



HAL
open science

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. 2016. hal-01266014v1

HAL Id: hal-01266014

<https://inria.hal.science/hal-01266014v1>

Preprint submitted on 1 Feb 2016 (v1), last revised 12 May 2016 (v2)

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

Abstract

We give a Las Vegas algorithm which computes the shifted Popov form of a nonsingular polynomial matrix $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ in expected $\mathcal{O}^\sim(m^\omega \lceil \sigma(\mathbf{A})/m \rceil) \subseteq \mathcal{O}^\sim(m^\omega \deg(\mathbf{A}))$ operations in \mathbb{K} , where $\deg(\mathbf{A})$ is the degree of \mathbf{A} , $\sigma(\mathbf{A})$ is some quantity such that $\sigma(\mathbf{A})/m$ is bounded from above by both the average row degree and the average column degree of \mathbf{A} , ω is the exponent of matrix multiplication, and $\mathcal{O}^\sim(\cdot)$ indicates that logarithmic factors are omitted. This improves upon the cost bound of the fastest known algorithms for row reduction and Hermite form computation, which are deterministic. This is the first algorithm for shifted row reduction with cost bound $\mathcal{O}^\sim(m^\omega \deg(\mathbf{A}))$ for an arbitrary shift.

This algorithm uses partial linearization to reduce to the case $\deg(\mathbf{A}) \leq \sigma(\mathbf{A})$, and builds a system of modular equations whose solution set is the row space of \mathbf{A} . It remains to find the basis in shifted Popov form of this solution set: we give a deterministic algorithm for this problem in $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ 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 reduced [21, 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 [21, Section 6.7.2]. Given a matrix \mathbf{A} , the computation of a reduced form and of the Popov form of \mathbf{A} has received a lot of attention [25, 21, 6, 30, 24, 14, 26, 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 [1, 7], robust Private Information Retrieval [12], and more generally in polynomial versions of the Coppersmith method [9, 10]. In particular, there has been progress recently on the fast computation of the Hermite form [17, 15, 32], which is a specific shifted Popov form. The case of an arbitrary shift has been studied in [5].

For a *shift* $\mathbf{s} = (s_1, \dots, s_n) \in \mathbb{Z}^n$, the *\mathbf{s} -degree* of a row $\mathbf{p} = [p_1, \dots, p_n] \in \mathbb{K}[X]^{1 \times n}$ 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 $\mathbf{P} = [p_{ij}]_{ij}$ 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_{ij} .

Now, we assume that $m \leq n$ and \mathbf{P} has full rank. Then, \mathbf{P} is said to be *\mathbf{s} -reduced* [21, 5] if $\text{lm}_{\mathbf{s}}(\mathbf{P})$ has full rank. For a full rank $\mathbf{A} \in \mathbb{K}[X]^{m \times n}$ and a shift \mathbf{s} , 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.

In particular, \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* [21, 4, 5], 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 of a row). *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 column index $j \in \{1, \dots, n\}$ such that $\text{rdeg}_{\mathbf{s}}(\mathbf{p}) = \deg(p_j) + s_j$. Then, p_j and $\deg(p_j)$ are called the \mathbf{s} -pivot entry and the \mathbf{s} -pivot degree of \mathbf{p} , respectively.*

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, we focus on computing shifted Popov forms of *square nonsingular matrices*; for the general case, studied in [5], a fast solution would require further developments. (We still need the definition in the rectangular case for nullspace bases in Section 2.) 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. Here is our main problem.

Problem 1 (Shifted Popov normal form).

Input:

- the base field \mathbb{K} ,
- the dimension m ,
- 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 particular cases are the Popov form for the *uniform* shift $\mathbf{s} = \mathbf{0}$ [25, 21], and the Hermite form [18, 21] for the shift $\mathbf{h} = (0, \delta, 2\delta, \dots, (m-1)\delta) \in \mathbb{N}^m$ with $\delta = m \deg(\mathbf{A})$ [5, Lemma 2.6]. For a broader perspective on shifted reduced forms, we refer the reader to [5].

For this kind of 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: $\mathcal{O}^\sim(m^\omega d)$ operations in \mathbb{K} . Here, ω is so that we can multiply $m \times m$ matrices in $\mathcal{O}(m^\omega)$ ring operations on any ring, the best known bound being $\omega < 2.38$ [11, 23]. For example, given a nonsingular $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$, in $\mathcal{O}^\sim(m^\omega \deg(\mathbf{A}))$ operations one can compute a $\mathbf{0}$ -reduced form of \mathbf{A} [14, 16], the $\mathbf{0}$ -Popov form of \mathbf{A} [26], and the Hermite form of \mathbf{A} [15, 32].

Nevertheless, in some cases $\deg(\mathbf{A})$ is significantly larger than the average degree of the entries of \mathbf{A} , in which case the cost $\mathcal{O}^\sim(m^\omega \deg(\mathbf{A}))$ may be unsatisfactory. Recently, for the computation of order bases, nullspace bases, interpolation bases, and matrix inversion [28, 31, 33, 34, 19, 20], fast algorithms do take into account the average row or column degree of the matrices rather than their degree. Here, in particular, we achieve a similar improvement for the computation of $\mathbf{0}$ -Popov and Hermite 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 sum 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 $\mathcal{O}^\sim(m^\omega \lceil \sigma(\mathbf{A})/m \rceil) \subseteq \mathcal{O}^\sim(m^\omega \deg(\mathbf{A}))$ operations in \mathbb{K} .*

We are mostly interested in the case $m \in \mathcal{O}(\sigma(\mathbf{A}))$; the cost bound above can then be written $\mathcal{O}^\sim(m^{\omega-1} \sigma(\mathbf{A}))$, which is both in $\mathcal{O}^\sim(m^{\omega-1} |\text{rdeg}(\mathbf{A})|)$ and $\mathcal{O}^\sim(m^{\omega-1} |\text{cdeg}(\mathbf{A})|)$.

Having $\sigma(\mathbf{A})$ small compared to m implies that \mathbf{A} has mostly constant entries, and the algorithm uses $\mathcal{O}^\sim(m^\omega)$ operations.

Previous work on $\mathbf{0}$ -reduction includes [6, 30, 1, 24, 14, 26, 16]. The fastest known algorithm for the $\mathbf{0}$ -Popov form is deterministic and has cost $\mathcal{O}^\sim(m^\omega \deg(\mathbf{A}))$; it computes first a $\mathbf{0}$ -reduced form [16], and then its $\mathbf{0}$ -Popov form via normalization [26]. Fast algorithms for the Hermite form have been given in [30, 17, 15, 32]; the cost $\mathcal{O}^\sim(m^\omega \deg(\mathbf{A}))$ was first achieved by a probabilistic algorithm in [15], and then deterministically in [32].

For an arbitrary \mathbf{s} , the algorithm in [5] is fraction-free and uses a number of operations that is, depending on \mathbf{s} , at least quintic in m and quadratic in $\deg(\mathbf{A})$.

A folklore solution is based on the fact that, assuming $\min(\mathbf{s}) \geq 0$, \mathbf{Q} is in \mathbf{s} -Popov form if and only if \mathbf{QD} is in $\mathbf{0}$ -Popov form, with $\mathbf{D} = \text{Diag}(X^{s_1}, \dots, X^{s_m})$. Then, this solution computes the $\mathbf{0}$ -Popov form \mathbf{P} of \mathbf{AD} using [16, 26], and returns \mathbf{PD}^{-1} . In general, this approach uses $\mathcal{O}^\sim(m^\omega(\deg(\mathbf{A}) + \max(\mathbf{s}) - \min(\mathbf{s})))$ operations, which is not satisfactory when $\max(\mathbf{s})$ is large. For example, its cost for computing the Hermite form is $\mathcal{O}^\sim(m^{\omega+2} \deg(\mathbf{A}))$. (This is the worst case since, as recalled in Appendix A, one can assume without loss of generality that $\max(\mathbf{s}) - \min(\mathbf{s}) \in \mathcal{O}(m^2 \deg(\mathbf{A}))$.)

Here we obtain, to the best of our knowledge, the best known cost bound $\mathcal{O}^\sim(m^\omega \lceil \sigma(\mathbf{A})/m \rceil) \subseteq \mathcal{O}^\sim(m^\omega \deg(\mathbf{A}))$ for Problem 1. In general, this removes the dependency in $\max(\mathbf{s}) - \min(\mathbf{s})$, which means in some cases a speedup by a factor m^2 . Besides, this is also an improvement for both cases $\mathbf{s} = \mathbf{0}$ and $\mathbf{s} = \mathbf{h}$ when \mathbf{A} has unbalanced degrees.

We note that, in order to obtain a cost which depends on $\sigma(\mathbf{A})$ rather than $\deg(\mathbf{A})$, when the initial problem is for the uniform shift we still reduce it to another instance of Problem 1 with a non-uniform shift which has large entries.

One of the main difficulties in row reduction algorithms is to control the size of the matrices, that is, the number of coefficients from \mathbb{K} needed for their dense representation. A first remark when dealing with arbitrary shifts is that, for some input (\mathbf{A}, \mathbf{s}) , the size of an \mathbf{s} -reduced form of \mathbf{A} may be beyond our target cost.

For example, consider $\mathbf{A} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{B} \end{bmatrix} \in \mathbb{K}[X]^{2m \times 2m}$ for any $\mathbf{0}$ -reduced $\mathbf{B} \in \mathbb{K}[X]^{m \times m}$. Then, for $\mathbf{s} = (0, \dots, 0, d, \dots, d)$ with $d > 0$, $\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{C} & \mathbf{B} \end{bmatrix}$ 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 $\deg(\mathbf{A})$.

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} .

Furthermore, it is well-known that the size of the unimodular transformation leading from \mathbf{A} to \mathbf{P} can be beyond the target cost: in particular, fast algorithms for $\mathbf{0}$ -reduction and Hermite form [14, 16, 17, 15, 32] do not directly perform unimodular transformations on \mathbf{A} to reduce its degrees.

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. In both cases, the mentioned algorithms take advantage of the shift in input.

In [14, Section 3], the fact that a $\mathbf{0}$ -reduced form of \mathbf{A} has degree at most $\deg(\mathbf{A})$ allows to compute it as an order basis with the uniform shift. We have no similar property for an arbitrary \mathbf{s} , even for the \mathbf{s} -Popov form of \mathbf{A} : its column degree may be any tuple whose entries add up to $\deg(\det(\mathbf{A}))$.

As for the Hermite form \mathbf{H} of \mathbf{A} , its triangular shape is used to pre-compute its \mathbf{h} -pivot degree [17, Section 2.2] [32, Section 3]: for an arbitrary \mathbf{s} , it is unclear to us how to pre-compute the \mathbf{s} -pivot degree of the output efficiently.

As in [17], our algorithm first computes the Smith form \mathbf{S} of \mathbf{A} along with 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: it remains to compute a basis of solutions in \mathbf{s} -Popov form.

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 solutions for $(\mathfrak{M}, \mathbf{F})$ 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 [22, 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 $(\mathfrak{M}, \mathbf{F})$* .

Furthermore, for example for the application to Problem 1, we are interested in such a \mathbf{P} which is \mathbf{s} -reduced, in which case \mathbf{P} is said to be an *\mathbf{s} -minimal solution basis for $(\mathfrak{M}, \mathbf{F})$* . In particular, the unique such basis which is in \mathbf{s} -Popov form is called the *\mathbf{s} -Popov solution basis for $(\mathfrak{M}, \mathbf{F})$* .

Problem 2 (Minimal solution basis).

Input:

- the base field \mathbb{K} ,
- dimensions m and n ,
- 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: an \mathbf{s} -minimal solution basis for $(\mathfrak{M}, \mathbf{F})$.

A particular case of this problem is Hermite-Padé approximation, for which the moduli are powers of X . More precisely, following [31, Definition 2.2], an *\mathbf{s} -order basis for \mathbf{F} and*

$(\sigma_1, \dots, \sigma_n)$ is an \mathbf{s} -minimal solution basis for $(\mathfrak{M}, \mathbf{F})$ with $\mathfrak{M} = (X^{\sigma_1}, \dots, X^{\sigma_n})$. When $\sigma_1 = \dots = \sigma_n$, fast algorithms have been studied in [3, 14, 28, 31].

More generally, when the moduli in \mathfrak{M} are products of known linear factors, Problem 2 is known as M-Padé approximation [2, 29, 4, 19]. For this problem, the algorithm given recently in [20] computes the \mathbf{s} -Popov solution basis using $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ operations, with $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n)$; it covers in particular the computation of the \mathbf{s} -Popov order basis for any $\sigma_1, \dots, \sigma_n$ [20, Theorem 1.4]. 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 $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ operations in \mathbb{K} , with $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n)$, and returns the \mathbf{s} -Popov solution basis for $(\mathfrak{M}, \mathbf{F})$.*

A similar cost bound for computing *one* solution of small \mathbf{s} -degree was obtained in [8, Theorem 2] with a probabilistic algorithm based on fast structured linear algebra over \mathbb{K} .

When the moduli are known as products of linear factors, one can rely on recurrence relations to solve the problem iteratively [2, 29, 3, 4], using at each step a multiplication by one of the linear factors. The fast algorithms in [3, 14, 31, 19, 20], beyond the techniques used to achieve efficiency, are essentially divide-and-conquer versions of the iterative solution and are thus based on the same recurrence relations.

However, for arbitrary moduli there are no such recurrence relations in general. Then, a key idea is to relate solution bases to nullspace bases: Problem 2 asks to find \mathbf{P} such that there is \mathbf{Q} with $[\mathbf{P}|\mathbf{Q}]\mathbf{N} = \mathbf{0}$ for $\mathbf{N} = [\mathbf{F}^\top] - \text{Diag}(\mathfrak{M})^\top$. More precisely, we will see that $[\mathbf{P}|\mathbf{Q}]$ can be obtained as a \mathbf{u} -minimal nullspace basis of \mathbf{N} for $\mathbf{u} = (\mathbf{s} - \min(\mathbf{s}), \mathbf{0}) \in \mathbb{N}^{m+n}$.

Still, this is not enough to solve Problem 2 efficiently: the algorithm in [33] performs this nullspace computation in $\mathcal{O}^\sim((m+n)^\omega(\sigma_{\max} + \eta))$ operations, where η is the average of the largest n entries of \mathbf{u} and $\sigma_{\max} = \max_j(\deg(\mathbf{m}_j))$. In particular, for $n = 1$ this is $\mathcal{O}^\sim(m^\omega(\sigma + \max(\mathbf{s}) - \min(\mathbf{s})))$.

Here, for $n = 1$, we solve this nullspace basis problem in $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ operations 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 obtained from an \mathbf{u} -Popov order basis for \mathbf{N} and order $\mathcal{O}(\sigma)$, computed in $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ using [20]. Now, when $\max(\mathbf{u})$ is large compared to σ , and assuming that \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 similar to those in [3, 14], although here the recursion is on n . Once two solution bases $\mathbf{P}^{(1)}$ and $\mathbf{P}^{(2)}$ in shifted Popov form have been obtained recursively, their product $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$ is an \mathbf{s} -minimal solution basis for $(\mathfrak{M}, \mathbf{F})$; however it may not be in \mathbf{s} -Popov form and may even have size beyond our target cost. 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 δ of \mathbf{P} . Using ideas similar to those in [20], and results therein, to obtain \mathbf{P} we give a fast algorithm for computing

a shifted Popov nullspace basis when we have *a priori* knowledge about the degrees and locations of the pivots in the output. Assuming $n \in \mathcal{O}(m)$, and using partial linearization to reduce the degrees in the output, this allows to recover \mathbf{P} from $\boldsymbol{\delta}$ using $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ operations.

We note that Problem 2 is a minimal interpolation basis problem [4, 19] when the so-called *multiplication matrix* \mathbf{M} is in Frobenius form, as detailed below. Previous algorithms in $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ for minimal interpolation bases are known in two cases: when \mathbf{M} is in Jordan form [20, Theorem 1.3], which corresponds to M-Padé approximation, and when $\sigma \in \mathcal{O}(m)$ for any \mathbf{M} [19, Theorem 1.4]. Here, we use the former case for order bases computations, and the latter case when $\sigma \in \mathcal{O}(m)$, which may arise in particular in recursive calls.

The correspondence is as follows: \mathbf{p} is a solution for $(\mathfrak{M}, \mathbf{F})$ if and only if \mathbf{p} is an *interpolant* for (\mathbf{E}, \mathbf{M}) [19, 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 to \mathbf{m}_j . In this context, the multiplication $\mathbf{p} \cdot \mathbf{E}$ defined by \mathbf{M} as in [4, 19] precisely corresponds to $\mathbf{p}\mathbf{F} \bmod \mathfrak{M}$.

2 Fast computation of the shifted Popov solution basis

In what follows, we call *s-minimal degree of $(\mathfrak{M}, \mathbf{F})$* the *s-pivot degree* $\boldsymbol{\delta}$ of the *s-Popov* solution basis for $(\mathfrak{M}, \mathbf{F})$. The sum of its entries $|\boldsymbol{\delta}|$ is at most $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n)$; this is classical for Hermite-Padé and M-Padé approximation (see for example [29, Theorem 4.1]), and this is proved for minimal interpolation bases in general in [19, Lemma 7.17].

2.1 Solution bases from nullspace bases

First, we 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 well-chosen $\mathbf{u} \in \mathbb{Z}^{m+n}$.

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 $\mathbf{w} \in \mathbb{Z}^n$ with $\max(\mathbf{w}) \leq \min(\mathbf{s})$. 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* $(1, \dots, m)$.*

Proof. We write $\mathbf{D} = \text{Diag}(\mathfrak{M})$, and $\mathbf{N} = [\mathbf{F}^\top | \mathbf{D}]^\top$.

First, we assume that \mathbf{P} is a solution basis for $(\mathfrak{M}, \mathbf{F})$. Then, $\mathbf{Q} = -\mathbf{P}\mathbf{F}\mathbf{D}^{-1} \in \mathbb{K}[X]^{m \times n}$ is such that $[\mathbf{P} | \mathbf{Q}]\mathbf{N} = 0$. Let us consider $[\mathbf{p} | \mathbf{q}] \in \mathbb{K}[X]^{1 \times m+n}$ in the nullspace of \mathbf{N} and show that it is a $\mathbb{K}[X]$ -linear combination of the rows of $[\mathbf{P} | \mathbf{Q}]$. Indeed, \mathbf{p} is a solution for $(\mathfrak{M}, \mathbf{F})$, so that $\mathbf{p} = \lambda\mathbf{P}$ for some $\lambda \in \mathbb{K}[X]^{1 \times m}$. Then, $-\mathbf{q}\mathbf{D} = \mathbf{p}\mathbf{F} = \lambda\mathbf{P}\mathbf{F} = -\lambda\mathbf{Q}\mathbf{D}$ yields $\mathbf{q} = \lambda\mathbf{Q}$, so that $[\mathbf{p} | \mathbf{q}] = \lambda[\mathbf{P} | \mathbf{Q}]$.

Now, we assume that there is $\mathbf{Q} \in \mathbb{K}[X]^{m \times n}$ such that $[\mathbf{P} | \mathbf{Q}]$ is a nullspace basis of \mathbf{N} . Then, $\mathbf{P}\mathbf{F} = 0 \bmod \mathfrak{M}$. We consider a solution \mathbf{p} for $(\mathfrak{M}, \mathbf{F})$ and show that it is a $\mathbb{K}[X]$ -linear combination of the rows of \mathbf{P} . Defining $\mathbf{q} = -\mathbf{p}\mathbf{F}\mathbf{D}^{-1}$, we have that $[\mathbf{p} | \mathbf{q}]\mathbf{N} = 0$. Thus, there is some $\lambda \in \mathbb{K}[X]^{1 \times m}$ such that $[\mathbf{p} | \mathbf{q}] = \lambda[\mathbf{P} | \mathbf{Q}]$, which gives $\mathbf{p} = \lambda\mathbf{P}$.

Finally, using $\deg(\mathbf{F}_{*,j}) < \deg(\mathbf{m}_j)$ we obtain 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 $\max(\mathbf{w}) \leq \min(\mathbf{s})$ this implies $\text{rdeg}_{\mathbf{w}}(\mathbf{q}) < \text{rdeg}_{\mathbf{s}}(\mathbf{p})$. 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 $\text{lm}_{\mathbf{u}}([\mathbf{P}|\mathbf{Q}]) = [\text{lm}_{\mathbf{s}}(\mathbf{P})|0]$. This implies that \mathbf{P} is in \mathbf{s} -Popov form if and only if $[\mathbf{P}|\mathbf{Q}]$ is in \mathbf{u} -Popov form with \mathbf{s} -pivot indices $\{1, \dots, m\}$. \square

2.2 Fast solution for known minimal degree

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

Lemma 2.2. *Let $\mathbf{s} \in \mathbb{Z}^{m+n}$, and $\mathbf{F} \in \mathbb{K}[X]^{m+n \times n}$ of full rank with column degree $(\sigma_1, \dots, \sigma_n)$. Let $\mathbf{B} \in \mathbb{K}[X]^{m \times m+n}$ be the \mathbf{s} -Popov nullspace basis for \mathbf{F} , (π_1, \dots, π_m) its \mathbf{s} -pivot index, $(\delta_1, \dots, \delta_m)$ its \mathbf{s} -pivot degree, and $\delta \geq \deg(\mathbf{B})$ a degree bound. Then, let $\mathbf{u} = (u_1, \dots, u_{m+n}) \in \mathbb{Z}_{\leq 0}^{m+n}$ with*

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

Let $\tau_j = \sigma_j + \delta + 1$, for $1 \leq j \leq n$, and let \mathbf{A} be the \mathbf{u} -Popov order basis for \mathbf{F} 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 \mathbf{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 $\deg(\mathbf{p}\mathbf{F}_{*,j}) < \sigma_j + \delta + 1 = \tau_j$, and from $\mathbf{p}\mathbf{F}_{*,j} = 0 \pmod{X^{\tau_j}}$ we obtain $\mathbf{p}\mathbf{F}_{*,j} = 0$, 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, $\mathbf{C}\mathbf{F} = 0 \pmod{(X_1^{\tau_1}, \dots, X_n^{\tau_n})}$ and since \mathbf{C} has its \mathbf{u} -pivot entries on the diagonal, it is \mathbf{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 \mathbf{u} -Popov form, hence $\mathbf{C} = \mathbf{A}$. \square

In particular, computing the \mathbf{s} -Popov nullspace basis \mathbf{B} , when its \mathbf{s} -pivot index, its \mathbf{s} -pivot degree, and $\delta \geq \deg(\mathbf{B})$ are known, can be done in $\mathcal{O}^{\sim}(m^{\omega-1}(\sigma + n\delta))$ with $\sigma = \sigma_1 + \dots + \sigma_n$ using the order basis algorithm in [20].

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 \mathbf{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 $\mathcal{O}^{\sim}(m^{\omega-1}(\sigma + n\sigma))$, which exceeds our target $\mathcal{O}^{\sim}(m^{\omega-1}\sigma)$.

One issue is that \mathbf{Q} has size $\mathcal{O}(mn\sigma)$ when \mathbf{P} has columns of large degree; yet in Problem 2 we are not interested in \mathbf{Q} . Then, we can find \mathbf{P} using the partial linearization approach in the following result, which consists in reducing the degrees of the columns of large degree of \mathbf{P} and is proved in general for interpolation bases in [20, Lemma 4.2].

Lemma 2.3. Let $\mathfrak{M} = (\mathbf{m}_1, \dots, \mathbf{m}_n) \in \mathbb{K}[X]_{\neq 0}^n$ of degrees $(\sigma_1, \dots, \sigma_n)$, $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$, and $\mathbf{s} \in \mathbb{Z}^m$. Furthermore, let $\boldsymbol{\delta} = (\delta_1, \dots, \delta_m)$ denote the \mathbf{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 $\tilde{\boldsymbol{\delta}} \in \mathbb{N}^{\tilde{m}}$ as

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

and the expansion-compression matrix $\mathcal{E} \in \mathbb{K}[X]^{\tilde{m} \times m}$ as in [20, Equation (3)]. Let $\mathbf{d} = -\tilde{\boldsymbol{\delta}} \in \mathbb{Z}^{\tilde{m}}$ and $\mathbf{P} \in \mathbb{K}[X]^{\tilde{m} \times \tilde{m}}$ be the \mathbf{d} -Popov solution basis for $(\mathfrak{M}, \mathcal{E}\mathbf{F} \bmod \mathfrak{M})$.

Then, \mathbf{P} has \mathbf{d} -pivot degree $\tilde{\boldsymbol{\delta}}$ and the \mathbf{s} -Popov solution basis for $(\mathfrak{M}, \mathbf{F})$ is the submatrix of $\mathbf{P}\mathcal{E}$ formed by the rows at indices $\alpha_1 + \dots + \alpha_i$ for $1 \leq i \leq m$.

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 $\mathbf{s} \in \mathbb{Z}^m$,
- $\boldsymbol{\delta} = (\delta_1, \dots, \delta_m)$ the \mathbf{s} -minimal degree of $(\mathfrak{M}, \mathbf{F})$.

Output: the \mathbf{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 $1 \leq i \leq m$, $\tilde{m} \leftarrow \alpha_1 + \dots + \alpha_m$,
 $\tilde{\boldsymbol{\delta}} \in \mathbb{N}^{\tilde{m}} \leftarrow$ as in (2), $\mathcal{E} \in \mathbb{K}[X]^{\tilde{m} \times m}$ as in [20, Eq. (3)],
 $\tilde{\mathbf{F}} \leftarrow \mathcal{E}\mathbf{F} \bmod \mathfrak{M}$
2. $\mathbf{u} \leftarrow (-\tilde{\boldsymbol{\delta}}, -\delta - 1, \dots, -\delta - 1) \in \mathbb{Z}^{\tilde{m}+n}$
 $\boldsymbol{\tau} \leftarrow (\deg(\mathbf{m}_j) + \delta + 1)_{1 \leq j \leq n}$
3. $\tilde{\mathbf{P}} \leftarrow$ the \mathbf{u} -Popov order basis for $[\tilde{\mathbf{F}}^\top | \text{Diag}(\mathfrak{M})]^\top$ and $\boldsymbol{\tau}$
 $\mathbf{P} \leftarrow$ 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 $1 \leq i \leq m$

Lemma 2.4. The computation of $\mathcal{E}\mathbf{F} \bmod \mathfrak{M}$ at Step 1 of Algorithm 1 can be done in $\mathcal{O}(m\sigma)$ operations in \mathbb{K} .

Proof. The matrix \mathcal{E} has $\tilde{m} = \alpha_1 + \cdots + \alpha_m$ rows, with $\alpha_i = 1 + \lfloor \delta_i/\delta \rfloor \leq 1 + m\delta_i/\sigma$, hence $\tilde{m} \leq 2m$. Now, write $\mathbf{F} = [f_{ij}]_{ij}$ and let $j \in \{1, \dots, n\}$. The column $\mathcal{E}\mathbf{F}_{*,j} \bmod \mathbf{m}_j$ is formed by the m subcolumns $[X^{k\delta} f_{ij} \bmod \mathbf{m}_j]_{0 \leq k < \alpha_i}^\top$ for each $i \in \{1, \dots, m\}$. The entries of the i -th subcolumn are computed iteratively in a total of $\mathcal{O}(\alpha_i \sigma_j)$ operations using fast polynomial division with remainder [13, Chapter 9]. Summing these costs over i and j gives the conclusion. \square

Proposition 2.5. *Algorithm KNOWNDEGPOLMODSYS is correct. Writing $\sigma = \deg(\mathbf{m}_1) + \cdots + \deg(\mathbf{m}_n)$ and assuming $\sigma \geq m \geq n$, it uses $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ 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 follows from Lemma 2.3, and the cost bound follows from Proposition 2.2 and [20, Theorem 1.4], since $\tau_1 + \cdots + \tau_n = \sigma + n(1 + \lceil \sigma/m \rceil) \in \mathcal{O}(\sigma)$. \square

2.3 The case of one equation

Now focusing on $n = 1$, 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 2.6. *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 $\mathbf{A} = \begin{bmatrix} \mathbf{P} & \mathbf{q} \\ \mathbf{p} & q \end{bmatrix}$ 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 that \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 degree at most σ . 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 $\mathbf{s}_{[1:j]}$ of \mathbf{s} and $\mathbf{F}_{[1:j]}$ of \mathbf{F} , where j is such that $\text{amp}(\mathbf{s}_{[1:j]})$ is about half of $\text{amp}(\mathbf{s})$. This reveals $i \leq j$ entries $\delta_{[1:i]}$ of δ and the first i rows $\mathbf{P}_{[1:i,:]} = [\mathbf{P}^{(1)} | \mathbf{0}]$ of \mathbf{P} , with $\mathbf{P}^{(1)} \in \mathbb{K}[X]^{i \times i}$. A central point is that $\text{amp}(\mathbf{s}_{[i+1:]})$ is about half of $\text{amp}(\mathbf{s})$ as well, where $\mathbf{s}_{[i+1:]}$ is the tail of \mathbf{s} starting at the entry $i + 1$.

Then, knowing the degrees $\delta_{[1:i]}$ allows us to set up an order basis computation in order to obtain a *residual*; that is, a column $\mathbf{G} \in \mathbb{K}[X]^{m-i \times 1}$ and a modulus \mathbf{n} so that we can continue the computation of \mathbf{P} via a second recursive call which computes the $\mathbf{s}_{[i+1:]}$ -Popov solution basis for (\mathbf{n}, \mathbf{G}) . Finally, from these two bases obtained recursively we deduce the \mathbf{s} -pivot degree of \mathbf{P} , and then \mathbf{P} using Algorithm 1.

Now, we give 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 a tuple $\mathbf{t} = (t_1, \dots, t_m) \in \mathbb{Z}^m$, we write $\mathbf{t}_{[i:j]}$ for the subtuple of \mathbf{t} of its entries at indices $\{\pi_i, \dots, \pi_j\}$; 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 $\{\pi_k, \pi_{k+1}, \dots, \pi_l\}$. The main ideas of the rest of this subsection can be understood by focusing on the case of a non-decreasing \mathbf{s} , taking $\pi_i = i$ for all i .

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

Definition 2.7 (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 $d_{\mu_j} + t_{\mu_j} - t_{\mu_{i+1}} < 0$ for all $j \in \{1, \dots, i\}$.*

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

To find a splitting index, we use 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 build subtuples of \mathbf{s} which all contain a splitting index, as follows.

First, given $\alpha \in \mathbb{Z}_{>0}$, we let $\ell = 1 + \lfloor (\max(\mathbf{s}) - \min(\mathbf{s})) / \alpha \rfloor$, 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 choose $\alpha \geq 2\sigma$ and we assume $s_{\pi_{i+1}} - s_{\pi_i} \leq \sigma$ for each $1 \leq i < m$; this is without loss of generality as explained in 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 are splitting indices for (δ, \mathbf{s}) . Thus, we will 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 allows us to divide the computation into two subproblems, the first of which we have already solved by computing $\mathbf{P}^{(0)}$.

For $j \in \{2, \dots, m\}$, let $\mathbf{s}^{(0)} = \mathbf{s}_{[1:j]}$, $\mathbf{P}^{(0)}$ the $\mathbf{s}^{(0)}$ -Popov solution basis for $(\mathbf{m}, \mathbf{F}_{[1:j]})$, and $\delta^{(0)}$ its $\mathbf{s}^{(0)}$ -pivot degree. Suppose that there is a splitting index $i \leq j$ for $(\delta^{(0)}, \mathbf{s}^{(0)})$.

Lemma 2.8. 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 $\delta^{(1)}$ be its $\mathbf{s}^{(1)}$ -pivot degree. Then, i is a splitting index for (δ, \mathbf{s}) , and $[\mathbf{P}_{[i,i]} | \mathbf{P}_{[i,i+1]}] = [\mathbf{P}^{(1)} | \mathbf{0}]$.

Proof. Since i is a splitting index for $(\delta^{(0)}, \mathbf{s}^{(0)})$, we have $[\mathbf{P}_{[i,i]}^{(0)} | \mathbf{P}_{[i,i+1]}^{(0)}] = [\mathbf{Q} | \mathbf{0}]$ for some $\mathbf{Q} \in \mathbb{K}[X]^{i \times i}$. Then \mathbf{Q} is in $\mathbf{s}^{(1)}$ -Popov form and $\mathbf{Q}\mathbf{F}_{[i]} = \mathbf{0} \bmod \mathbf{m}$. Let $\mathbf{p} \in \mathbb{K}[X]^{1 \times i}$ be a solution for $(\mathbf{m}, \mathbf{F}_{[i]})$; we will prove that \mathbf{p} is a combination of the rows of \mathbf{Q} , so that $\mathbf{Q} = \mathbf{P}^{(1)}$. Let $\mathbf{q} \in \mathbb{K}[X]^{1 \times j}$ defined by $[\mathbf{q}_{[i,i]} | \mathbf{q}_{[i+1,i]}] = [\mathbf{p} | \mathbf{0}]$: \mathbf{q} is a solution for $(\mathbf{m}, \mathbf{F}_{[j]})$ and then $\mathbf{q} = \lambda \mathbf{P}^{(0)}$ for some $\lambda \in \mathbb{K}[X]^{1 \times j}$. Besides, since $\mathbf{P}^{(0)}$ is nonsingular, $\mathbf{P}_{[i,i+1]}^{(0)} = \mathbf{0}$ implies that $[\lambda_{[i,i]} | \lambda_{[i+1,i]}] = [\mu | \mathbf{0}]$ with $\mu \in \mathbb{K}[X]^{1 \times i}$. Thus, $\mathbf{p} = \mu \mathbf{Q}$.

Now, the matrix $\mathbf{B} \in \mathbb{K}[X]^{i \times m}$ with $[\mathbf{B}_{[i,i]} | \mathbf{B}_{[i,i+1]}] = [\mathbf{P}^{(1)} | \mathbf{0}]$ is in \mathbf{s} -Popov form and its rows are solutions for $(\mathfrak{M}, \mathbf{F})$. Thus by minimality of \mathbf{P} , $\mathbf{P}_{[i,:]}$ has \mathbf{s} -pivot degree at most $\delta^{(1)}$ componentwise. This implies that i is also a splitting index for (δ, \mathbf{s}) , so that $[\mathbf{P}_{[i,i]} | \mathbf{P}_{[i,i+1]}] = [\mathbf{R} | \mathbf{0}]$ for some $\mathbf{R} \in \mathbb{K}[X]^{i \times i}$, hence the conclusion. \square

The next two lemmas show that this knowledge of $\delta^{(1)}$ allows us to compute a so-called *residual*, from which we can continue the computation of \mathbf{P} .

Lemma 2.9. Let $\mathbf{s}^{(2)} = \mathbf{s}_{[i+1]}$, $\mathbf{d} = -\delta^{(1)} + \min(\mathbf{s}^{(2)}) - 2\sigma \in \mathbb{Z}^i$, $\mathbf{u} = (\mathbf{v}, \min(\mathbf{d})) \in \mathbb{Z}^{m+1}$ where $\mathbf{v}_{[i]} = \mathbf{d}$ and $\mathbf{v}_{[i+1]} = \mathbf{s}^{(2)}$. 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 , with $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ and $q \in \mathbb{K}[X]$. Then, $\deg(q) \geq \sigma$, $\mathbf{A}_{[i,i+1]} = \mathbf{0}$, $\mathbf{p}_{[i+1]} = \mathbf{0}$, and $[\mathbf{A}_{[i,i]} | \mathbf{q}_{[i]}] = [\mathbf{P}^{(1)} | \mathbf{q}^{(1)}]$ with $\mathbf{q}^{(1)} = -\mathbf{P}^{(1)}\mathbf{F}^{(1)}/\mathbf{m}$.

Proof. By choice of \mathbf{u} , we have $\deg(\mathbf{p}) \leq \deg(q)$. Since $\deg(\mathbf{F}) < \deg(\mathbf{m}) = \sigma$, the degree of $\mathbf{p}\mathbf{F} + q\mathbf{m}$ is $\deg(q) + \sigma$; then $\mathbf{p}\mathbf{F} + q\mathbf{m} = \mathbf{0} \bmod X^{2\sigma}$ implies that $\deg(q) + \sigma \geq 2\sigma$.

Besides, since \mathbf{A} is in \mathbf{u} -Popov form with sum of \mathbf{u} -pivot degrees at most the order 2σ , from $\min(\mathbf{s}^{(2)}) \geq \max(\mathbf{d}) + 2\sigma$ we obtain $\mathbf{A}_{[i,i+1]} = \mathbf{0}$ and $\mathbf{p}_{[i+1]} = \mathbf{0}$.

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

Thus, up to row and column permutations this order basis is $\begin{bmatrix} \mathbf{P}^{(1)} & \mathbf{0} & \mathbf{q}^{(1)} \\ * & \mathbf{P}^{(2)} & * \\ * & \mathbf{0} & q \end{bmatrix}$ 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 $\delta^{(2)}$ denote its $\mathbf{s}^{(2)}$ -pivot degree.

Lemma 2.10. Let $\mathbf{n} = X^{-2\sigma}(\mathbf{p}_{[i+1]} \mathbf{F}_{[i+1]} + q\mathbf{m}) \in \mathbb{K}[X]$ and $\mathbf{G} = X^{-2\sigma}(\mathbf{A}_{[i+1,i+1]} \mathbf{F}_{[i+1]} + \mathbf{q}_{[i+1]} \mathbf{m}) \in \mathbb{K}[X]^{m-i \times 1}$. Then $\deg(\mathbf{G}) < \deg(\mathbf{n}) \leq \sigma - |\delta^{(1)}| - |\delta^{(2)}|$. Let $\mathbf{P}^{(3)}$ the \mathbf{t} -Popov solution basis for (\mathbf{n}, \mathbf{G}) with $\mathbf{t} = \text{rdeg}_{\mathbf{s}^{(2)}}(\mathbf{P}^{(2)})$ and $\delta^{(3)}$ its \mathbf{t} -pivot degree, then $(\delta_{[i]}, \delta_{[i+1]}) = (\delta^{(1)}, \delta^{(2)} + \delta^{(3)})$.

Proof. The sum $|\delta^{(1)}| + |\delta^{(2)}| + \deg(q)$ 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, $\deg(\mathbf{n}) = \deg(q) - \sigma \leq \sigma - |\delta^{(1)}| - |\delta^{(2)}|$. On the other hand, we have $\deg(\mathbf{A}_{[i+1,i+1]}) <$

$|\delta^{(1)}| \leq \sigma$, $\deg(\mathbf{A}_{[i+1:,i+1:]}) \leq |\delta^{(2)}| \leq 2\sigma - |\delta^{(1)}| - \deg(q) \leq \sigma - |\delta^{(1)}| \leq \sigma$, and $\deg(\mathbf{q}_{[i+1:]}) < \deg(q)$; thus, $\deg(\mathbf{G}) < \deg(q) - \sigma = \deg(\mathbf{n})$.

Let $\mathbf{q}^{(3)} = -\mathbf{P}^{(3)}\mathbf{G}/\mathbf{n}$ and $t = \text{rdeg}_{\mathbf{u}}([\mathbf{p}|q])$; then, $t = \deg(q) - \max(\delta^{(1)}) + \min(\mathbf{s}^{(2)}) - 2\sigma \leq \min(\mathbf{s}^{(2)}) \leq \min(\mathbf{t})$. By Lemma 2.1, $[\mathbf{P}^{(3)}|\mathbf{q}^{(3)}]$ is the (\mathbf{t}, t) -Popov nullspace basis for $[\mathbf{G}^\top|\mathbf{n}]^\top$. Thus, defining $\mathbf{B} \in \mathbb{K}[X]^{m \times m}$ and $\mathbf{c} \in \mathbb{K}[X]^{m \times 1}$ by $\begin{bmatrix} \mathbf{B}_{[:i,:i]} & \mathbf{B}_{[:i,i+1:]} & \mathbf{c}_{[:i]} \\ \mathbf{B}_{[i+1:,i]} & \mathbf{B}_{[i+1:,i+1:]} & \mathbf{c}_{[i+1:]} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{P}^{(3)} & \mathbf{q}^{(3)} \end{bmatrix}$, by [33, Theorem 3.9], $\begin{bmatrix} \mathbf{B} & \mathbf{c} \\ \mathbf{A} & \mathbf{q} \end{bmatrix}$ is a \mathbf{u} -minimal nullspace basis of $[\mathbf{F}^\top|\mathbf{m}]^\top$. Lemma 2.1 implies that $\bar{\mathbf{P}} = \begin{bmatrix} \mathbf{B} & \mathbf{c} \\ \mathbf{A} & \mathbf{q} \end{bmatrix}$ is a \mathbf{v} -minimal solution basis for (\mathbf{m}, \mathbf{F}) .

It is easily checked that $\bar{\mathbf{P}}$ is in \mathbf{v} -Popov form, so that the \mathbf{v} -Popov form of $\bar{\mathbf{P}}$ is \mathbf{P} its \mathbf{v} -pivot degree is δ . Besides $\begin{bmatrix} \bar{\mathbf{P}}_{[:i,:i]} & \bar{\mathbf{P}}_{[:i,i+1:]} \\ \bar{\mathbf{P}}_{[i+1:,i]} & \bar{\mathbf{P}}_{[i+1:,i+1:]} \end{bmatrix} = \begin{bmatrix} \mathbf{P}^{(1)} & \mathbf{P}^{(1)} \\ \mathbf{P}^{(3)}\mathbf{A}_{2,1+\mathbf{q}^{(3)}}\mathbf{A}_{3,1} & \mathbf{P}^{(3)}\mathbf{P}^{(2)} \end{bmatrix}$, so that $(\delta_{[:i]}, \delta_{[i+1:]}) = (\delta^{(1)}, \delta^{(2)} + \delta^{(3)})$ [20, Lemmas 3.2 and 3.4]. \square

Here is the resulting algorithm. As an input, we have a parameter α as above, which dictates the amplitude of the subtuples used to partition the shift \mathbf{s} into sub-shifts which all contain a splitting index. As explained above, the initial call can be made with $\alpha = 2\sigma$.

Proposition 2.11. *Algorithm POLMODSYSONE is correct and uses $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ operations in \mathbb{K} .*

Proof. The correctness follows from the results in this subsection. By [20, Theorem 1.4], the order bases computations at Steps 1.a and 2.c use $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ operations. Using partial linearization as in Lemma 2.12 below, the computation of \mathbf{G} and \mathbf{n} at Step 2.c can be done in $\mathcal{O}^\sim(m^{\omega-1}\sigma)$; from Proposition 2.5, Step e uses $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ operations.

Running the algorithm with initial input $\alpha = 2\sigma$, the recursive tree has depth $\mathcal{O}(\log(\ell)) = \mathcal{O}(\log(\text{amp}(\mathbf{s})/2\sigma))$, which is in $\mathcal{O}(m^2)$. Besides, all sub-problems are for a modulus of degree at most σ , and on a given level of the tree, the sum of the dimensions of the column vector in input of each sub-problem is in $\mathcal{O}(m)$. Thus, since $a^{\omega-1} + b^{\omega-1} \leq (a+b)^{\omega-1}$ for $a, b > 0$, each level of the tree uses a total of $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ operations. \square

2.4 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 . From the two bases obtained recursively, it first deduces the \mathbf{s} -pivot degree of the sought basis \mathbf{P} and then uses this additional knowledge to compute \mathbf{P} with Algorithm 1.

When the sum $\sigma = \deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n)$ is in $\mathcal{O}(m)$, we use the algorithm LINEARIZATIONMIB from [19, Section 7].

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

Lemma 2.12. *Let $\mathfrak{M} = (\mathbf{m}_1, \dots, \mathbf{m}_n) \in \mathbb{K}[X]_{\neq 0}^n$ of degrees $(\sigma_1, \dots, \sigma_n)$, $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$, and $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$ with $m \geq n$ and $\deg(\mathbf{F}_{*,j}) < \sigma_j$. Let $\sigma \geq m$ such that $\sigma \geq \sigma_1 + \dots + \sigma_n$ and $|\text{cdeg}(\mathbf{P})| \leq \sigma$. Then, one can compute $\mathbf{P}\mathbf{F} \bmod \mathfrak{M}$ using $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ operations in \mathbb{K} .*

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 $\deg(\mathbf{F}) < \deg(\mathbf{m})$,
- 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 $\max(\mathbf{s}) - \min(\mathbf{s}) \leq 4\sigma$:

- a. $\mathbf{A} \leftarrow$ the $(\mathbf{s}, \min(\mathbf{s}))$ -Popov order basis for $[\mathbf{F}^\top | \mathbf{m}]^\top$ and 6σ ; return the principal $m \times m$ submatrix of \mathbf{A} and the degrees of its diagonal entries

2. Else:

- a. $(\mathbf{s}_1, \dots, \mathbf{s}_\ell) \leftarrow \text{PARTITION}(\mathbf{s}, \alpha)$,
 $j \leftarrow$ sum of the lengths of $\mathbf{s}_1, \dots, \mathbf{s}_{[\ell/2]}$, $\mathbf{s}^{(0)} \leftarrow \mathbf{s}_{[:j]}$,
 $(\mathbf{P}^{(0)}, \boldsymbol{\delta}^{(0)}) \leftarrow \text{POLMODSYSONE}(\mathbf{m}, \mathbf{F}_{[:j]}, \mathbf{s}^{(0)}, \alpha)$
- 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 $\mathbf{v}_{[:i]} = \mathbf{d}$ and $\mathbf{v}_{[i+1:]} = \mathbf{s}^{(2)}$, $\mathbf{u} = (\mathbf{v}, \min(\mathbf{d}))$
- c. $\begin{bmatrix} \mathbf{A} & \mathbf{q} \\ \mathbf{p} & \mathbf{q} \end{bmatrix} \leftarrow$ \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 $\mathbf{A}_{[i+1:, i+1:]}$,
 $\mathbf{G} \leftarrow X^{-2\sigma}(\mathbf{A}_{[i+1:, :]} \mathbf{F}_{[i+1:]} + \mathbf{q}_{[i+1:]} \mathbf{m})$,
 $\mathbf{n} \leftarrow X^{-2\sigma}(\mathbf{p}_{[i+1:]} \mathbf{F}_{[i+1:]} + \mathbf{q} \mathbf{m})$
- 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 $(\mathbf{P}, \boldsymbol{\delta})$

Proof. Using notation from Lemma 2.3, let $\tilde{\mathbf{P}} \in \mathbb{K}[X]^{m \times \tilde{m}}$ such that $\mathbf{P} = \tilde{\mathbf{P}}\mathcal{E}$ and $\deg(\tilde{\mathbf{P}}) < \lceil |\text{cdeg}(\mathbf{P})|/m \rceil$. Then, by Lemma 2.4 $\tilde{\mathbf{F}} = \mathcal{E}\mathbf{F} \bmod \mathfrak{M}$ can be computed using $\mathcal{O}(\tilde{m}\sigma)$ operations; we want to compute $\mathbf{P}\mathbf{F} \bmod \mathfrak{M} = \tilde{\mathbf{P}}\tilde{\mathbf{F}} \bmod \mathfrak{M}$.

Algorithm 3 (POLMODSYS).

Input:

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

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

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 $\mathcal{O}^{\sim}(m^{\omega-1}\sigma)$ 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 \mathfrak{m}_j in $\mathcal{O}^{\sim}(\sigma + m\sigma_j)$ operations [13, Chapter 9]; summing over $1 \leq j \leq n$ with $n \leq m$, this is in $\mathcal{O}^{\sim}(m\sigma)$. \square

Proof of Theorem 1.4. The correctness and the cost $\mathcal{O}^{\sim}(m^{\omega-1}\sigma)$ for Steps 1 and 2 follow from [19, Theorem 1.4] and Proposition 2.11. With the cost of Steps 3.c and 3.e in Proposition 2.5 and Lemma 2.12, 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 $\mathbf{p} = \boldsymbol{\lambda}\mathbf{P}^{(1)}$ for some $\boldsymbol{\lambda}$, and it is one for $(\mathfrak{M}^{(2)}, \mathbf{F}^{(2)})$ so that $\mathbf{p}\mathbf{F}^{(2)} = \boldsymbol{\lambda}\mathbf{P}^{(1)}\mathbf{F}^{(2)} = \boldsymbol{\lambda}\mathbf{R} = \mathbf{0} \bmod \mathfrak{M}^{(2)}$ and thus $\boldsymbol{\lambda} = \boldsymbol{\mu}\mathbf{P}^{(2)}$ for some $\boldsymbol{\mu}$; then $\mathbf{p} = \boldsymbol{\mu}\mathbf{P}^{(2)}\mathbf{P}^{(1)}$.

- the first m rows of $\mathcal{L}_\delta^c(\mathbf{A})$ form the unique matrix $\tilde{\mathbf{A}} \in \mathbb{K}[X]^{m \times \tilde{m}}$ which is such that $\mathbf{A} = \tilde{\mathbf{A}}\mathcal{E}$ and 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 at index $m + (\alpha_1 - 1) + \dots + (\alpha_{i-1} - 1) + 1$ of $\mathcal{L}_\delta^c(\mathbf{A})$ is $[0, \dots, 0, -X^\delta, 0, \dots, 0, 1, 0, \dots, 0]$ with the entry $-X^\delta$ at column index i and the entry 1 on the diagonal,
- for $1 \leq i \leq m$ and $2 \leq j \leq \alpha_i - 1$, the row at index $m + (\alpha_1 - 1) + \dots + (\alpha_{i-1} - 1) + j$ of $\mathcal{L}_\delta^c(\mathbf{A})$ is $[0, \dots, 0, -X^\delta, 1, 0, \dots, 0]$ with the entry 1 on the diagonal.

We will use the following properties from [16, Section 6].

Lemma 3.3. *Let $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ and $\delta \in \mathbb{N}^m$. Then,*

- $\mathcal{L}_\delta^c(\mathbf{A})$ is right-unimodularly equivalent to $\begin{bmatrix} \mathbf{A} & * \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$,
- the Smith form of $\mathcal{L}_\delta^c(\mathbf{A})$ is $\mathbf{S} = \text{Diag}(\mathbf{I}, \mathbf{S})$ where \mathbf{S} is the Smith form of \mathbf{A} .

Similarly, for a matrix $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ and $\delta \in \mathbb{N}^m$, the row partial linearization $\mathcal{L}_\delta^r(\mathbf{A})$ is defined as the transpose of $\mathcal{L}_\delta^c(\mathbf{A}^\top)$; translating the lemma above for the row partial linearization is straightforward.

3.3 Reducing to small average column degree

To obtain a cost bound involving $\sigma(\mathbf{A})$, we use in particular the following remark from [16]. Let π_1, π_2 be permutation matrices such that $\mathbf{B} = \pi_1 \mathbf{A} \pi_2 = [b_{ij}]_{i,j}$ satisfies

$$\deg(b_{ii}) \geq \deg(\mathbf{B}_{i\dots m, i\dots m}) \text{ for every } i \in \{1, \dots, m\}. \quad (4)$$

Then, we define $\mathbf{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$ by

$$d_i = \begin{cases} \deg(b_{ii}) & \text{if } b_{ii} \neq 0, \\ 0 & \text{otherwise;} \end{cases} \quad (5)$$

we have $d_1 + \dots + d_m = \sigma(\mathbf{A})$ by definition of $\sigma(\mathbf{A})$ in (1). We consider $\delta = \pi_1^{-1} \mathbf{d}$, where \mathbf{d} is seen as a column vector.

Now we show how to obtain the \mathbf{s} -Popov form of \mathbf{A} from the \mathbf{u} -Popov form of $\mathcal{L}_\delta^r(\mathbf{A})$ for a well-chosen \mathbf{u} . Besides, $\mathcal{L}_\delta^r(\mathbf{A})$ has dimension at most twice that of \mathbf{A} , and its average column degree is in $\mathcal{O}([\sigma(\mathbf{A})/m])$.

Lemma 3.4. *Let $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ nonsingular and $\mathbf{s} \in \mathbb{Z}^m$. Let $\delta \in \mathbb{N}^m$ as above and let \bar{m} be the dimension of $\mathcal{L}_\delta^r(\mathbf{A})$. Then, $\bar{m} < 2m$ and the sum of the column degrees of $\mathcal{L}_\delta^r(\mathbf{A})$ is less than $m + 2\sigma(\mathbf{A})$. Besides, let $\mathbf{u} = (\mathbf{s}, t, \dots, t) \in \mathbb{Z}^{\bar{m}}$ with $t = \max(\mathbf{s}) + m \deg(\mathbf{A})$ and \mathbf{P} the \mathbf{s} -Popov form of \mathbf{A} . Then the \mathbf{u} -Popov form of $\mathcal{L}_\delta^r(\mathbf{A})$ is $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ * & \mathbf{I} \end{bmatrix}$.*

Proof. By definition, $\alpha_i = 1 + \lfloor \delta_i/\delta \rfloor < 1 + \delta_i m/|\delta|$; summing over i we obtain $\bar{m} < 2m$. By construction, the first m columns of $\mathcal{L}_\delta^r(\mathbf{A})$ have column degree at most $\mathbf{d}\pi_2^{-1}$ componentwise (where \mathbf{d} is seen as a row vector), and its last $\bar{m} - m < m$ columns have degree $\delta \leq 1 + |\mathbf{d}|/m$. Thus, $|\text{cdeg}(\mathcal{L}_\delta^r(\mathbf{A}))| < |\mathbf{d}| + m\delta \leq m + 2|\mathbf{d}| = m + 2\sigma(\mathbf{A})$.

Lemma 3.3 states that there is $\mathbf{B} \in \mathbb{K}[X]^{\bar{m}-m \times m}$ such that $\mathcal{L}_\delta^r(\mathbf{A})$ is left-unimodularly equivalent to $\begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{B} & \mathbf{I} \end{bmatrix}$. Then, let \mathbf{R} be the remainder of \mathbf{B} modulo \mathbf{P} , that is, the unique matrix in $\mathbb{K}[X]^{\bar{m}-m \times m}$ 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 instance [21, 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 \mathbf{P} is in \mathbf{s} -Popov form we have $\deg(\mathbf{P}) \leq \deg(\det(\mathbf{P})) = \deg(\det(\mathbf{A})) \leq m \deg(\mathbf{A})$; hence $\deg(\mathbf{R}) < m \deg(\mathbf{A})$, and $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{R} & \mathbf{I} \end{bmatrix}$ is in \mathbf{u} -Popov form. \square

3.4 Reducing to uniform column degree

Now, we show how to combine column linearization and the reduction to Problem 2 in Subsection 3.1 in order to reduce to the case of a matrix \mathbf{A} which has column degree close to uniform; or in other words, with $\deg(\mathbf{A})$ close to the average column degree $|\text{cdeg}(\mathbf{A})|/m$.

Lemma 3.5. *Let $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ be nonsingular, let $\mathbf{s} \in \mathbb{Z}^m$, let $\delta = \text{cdeg}(\mathbf{A}) \in \mathbb{N}^m$ be the column degree of \mathbf{A} , let \tilde{m} be the dimension of $\mathcal{L}_\delta^c(\mathbf{A})$, let $\mathcal{E} = [\mathbf{I}|\mathbf{E}^\top]^\top$ as in (3), and consider the shift $\mathbf{u} = (\mathbf{s}, \mathbf{t}) \in \mathbb{Z}^{\tilde{m}}$ for any $\mathbf{t} \in \mathbb{Z}^{\tilde{m}-m}$. Then,*

- $m \leq \tilde{m} \leq 2m$ and $\deg(\mathcal{L}_\delta^c(\mathbf{A})) \leq 1 + \lfloor |\text{cdeg}(\mathbf{A})|/m \rfloor$,
- the \mathbf{u} -Popov form of $\mathcal{L}_\delta^c(\mathbf{A}) \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{E} & \mathbf{I} \end{bmatrix}$ is $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$, where \mathbf{P} is the \mathbf{s} -Popov form of \mathbf{A} ,
- let $\mathbf{S} = \mathbf{U}\mathcal{L}_\delta^c(\mathbf{A})\mathbf{V}$ be the Smith form of $\mathcal{L}_\delta^r(\mathbf{A})$, where \mathbf{U} and \mathbf{V} are unimodular. Writing $\mathbf{S} = \text{Diag}(\mathfrak{M})$ for some $\mathfrak{M} \in \mathbb{K}[X]_{\neq 0}^m$, let $\mathbf{F} \in \mathbb{K}[X]^{\tilde{m} \times n}$ be any matrix such that $\mathbf{F} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{E} & \mathbf{I} \end{bmatrix} \mathbf{V} \bmod \mathfrak{M}$. Then, the \mathbf{u} -Popov solution basis for $(\mathfrak{M}, \mathbf{F})$ is $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$.

Proof. The first item can be found in [16, Corollary 2]. Concerning the second item, the matrix $\begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$ is obviously in \mathbf{u} -Popov form: we want 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 $\mathcal{L}_\delta^c(\mathbf{A}) \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{E} & \mathbf{I} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{0} & \mathbf{T} \end{bmatrix} = \begin{bmatrix} \mathbf{W} & \mathbf{B} \\ \mathbf{0} & \mathbf{T} \end{bmatrix} \begin{bmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$. Now, the third item directly follows from the second and Lemma 3.1, since the Smith form of $\mathcal{L}_\delta^c(\mathbf{A}) \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{E} & \mathbf{I} \end{bmatrix}$ can be written $\mathbf{S} = \mathbf{U}\mathcal{L}_\delta^c(\mathbf{A}) \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{E} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{E} & \mathbf{I} \end{bmatrix} \mathbf{V}$. \square

3.5 Fast shifted Popov form algorithm

Here, we gather the results of the previous subsections to give a fast algorithm for Problem 1.

Proof of Theorem 1.3. According to Lemma 3.4, the \mathbf{s} -Popov form \mathbf{P} of \mathbf{A} is the principal $m \times m$ submatrix of the \mathbf{u} -Popov form of $\mathcal{L}_d^r(\mathbf{A})$ (as in Step 1). Then, Lemma 3.5 states that the \mathbf{u} -Popov form of $\mathcal{L}_d^r(\mathbf{A})$ is the principal $\bar{m} \times \bar{m}$ submatrix of $\tilde{\mathbf{P}}$ (as in Step 4), hence the correctness.

Furthermore, those two lemmas also imply that $\deg(\tilde{\mathbf{A}}) \leq 1 + \lfloor |\delta|/m \rfloor \leq 2(1 + \sigma(\mathbf{A})/m)$, and $\tilde{m} < 4m$. Then, the computation of \mathbf{S} and \mathbf{F} in Step 3 can be done in $\mathcal{O}(m^{\omega-1}\sigma(\mathbf{A}))$ operations using the algorithm in [15, Figure 3.2] with the preconditioning as in [15, Figure 6.1] (see also [17, Lemma 2]); this relies in particular on [27, Algorithm 12]. Then, since $\deg(\mathbf{m}_1) + \dots + \deg(\mathbf{m}_n) = \deg(\det(\mathbf{A}))$, by Lemma 2.4 and Theorem 1.4 Step 4 uses $\mathcal{O}(m^{\omega-1} \deg(\det(\mathbf{A})))$ operations, with $\deg(\det(\mathbf{A})) \leq \sigma(\mathbf{A})$. \square

Algorithm 4 (SHIFTEDPOPOVFORM).

Input:

- 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} .

1. $\pi_1, \pi_2 \in \mathbb{K}^{m \times m} \leftarrow$ permutation matrices such that $\mathbf{B} = \pi_1 \mathbf{A} \pi_2$ satisfies (4), $\mathbf{d} \leftarrow$ as in (5), $\mathbf{d} \leftarrow \pi_1^{-1} \mathbf{d}$, $\bar{m} \leftarrow$ dimension of $\mathcal{L}_d^r(\mathbf{A})$, $t \leftarrow \max(\mathbf{s}) + m \deg(\mathbf{A})$ and $\mathbf{u} \leftarrow (\mathbf{s}, t, \dots, t) \in \mathbb{Z}^{\bar{m}}$
2. $\delta \leftarrow \text{cdeg}(\mathcal{L}_d^r(\mathbf{A}))$, $\tilde{\mathbf{A}} \leftarrow \mathcal{L}_\delta^c(\mathcal{L}_d^r(\mathbf{A}))$, $\tilde{m} \leftarrow$ dimension of $\tilde{\mathbf{A}}$, $\mathcal{E} \leftarrow [\mathbf{I} | \mathbf{E}^T]^T$ as in (3), $\mathbf{v} \leftarrow (\mathbf{u}, 0, \dots, 0) \in \mathbb{Z}^{\tilde{m}}$
3. $\mathbf{S} \leftarrow$ the Smith form of $\tilde{\mathbf{A}}$, write $\mathbf{S} = \text{Diag}(1, \dots, 1, \mathfrak{M})$ with nonconstant polynomials in $\mathfrak{M} = (\mathbf{m}_1, \dots, \mathbf{m}_n)$, and compute $\mathbf{F} \in \mathbb{K}[X]^{\tilde{m} \times n}$ with $\deg(\mathbf{F}_{*,j}) < \deg(\mathbf{m}_j)$ and $[\mathbf{0} | \mathbf{F}] = \mathbf{V} \bmod \mathbf{S}$ for some right unimodular Smith transform \mathbf{V}
4. $\tilde{\mathbf{P}} \leftarrow \mathbf{v}$ -Popov solution basis for $(\mathfrak{M}, [\begin{smallmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{E} & \mathbf{I} \end{smallmatrix}] \mathbf{F} \bmod \mathfrak{M})$
5. Return the principal $m \times m$ submatrix of $\tilde{\mathbf{P}}$

Acknowledgments. The author would like to thank C.-P. Jeannerod, G. Labahn, É. Schost, A. Storjohann, and G. Villard for valuable comments and discussions. The author was supported by the international mobility grants Explo'ra Doc from *Région Rhône-Alpes*, *PALSE*, and *Mitacs Globalink - Inria*.

References

- [1] M. Alekhovich. Linear Diophantine equations over polynomials and soft decoding of Reed-Solomon codes. *IEEE Trans. Inf. Theory*, 51(7):2257–2265, July 2005.
- [2] B. Beckermann. A reliable method for computing M-*Padé* approximants on arbitrary staircases. *J. Comput. Appl. Math.*, 40(1):19–42, 1992.
- [3] 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.
- [4] 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.
- [5] B. Beckermann, G. Labahn, and G. Villard. Normal forms for general polynomial matrices. *J. Symbolic Comput.*, 41(6):708–737, 2006.
- [6] Th.G.J. Beelen, G.J. van den Hurk, and C. Praagman. A new method for computing a column reduced polynomial matrix. *Systems and Control Letters*, 10(4):217 – 224, 1988.
- [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. IS-SAC’11, pages 155–162. ACM, 2011.
- [18] C. Hermite. Sur l’introduction des variables continues dans la théorie des nombres. *Journal für die reine und angewandte Mathematik*, 41:191–216, 1851.
- [19] C.-P. Jeannerod, V. Neiger, E. Schost, and G. Villard. Computing minimal nullspace bases, 2015.
- [20] C.-P. Jeannerod, V. Neiger, E. Schost, and G. Villard. Fast computation of minimal nullspace bases in popov form for arbitrary shifts, 2016.
- [21] T. Kailath. *Linear Systems*. Prentice-Hall, 1980.
- [22] S. Lang. *Algebra (Revised Third Edition)*. Springer-Verlag New-York Inc., 2002.
- [23] F. Le Gall. Powers of tensors and fast matrix multiplication. In *ISSAC’14*, pages 296–303. ACM, 2014.
- [24] T. Mulders and A. Storjohann. On lattice reduction for polynomial matrices. *J. Symb. Comput.*, 35:377–401, 2003.
- [25] V. M. Popov. Invariant description of linear, time-invariant controllable systems. *SIAM Journal on Control*, 10(2):252–264, 1972.
- [26] S. Sarkar and A. Storjohann. Normalization of row reduced matrices. In *ISSAC’11*, pages 297–304. ACM, 2011.
- [27] A. Storjohann. High-order lifting and integrality certification. *J. Symbolic Comput.*, 36(3-4):613–648, 2003.
- [28] A. Storjohann. Notes on computing minimal approximant bases. In *Dagstuhl Seminar Proceedings*, 2006.
- [29] 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.
- [30] G. Villard. Computing Popov and Hermite forms of polynomial matrices. ISSAC’96, pages 250–258. ACM, 1996.

- [31] W. Zhou and G. Labahn. Efficient algorithms for order basis computation. *J. Symbolic Comput.*, 47(7):793–819, 2012.
- [32] W. Zhou and G. Labahn. A fast, deterministic algorithm for computing a Hermite normal form of a polynomial matrix, Preprint 2016.
- [33] W. Zhou, G. Labahn, and A. Storjohann. Computing minimal nullspace bases. In *ISSAC'12*, pages 366–373. ACM, 2012.
- [34] W. Zhou, G. Labahn, and A. Storjohann. A deterministic algorithm for inverting a polynomial matrix. *J. Complexity*, 31(2):162 – 173, 2015.

A Reducing the entries of the shift

Let $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ be nonsingular, let $\sigma \in \mathbb{N}$ such that $\sigma > \deg(\det(\mathbf{A}))$, and $\mathbf{s} \in \mathbb{Z}^m$. Here, we show a shift $\mathbf{t} \in \mathbb{N}^m$ such that the \mathbf{s} -Popov form \mathbf{P} of \mathbf{A} equals the \mathbf{t} -Popov form of \mathbf{A} , and $t_{i+1} - t_i \leq \sigma$; thus, in particular $|\mathbf{t}| \in \mathcal{O}(m^2\sigma)$.

Denote $\hat{\mathbf{s}} = (s_{\pi(1)}, \dots, s_{\pi(m)})$ where π is a permutation of $\{1, \dots, m\}$ such that $\hat{\mathbf{s}}$ is non-decreasing. Then, we define $\hat{\mathbf{t}} = (\hat{t}_1, \dots, \hat{t}_m)$ by $\hat{t}_1 = 0$ and, for each $i \in \{2, \dots, m\}$,

$$\hat{t}_i - \hat{t}_{i-1} = \begin{cases} \sigma & \text{if } s_i - s_{i-1} \geq \sigma \\ s_i - s_{i-1} & \text{otherwise} \end{cases} .$$

Then, let $\mathbf{t} = (\hat{t}_{\pi^{-1}(1)}, \dots, \hat{t}_{\pi^{-1}(m)})$. Since every entry on the diagonal of \mathbf{P} is of degree at most $\deg(\det(\mathbf{A})) < \sigma$, we have that \mathbf{P} is also in \mathbf{t} -Popov form.