{\mathbf{A}}$ has all columns of degree less than δ except possibly those at indices $m + (\alpha_1 - 1) + \dots + (\alpha_i - 1)$ for $1 \leq i \leq m$,
- for $1 \leq i \leq m$, the row $m + (\alpha_1 - 1) + \dots + (\alpha_{i-1} - 1) + 1$ of $\mathcal{L}_\delta^c(\mathbf{A})$ is $[0, \dots, 0, -X^\delta, 0, \dots, 0, 1, 0, \dots, 0]$ where $-X^\delta$ is at index i and 1 is on the diagonal,

- for $1 \leq i \leq m$ and $2 \leq j \leq \alpha_i - 1$, the row $m + (\alpha_1 - 1) + \dots + (\alpha_{i-1} - 1) + j$ of $\mathcal{L}_\delta^c(\mathbf{A})$ is $[0, \dots, 0, -X^\delta, 1, 0, \dots, 0]$ where 1 is on the diagonal.

Defining the row partial linearization $\mathcal{L}_\delta^r(\mathbf{A})$ of \mathbf{A} similarly, both linearizations are related by $\mathcal{L}_\delta^r(\mathbf{A}) = \mathcal{L}_\delta^c(\mathbf{A}^\top)^\top$.

Now we show that for a well-chosen \mathbf{u} , one can directly read the \mathbf{s} -Popov form of \mathbf{A} as a submatrix of the \mathbf{u} -Popov form of $\mathcal{L}_\delta^r(\mathbf{A})$ (resp. $\mathcal{L}_\delta^c(\mathbf{A})$).

LEMMA 3.4. Let $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ be nonsingular, $\mathbf{s} \in \mathbb{Z}^m$, \mathbf{P} be the \mathbf{s} -Popov form of \mathbf{A} , and $\delta \in \mathbb{N}^m$. We have that:

- (i) if \tilde{m} is the dimension of $\mathcal{L}_\delta^r(\mathbf{A})$ and $\mathbf{u} = (\mathbf{s}, t, \dots, t)$ is in $\mathbb{Z}^{\tilde{m}}$ with $t \geq \max(\mathbf{s}) + \deg(\mathbf{P})$, then the \mathbf{u} -Popov form of $\mathcal{L}_\delta^r(\mathbf{A})$ is $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ * & \mathbf{I} \end{bmatrix}$;
- (ii) if \tilde{m} is the dimension of $\mathcal{L}_\delta^c(\mathbf{A})$, \mathbf{E} is as in (4), and $\mathbf{u} = (\mathbf{s}, \mathbf{t}) \in \mathbb{Z}^{\tilde{m}}$ for any $\mathbf{t} \in \mathbb{Z}^{\tilde{m}-m}$, then the \mathbf{u} -Popov form of $\mathcal{L}_\delta^c(\mathbf{A})$ is $\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{E} & \mathbf{I} \end{bmatrix}$ is $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$;
- (iii) if \tilde{m} is the dimension of $\mathcal{L}_\delta^c(\mathbf{A})$ and $\mathbf{u} = (\mathbf{s}, t, \dots, t)$ is in $\mathbb{Z}^{\tilde{m}}$ with $t \geq \max(\mathbf{s}) + \deg(\mathbf{P})$, then the \mathbf{u} -Popov form of $\mathcal{L}_\delta^c(\mathbf{A})$ is $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ * & \mathbf{I} \end{bmatrix}$.

PROOF. (i) $\mathcal{L}_\delta^r(\mathbf{A})$ is left-unimodularly equivalent to $\begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{B} & \mathbf{I} \end{bmatrix}$ for some $\mathbf{B} \in \mathbb{K}[X]^{(\tilde{m}-m) \times m}$ [16, Theorem 10 (i)]. Then, let \mathbf{R} be the remainder of \mathbf{B} modulo \mathbf{P} , that is, the unique matrix in $\mathbb{K}[X]^{(\tilde{m}-m) \times m}$ which has column degree bounded by the column degree of \mathbf{P} componentwise and such that $\mathbf{R} = \mathbf{B} + \mathbf{Q}\mathbf{P}$ for some matrix \mathbf{Q} (see for example [22, Theorem 6.3-15], noting that \mathbf{P} is $\mathbf{0}$ -column reduced).

Let \mathbf{W} denote the unimodular matrix such that $\mathbf{P} = \mathbf{W}\mathbf{A}$. Then, $\begin{bmatrix} \mathbf{W} & \mathbf{0} \\ \mathbf{Q}\mathbf{W} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{B} & \mathbf{I} \end{bmatrix} = \begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{R} & \mathbf{I} \end{bmatrix}$ is left-unimodularly equivalent to $\mathcal{L}_\delta^r(\mathbf{A})$. Besides, since $\deg(\mathbf{R}) < \deg(\mathbf{P})$, we have that $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{R} & \mathbf{I} \end{bmatrix}$ is in \mathbf{u} -Popov form by choice of t .

(ii) The matrix $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$ is obviously in \mathbf{u} -Popov form: it remains to prove that it is left-unimodularly equivalent to $\mathcal{L}_\delta^c(\mathbf{A}) \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{E} & \mathbf{I} \end{bmatrix}$. Let \mathbf{T} denote the trailing principal submatrix $\mathbf{T} = \mathcal{L}_\delta^c(\mathbf{A})_{m+1 \dots \tilde{m}, m+1 \dots \tilde{m}}$, and let \mathbf{W} be the unimodular matrix such that $\mathbf{W}\mathbf{P} = \mathbf{A}$. Then, \mathbf{T} is unit lower triangular, thus unimodular, and by construction of $\mathcal{L}_\delta^c(\mathbf{A})$, for some matrix \mathbf{B} we have $\mathcal{L}_\delta^c(\mathbf{A}) \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{E} & \mathbf{I} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{0} & \mathbf{T} \end{bmatrix} = \begin{bmatrix} \mathbf{W} & \mathbf{0} \\ \mathbf{0} & \mathbf{T} \end{bmatrix} \begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$.

(iii) From (ii), $\mathcal{L}_\delta^c(\mathbf{A})$ is left-unimodularly equivalent to $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{E} & \mathbf{I} \end{bmatrix} = \begin{bmatrix} \mathbf{P} & \mathbf{0} \\ -\mathbf{E} & \mathbf{I} \end{bmatrix}$. Using arguments in the proof of (i) above, by choice of t the \mathbf{u} -Popov form of $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ -\mathbf{E} & \mathbf{I} \end{bmatrix}$ is $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{R} & \mathbf{I} \end{bmatrix}$ with \mathbf{R} the remainder of $-\mathbf{E}$ modulo \mathbf{P} . \square

In the usual case where $\deg(\mathbf{P})$ is not known *a priori*, one may choose t using the inequality $\deg(\mathbf{P}) \leq \deg(\det(\mathbf{P})) = \deg(\det(\mathbf{A})) \leq m \deg(\mathbf{A})$.

This result implies Proposition 3.2 thanks to the following remark from [16]. Let π_1, π_2 be permutation matrices such that $\mathbf{B} = \pi_1 \mathbf{A} \pi_2 = [b_{i,j}]_{i,j}$ satisfies $\deg(b_{i,i}) \geq \deg(b_{j,k})$ for all $j, k \geq i$ and $1 \leq i \leq m$. Defining $\mathbf{d} = (d_i)_i \in \mathbb{N}^m$ by $d_i = \overline{\deg(b_{i,i})} = \begin{cases} \deg(b_{i,i}) & \text{if } b_{i,i} \neq 0 \\ 0 & \text{otherwise} \end{cases}$, we have $d_1 + \dots + d_m \leq \sigma(\mathbf{A})$ by definition of $\sigma(\mathbf{A})$ in (1). Let $\delta = \pi_1^{-1} \mathbf{d}$, where \mathbf{d} is seen as a column vector, and $\gamma = \text{cdeg}(\mathcal{L}_\delta^r(\mathbf{A}))$. Then the matrix $\tilde{\mathbf{A}} = \mathcal{L}_\gamma^r(\mathcal{L}_\delta^r(\mathbf{A}))$ is $\tilde{m} \times \tilde{m}$ with $\tilde{m} < 3m$, and we have $\deg(\tilde{\mathbf{A}}) \leq \lceil \sigma(\mathbf{A})/m \rceil$ [16, Corollary 3]. Lemma 3.4 further shows that the \mathbf{s} -Popov form of \mathbf{A} is the principal $m \times m$ submatrix of the \mathbf{u} -Popov form of $\tilde{\mathbf{A}}$, for the shift $\mathbf{u} = (\mathbf{s}, t, \dots, t) \in \mathbb{Z}^{\tilde{m}}$ with $t = \max(\mathbf{s}) + m \deg(\mathbf{A})$.

Acknowledgements. The author sincerely thanks the anonymous reviewers for their careful reading and detailed comments, which were very helpful for preparing the final version of this paper. He also thanks C.-P. Jeannerod, G. Labahn, É. Schost, A. Storjohann, and G. Villard for their interesting and useful comments. The author gratefully acknowledges financial support provided through the international mobility grants *Explo'ra Doc from Région Rhône-Alpes, PALSE*, and *Mitacs Globalink - Inria*.

4. REFERENCES

- [1] M. Alekhovich. Linear Diophantine equations over polynomials and soft decoding of Reed-Solomon codes. In *FOCS'02*, pages 439–448. IEEE, 2002.
- [2] M. Alekhovich. Linear Diophantine equations over polynomials and soft decoding of Reed-Solomon codes. *IEEE Trans. Inf. Theory*, 51(7):2257–2265, July 2005.
- [3] B. Beckermann. A reliable method for computing M-Padé approximants on arbitrary staircases. *J. Comput. Appl. Math.*, 40(1):19–42, 1992.
- [4] 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.
- [5] B. Beckermann and G. Labahn. Fraction-free computation of matrix rational interpolants and matrix gcds. *SIAM J. Matrix Anal. Appl.*, 22(1):114–144, 2000.
- [6] B. Beckermann, G. Labahn, and G. Villard. Normal forms for general polynomial matrices. *J. Symbolic Comput.*, 41(6):708–737, 2006.
- [7] P. Busse. *Multivariate List Decoding of Evaluation Codes with a Gröbner Basis Perspective*. PhD thesis, University of Kentucky, 2008.
- [8] M. Chowdhury, C.-P. Jeannerod, V. Neiger, É. Schost, and G. Villard. Faster algorithms for multivariate interpolation with multiplicities and simultaneous polynomial approximations. *IEEE Trans. Inf. Theory*, 61(5):2370–2387, 2015.
- [9] H. Cohn and N. Heninger. Approximate common divisors via lattices. In *Tenth Algorithmic Number Theory Symposium*, pages 271–293. Mathematical Sciences Publishers (MSP), 2012-2013.
- [10] H. Cohn and N. Heninger. Ideal forms of Coppersmith's theorem and Guruswami-Sudan list decoding. *Advances in Mathematics of Communications*, 9(3):311–339, 2015.
- [11] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Comput.*, 9(3):251–280, 1990.
- [12] C. Devet, I. Goldberg, and N. Heninger. Optimally robust private information retrieval. *Cryptology ePrint Archive*, Report 2012/083, 2012.
- [13] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra (third edition)*. Cambridge University Press, 2013.
- [14] P. Giorgi, C.-P. Jeannerod, and G. Villard. On the complexity of polynomial matrix computations. In *ISSAC'03*, pages 135–142. ACM, 2003.
- [15] S. Gupta. Hermite forms of polynomial matrices. Master's thesis, University of Waterloo, 2011.
- [16] S. Gupta, S. Sarkar, A. Storjohann, and J. Valeriote. Triangular x -basis decompositions and derandomization of linear algebra algorithms over $K[x]$. *J. Symbolic Comput.*, 47(4):422–453, 2012.
- [17] S. Gupta and A. Storjohann. Computing Hermite forms of polynomial matrices. *ISSAC'11*, pages 155–162. ACM, 2011.
- [18] J. L. Hafner and K. McCurley. Asymptotically fast triangularization of matrices over rings. *SIAM Journal on Computing*, 20(6):1068–1083, 1991.
- [19] C. Hermite. Sur l'introduction des variables continues dans la théorie des nombres. *Journal für die reine und angewandte Mathematik*, 41:191–216, 1851.
- [20] C.-P. Jeannerod, V. Neiger, É. Schost, and G. Villard. Computing minimal interpolation bases. HAL Open archive - <https://hal.inria.fr/hal-01241781>, 2015.
- [21] C.-P. Jeannerod, V. Neiger, É. Schost, and G. Villard. Fast computation of minimal interpolation bases in Popov form for arbitrary shifts. HAL Open archive - <https://hal.inria.fr/hal-01265983>, 2016.
- [22] T. Kailath. *Linear Systems*. Prentice-Hall, 1980.
- [23] E. Kaltofen, M.S. Krishnamoorthy, and D. Saunders. Parallel algorithms for matrix normal forms. *Linear Algebra Appl.*, 136:189–208, 1990.
- [24] S. Lang. *Algebra (Revised Third Edition)*. Springer-Verlag New-York Inc., 2002.
- [25] F. Le Gall. Powers of tensors and fast matrix multiplication. In *ISSAC'14*, pages 296–303. ACM, 2014.
- [26] T. Mulders and A. Storjohann. On lattice reduction for polynomial matrices. *J. Symbolic Comput.*, 35:377–401, 2003.
- [27] V. M. Popov. Invariant description of linear, time-invariant controllable systems. *SIAM Journal on Control*, 10(2):252–264, 1972.
- [28] S. Sarkar and A. Storjohann. Normalization of row reduced matrices. In *ISSAC'11*, pages 297–304. ACM, 2011.
- [29] A. Storjohann. High-order lifting and integrality certification. *J. Symbolic Comput.*, 36(3-4):613–648, 2003.
- [30] A. Storjohann. Notes on computing minimal approximant bases. In *Dagstuhl Seminar Proceedings*, 2006.
- [31] A. Storjohann and G. Labahn. Asymptotically fast computation of Hermite normal forms of integer matrices. *ISSAC'96*, pages 259–266. ACM, 1996.
- [32] M. Van Barel and A. Bultheel. A general module theoretic framework for vector M-Padé and matrix rational interpolation. *Numer. Algorithms*, 3:451–462, 1992.
- [33] G. Villard. Computing Popov and Hermite forms of polynomial matrices. *ISSAC'96*, pages 250–258. ACM, 1996.
- [34] W. Zhou and G. Labahn. Efficient algorithms for order basis computation. *J. Symbolic Comput.*, 47(7):793–819, 2012.
- [35] W. Zhou and G. Labahn. A fast, deterministic algorithm for computing a Hermite normal form of a polynomial matrix. *arXiv e-Print archive* - <http://arxiv.org/abs/1602.02049>, 2016.
- [36] W. Zhou, G. Labahn, and A. Storjohann. Computing minimal nullspace bases. In *ISSAC'12*, pages 366–373. ACM, 2012.
- [37] W. Zhou, G. Labahn, and A. Storjohann. A deterministic algorithm for inverting a polynomial matrix. *J. Complexity*, 31(2):162–173, 2015.