

Deterministic computation of the characteristic polynomial in the time of matrix multiplication

Vincent Neiger, Clément Pernet

► **To cite this version:**

Vincent Neiger, Clément Pernet. Deterministic computation of the characteristic polynomial in the time of matrix multiplication. 2021. hal-02963147v2

HAL Id: hal-02963147

<https://hal-unilim.archives-ouvertes.fr/hal-02963147v2>

Preprint submitted on 9 Apr 2021

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.

Deterministic computation of the characteristic polynomial in the time of matrix multiplication

Vincent Neiger

Univ. Limoges, CNRS, XLIM, UMR 7252, F-87000 Limoges, France

Clément Pernet

*Université Grenoble Alpes, Laboratoire Jean Kuntzmann, CNRS, UMR 5224
700 avenue centrale, IMAG - CS 40700, 38058 Grenoble cedex 9, France*

Abstract

This paper describes an algorithm which computes the characteristic polynomial of a matrix over a field within the same asymptotic complexity, up to constant factors, as the multiplication of two square matrices. Previously, this was only achieved by resorting to genericity assumptions or randomization techniques, while the best known complexity bound with a general deterministic algorithm was obtained by Keller-Gehrig in 1985 and involves logarithmic factors. Our algorithm computes more generally the determinant of a univariate polynomial matrix in reduced form, and relies on new subroutines for transforming shifted reduced matrices into shifted weak Popov matrices, and shifted weak Popov matrices into shifted Popov matrices.

Keywords: Characteristic polynomial, polynomial matrices, determinant, fast linear algebra.

1. Introduction

The last five decades witnessed a constant effort towards computational reductions of linear algebra problems to matrix multiplication. It has been showed that most classical problems are not harder than multiplying two square matrices, such as matrix inversion, LU decomposition, nullspace basis computation, linear system solving, rank and determinant computations, etc. [7] [25] [8, Chap. 16]. In this context, one major challenge stands out: designing a similar reduction to matrix multiplication for the computation of characteristic polynomials and related objects such as minimal polynomials and Frobenius forms. For the characteristic polynomial, significant progress was achieved by Keller-Gehrig [31], and more recently by Pernet and Storjohann [39] who solved the problem if one allows randomization. This paper closes the problem by providing a deterministic algorithm with the same asymptotic complexity as matrix multiplication.

The characteristic polynomial of a square matrix over a field \mathbb{K} , say $\mathbf{M} \in \mathbb{K}^{m \times m}$, is defined as $\det(x\mathbf{I}_m - \mathbf{M})$. Specific algorithms exist for sparse or structured matrices; here we consider the classical, dense case. In this paper the complexity of an algorithm is measured as an upper bound on its arithmetic cost, that is, the number of basic field operations it uses to compute the output.

Theorem 1.1. *Let \mathbb{K} be a field. Using a subroutine which multiplies two matrices in $\mathbb{K}^{m \times m}$ in $O(m^\omega)$ field operations for some $\omega > 2$, the characteristic polynomial of a matrix in $\mathbb{K}^{m \times m}$ can be computed deterministically in $O(m^\omega)$ field operations.*

Outline. The rest of this introduction gives more details about our framework for complexity bounds (Section 1.1), summarizes previous work (Section 1.2), describes our contribution on polynomial matrix determinant computation (Section 1.3), gives an overview of our approach and of new tools that we designed to avoid logarithmic factors (Sections 1.4 and 1.5), and finally lists a few perspectives (Section 1.6). Section 2 introduces the notation, main definitions, and basic properties used in this paper. Then Section 3 presents the main algorithm of this paper along with a detailed complexity analysis. This algorithm uses two new technical tools described in Sections 4 and 5: the transformation of reduced forms into weak Popov forms and of weak Popov forms into Popov forms, in the case of shifted forms.

1.1. Framework for complexity bounds

In this paper, \mathbb{K} is any field and we seek upper bounds on the complexity of algorithms which operate on objects such as matrices and polynomials over \mathbb{K} . We consider the arithmetic cost of these algorithms, i.e. the number of basic operations in \mathbb{K} that are used to compute the output from some input of a given size. The basic operations are addition, subtraction, multiplication, and inversion in the field, as well as testing whether a given field element is zero.

As already highlighted in Theorem 1.1, in this paper we fix any $2 < \omega \leq 3$ as well as any algorithm which multiplies matrices in $\mathbb{K}^{m \times m}$ using $O(m^\omega)$ operations in \mathbb{K} : this algorithm is assumed to be the one used as a black box for all matrix multiplications arising in the algorithms we design. The current best known cost bounds ensure that any $\omega > 2.373$ is feasible [33]. In practice, one often considers a cubic algorithm with $\omega = 3$ or Strassen's algorithm with $\omega = \log_2(7)$ [47]. Our results hold with the only assumption that $2 < \omega \leq 3$.

In the computer algebra literature, this setting is classical and often implicit; we still emphasize it because here, and more generally when one studies the logarithmic factors in the cost bound of some algorithm, this clarification of how the underlying matrix multiplications are performed is of the utmost importance. Indeed, if one were allowed to use any matrix multiplication subroutine, then the question of logarithmic factors becomes void: for any exponent ω known to be feasible at the time of writing, it is known that $\omega - \varepsilon$ is feasible as well for a sufficiently small $\varepsilon > 0$; then one might rather rely on this faster subroutine, and apply Keller-Gehrig's algorithm to obtain the characteristic polynomial in $O(m^{\omega-\varepsilon} \log(m))$ operations in \mathbb{K} , which is in $O(m^\omega)$.

Similarly, we consider a nondecreasing function $d \mapsto M(d)$ and an algorithm which multiplies two polynomials in $\mathbb{K}[x]$ of degree at most d using at most $M(d)$ operations in \mathbb{K} ; our algorithms rely on this subroutine for polynomial multiplication. Here d is any nonnegative real number; it will often be a fraction D/m of positive integers; we assume that $M(d) = 1$ for $0 \leq d < 1$, so that $M(d) \geq 1$ for all $d \geq 0$. To help derive complexity upper bounds, we also consider the following assumptions \mathcal{H}_{sl} , \mathcal{H}_{sm} , and \mathcal{H}_ω .

\mathcal{H}_{sl} : $2M(d) \leq M(2d)$ for all $d \geq 1$ (superlinearity).

\mathcal{H}_{sm} : $M(d_1 d_2) \leq M(d_1) M(d_2)$ for all $d_1, d_2 \geq 0$ (submultiplicativity).

\mathcal{H}_ω : $M(d) \in O(d^{\omega-1-\varepsilon})$ for some $\varepsilon > 0$.

The first assumption is customary, see e.g. [19, Sec.8.3]; note that it implies $M(d) \geq d$ for all $d \geq 1$. The second and last assumptions are commonly made in complexity analyses for divide and conquer algorithms on polynomial matrices [45, 23]: we refer to [45, Sec.2] for further comments on these assumptions. They are satisfied by the cost bounds of polynomial multiplication algorithms such as the quasi-linear algorithm of Cantor and Kaltofen [9] and, for suitable fields \mathbb{K} , the quasi-linear algorithm of Harvey and van der Hoeven and Lecerf [24], and most of Toom-Cook subquadratic algorithms [50, 10]. For the latter only \mathcal{H}_ω might not

be satisfied, depending on ω and on the number of points used. Note that with the current estimates having $\omega > 2.373$, an order 5 Toom-Cook multiplication (requiring a field with at least 9 points) has exponent $\log(9)/\log(5) \approx 1.365 < \omega - 1$; thus for such exponents ω all Toom-Cook algorithms of order 5 or more satisfy all the above assumptions.

Following [45, 23], we also define a function $d \mapsto M'(d)$ related to the cost of divide and conquer methods such as the half-gcd algorithm: $M'(d) = \sum_{0 \leq i \leq \lceil \log_2(d) \rceil} 2^i M(2^{-i}d)$ for $d \geq 1$, and $M'(d) = 1$ for $0 \leq d \leq 1$. By definition one has $M'(d) \geq M(d) \geq 1$ for all $d \geq 0$, and the identity $M'(2d) = 2M'(d) + M(2d)$ for $d \geq 1$ ensures that $M'(d)$ is superlinear: $2M'(d) \leq M'(2d)$ for all $d \geq 1$. Assuming \mathcal{H}_{sl} yields the asymptotic bound $M'(d) \in O(M(d) \log(d))$ where the $\log(d)$ factor only occurs if a quasi-linear polynomial multiplication is used; in particular, \mathcal{H}_{sl} and \mathcal{H}_ω imply $M'(d) \in O(d^{\omega-1-\varepsilon})$ for some $\varepsilon > 0$. Furthermore, if one assumes \mathcal{H}_{sl} and \mathcal{H}_{sm} , then $M'(\cdot)$ is submultiplicative as well: $M'(d_1 d_2) \leq M'(d_1) M'(d_2)$ for all $d_1, d_2 \geq 0$.

In what follows we assume that two polynomial matrices in $\mathbb{K}[x]^{m \times m}$ of degree at most $d \geq 0$ can be multiplied in $O(m^\omega M(d))$ operations in \mathbb{K} . This is a very mild assumption: it holds as soon as $M(d)$ corresponds to one of the above-mentioned polynomial multiplication algorithms, and it also holds if the chosen matrix multiplication algorithm defining ω supports matrices over a commutative ring using only the operations $\{+, -, \times\}$ (so that one can use it to multiply $m \times m$ matrices over $\mathbb{K}[x]/(x^{2d+1})$). Note still that this bound $O(m^\omega M(d))$ is slightly worse than the best known ones [9, 24]; for example, Cantor and Kaltofen's algorithm performs polynomial matrix multiplication in $O(m^\omega d \log(d) + m^2 d \log(d) \log(\log(d)))$ field operations, which is finer than the bound $O(m^\omega M(d))$ with $M(d) = \Theta(d \log(d) \log(\log(d)))$ in that case. This simplification is frequent in the polynomial matrix literature, and it is made here for the sake of presentation, to improve the clarity of our main complexity results and of the analyses that lead to them.

1.2. Previous work

Previous algorithms based on linear algebra over \mathbb{K} for computing the characteristic polynomial of $\mathbf{M} \in \mathbb{K}^{m \times m}$ mainly fall in three types of methods.

Traces of powers: combining the traces of the first n powers of the input matrix using the Newton identities reveals the coefficients of the characteristic polynomial. Known as the Faddeev-LeVerrier algorithm, it was introduced in [34], refined and rediscovered in [43, 16, 18], and used in [11] to prove that the problem is in the NC^2 parallel complexity class.

Determinant expansion formula: introduced in [41] and improved in [6], this approach does not involve division, and is therefore well suited for computing over integral domains. Later developments in this field include [1, 30], the latter reaching the best known cost bound of $O(m^{2.6973} \log(m)^c)$ ring operations using a deterministic algorithm, for some constant $c > 0$.

Krylov methods: based on sequences of iterates of vectors under the application of the matrix: $(\mathbf{v}, \mathbf{M}\mathbf{v}, \mathbf{M}^2\mathbf{v}, \dots)$. These methods rely on the fact that the first linear dependency between these iterates defines a polynomial which divides the characteristic polynomial. Some algorithms construct the Krylov basis explicitly [31, 20, 14], while others can be interpreted as an implicit Krylov iteration with structured vectors [12, 39].

Methods based on traces of powers use $O(m^4)$ or $O(m^{\omega+1})$ field operations, and are mostly competitive for their parallel complexity. Methods based on determinant expansions use $O(m^4)$ or $O(m^{\omega+1})$ field operations and are relevant for division-free algorithms. Lastly, the Krylov methods run in $O(m^3)$ [12, 14] or $O(m^\omega \log m)$ [31] field operations with deterministic algorithms, or in $O(m^\omega)$ field operations with the Las Vegas randomized algorithm in [39].

Note that the characteristic polynomial of \mathbf{M} cannot be computed faster than the determinant of \mathbf{M} , since the latter is the constant coefficient of the former. Furthermore, under the model of computation trees, the determinant of $m \times m$ matrices cannot be computed faster than the product of two $m \times m$ matrices [8, Sec. 16.4], a consequence of Baur and Strassen’s theorem [2].

Another type of characteristic polynomial algorithms is based on operations on matrices over $\mathbb{K}[x]$, called polynomial matrices in what follows. Indeed the characteristic polynomial may be obtained by calling a determinant algorithm on the characteristic matrix $x\mathbf{I}_m - \mathbf{M}$, which is in $\mathbb{K}[x]^{m \times m}$. Existing algorithms, which accept any matrix in $\mathbb{K}[x]^{m \times m}$ of degree d as input, include

- the evaluation-interpolation method, which costs $O(m^{\omega+1}d + m^3 M'(d))$ field operations, requires that the field \mathbb{K} is large enough, and mainly relies on the computation of about md determinants of matrices in $\mathbb{K}^{m \times m}$;
- the algorithm of Mulders and Storjohann [36] based on weak Popov form computation, which uses $O(m^3 d^2)$ field operations;
- retrieving the determinant as the product of the diagonal entries of the Smith form, itself computed by a Las Vegas randomized algorithm in $O(m^\omega M'(d) \log(m)^2)$ field operations [45, Prop. 41], assuming $M'(d) \in O(d^{\omega-1})$;
- the algorithm based on unimodular triangularization in [32], which is deterministic and uses $O(m^\omega d \log(d)^a \log(m)^b)$ field operations for some constants $a, b \in \mathbb{Z}_{>0}$.

In the last two items the cost bound is, up to logarithmic factors, the same as the cost of multiplying matrices $\mathbb{K}[x]^{m \times m}$ of degree d by relying on both fast linear algebra over \mathbb{K} and fast arithmetic in $\mathbb{K}[x]$, as showed in [9]. The last two of these cost bounds do involve factors logarithmic in m , whereas the first two have an exponent on m which exceeds ω .

In summary, the fastest characteristic polynomial algorithms either are randomized or have a cost a logarithmic factor away from the lower bound. This paper, with Theorem 1.1, bridges this gap by proposing the first deterministic algorithm with cost $O(m^\omega)$.

1.3. A more general result: determinant of reduced polynomial matrices

Our algorithm falls within the category of polynomial matrix determinant computation. Yet unlike the above-listed approaches ours is tailored to a specific family of polynomial matrices, which contains the characteristic matrix $x\mathbf{I}_m - \mathbf{M}$: the family of *row reduced matrices* [52, 29]. Restricting to such matrices provides us with good control of the degrees in computations; as a typical example, it is easy to predict the degree of a vector-matrix product $\mathbf{v}(x\mathbf{I}_m - \mathbf{M})$ by observing the degrees in \mathbf{v} , without actually computing the product. As we explain below, this degree control allows us to avoid searches of degree profiles, which would add logarithmic terms to the cost. Although the characteristic matrix has other properties besides row reducedness (it has degree 1, and is in Popov form [40] hence column reduced), we do not exploit them.

When appropriate, the average row degree D/m , where D is the sum of the degrees of the rows of the matrix, is chosen as a measure of the input degree which refines the matrix degree d used above. This gives cost bounds more sensitive to the input degrees and also, most importantly, leverages the fact that even if the algorithm starts from a matrix with uniform degrees such as $x\mathbf{I}_m - \mathbf{M}$, it may end up handling matrices with unbalanced row degrees in the process.

Theorem 1.2. *Assuming \mathcal{H}_{sl} , \mathcal{H}_{sm} , and \mathcal{H}_ω (hence in particular $\omega > 2$), there is an algorithm which takes as input a row reduced matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times m}$ and computes its determinant using*

$$O(m^\omega M'(D/m)) \subseteq O(m^\omega M'(\deg(\mathbf{A})))$$

operations in \mathbb{K} , where $D = \deg(\det(\mathbf{A}))$ is equal to the sum of the degrees of the rows of \mathbf{A} .

The fact that $\deg(\det(\mathbf{A}))$ is the sum of row degrees is a consequence of row reducedness [29], and the cost bound inclusion follows from $\deg(\det(\mathbf{A})) \leq m \deg(\mathbf{A})$. Taking $\mathbf{A} = x\mathbf{I}_m - \mathbf{M}$ for $\mathbf{M} \in \mathbb{K}^{m \times m}$, Theorem 1.1 is a direct corollary of Theorem 1.2. The only assumption needed in Theorem 1.1 is $\omega > 2$, since it implies the existence of a polynomial multiplication algorithm such that \mathcal{H}_{sl} , \mathcal{H}_{sm} , and \mathcal{H}_ω hold, such as Cantor and Kaltofen’s algorithm [9].

Previous polynomial matrix determinant algorithms with costs of the order of $m^\omega \deg(\mathbf{A})$, up to logarithmic factors, have been listed above: a randomized one from [45], and a deterministic one from [32]. To our knowledge, this paper gives the first description of an algorithm achieving such a cost involving no factor logarithmic in m . Our approach partially follows the algorithm of [32], but also substantially differs from it in a way that allows us to benefit from the reducedness of \mathbf{A} . The cost bound $O(m^\omega M'(\deg(\mathbf{A})))$ has been obtained before in [21, Sec. 4.2.2] in the particular case of a “sufficiently generic” matrix \mathbf{A} . In that case, both the algorithm of [32] and the one here coincide and become the algorithm of [21, Sec. 4.2.2]; when \mathbf{A} is the characteristic matrix $x\mathbf{I}_m - \mathbf{M}$, this also relates to the fast algorithm in [31, Sec. 6] for a generic \mathbf{M} .

1.4. Approach, and existing tools

For the sake of presentation, suppose m is a power of 2. Writing $\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{A}_3 & \mathbf{A}_4 \end{bmatrix}$ with the \mathbf{A}_i ’s of dimensions $(m/2) \times (m/2)$, the algorithm of [32] is based on the block triangularization

$$\begin{bmatrix} * & * \\ \mathbf{K}_1 & \mathbf{K}_2 \end{bmatrix} \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{A}_3 & \mathbf{A}_4 \end{bmatrix} = \begin{bmatrix} \mathbf{R} & * \\ \mathbf{0} & \mathbf{B} \end{bmatrix}$$

where the entries “*” are not computed, $\mathbf{B} = \mathbf{K}_1\mathbf{A}_2 + \mathbf{K}_2\mathbf{A}_4$, and \mathbf{R} and $[\mathbf{K}_1 \ \mathbf{K}_2]$ are computed from $\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_3 \end{bmatrix}$ as a *row basis* and a *kernel basis*, respectively (see Section 2.2 for definitions). Then the leftmost matrix in the above identity is unimodular [32, Lem. 3.1] and thus, up to a constant factor, $\det(\mathbf{A})$ can be computed recursively as $\det(\mathbf{R}) \det(\mathbf{B})$.

A first observation is that neither the kernel basis computation nor the matrix multiplication giving \mathbf{B} is an obstacle towards a cost which is free of $\log(m)$. (The fastest known method for multiplying matrices with unbalanced degrees, such as in $\mathbf{B} = \mathbf{K}_1\mathbf{A}_2 + \mathbf{K}_2\mathbf{A}_4$, splits the computation into $O(\log(m))$ multiplications of smaller matrices with balanced degrees [58, Sec. 3.6], suggesting that its cost may involve a $\log(m)$ factor.) Indeed we show that, under the above assumptions on $M(\cdot)$, the cost of these operations is in $O(m^\omega M(D/m))$ and $O(m^\omega M'(D/m))$, thus only involving factors logarithmic in D/m . In previous work, cost bounds either hide logarithmic factors [58] or they are derived without assuming \mathcal{H}_{sm} and have the form $O(m^{\omega-1} M(D))$ and $O(m^{\omega-1} M'(D))$ [27], thus resulting in factors logarithmic in D . Proving this observation is straightforward from the analyses in [58, 27] (see Section 2.5). This is a first key towards our main result: the characteristic matrix has $D = m$, and $O(m^\omega M(1))$ is the same as $O(m^\omega)$ whereas $O(m^{\omega-1} M(m))$ involves factors logarithmic in m .

However, the computation of the row basis \mathbf{R} remains an obstacle which prevents the algorithm of [32] from being a candidate for Theorem 1.2. Indeed, among the row basis algorithms we are aware of, only one has a cost bound which fits into our target up to logarithmic factors: the one of [56]. It relies on three kernel bases computations, and while one of them is similar to the computation of $[\mathbf{K}_1 \ \mathbf{K}_2]$ and is handled via the algorithm of [58], the two others have different constraints on the input and were the subject of a specific algorithm described in [56, Sec. 4].

¹Precisely, if the upper triangular, row-wise Hermite normal form of \mathbf{A} has diagonal entries $(1, \dots, 1, \lambda \det(\mathbf{A}))$, for some $\lambda \in \mathbb{K} \setminus \{0\}$ making $\lambda \det(\mathbf{A})$ monic.

In this reference, cost bounds were given without showing logarithmic factors; our own analysis revealed the presence of a factor logarithmic in m . The algorithm has a loop over $\Theta(\log(m))$ iterations, each of them calling [55, Algo. 2] for minimal approximant bases with unbalanced input. This approximant basis algorithm may spend a logarithmic number of iterations for finding some degree profile of the output basis, in a way reminiscent of Keller-Gehrig’s algorithm in [31, Sec. 5] which finds the lengths of Krylov sequences (the link between the two situations becomes more explicit for approximant bases at small orders, see [27, Sec. 7]).

Our attempts at accelerating the row basis algorithm of [56] having not succeeded, the algorithm in this paper follows an approach which is more direct at first: remove the obstacle. Instead of computing a row basis \mathbf{R} and relying on the identity $\det(\mathbf{A}) = \det(\mathbf{R}) \det(\mathbf{B})$ (up to a constant), keep the first block row of \mathbf{A} :

$$\begin{bmatrix} \mathbf{I}_{m/2} & \mathbf{0} \\ \mathbf{K}_1 & \mathbf{K}_2 \end{bmatrix} \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{A}_3 & \mathbf{A}_4 \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{0} & \mathbf{B} \end{bmatrix} \quad (1)$$

and rely on the identity $\det(\mathbf{A}) = \det(\mathbf{A}_1) \det(\mathbf{B}) / \det(\mathbf{K}_2)$. The nonsingularity of \mathbf{A}_1 and \mathbf{K}_2 is easily ensured thanks to the assumption that \mathbf{A} is reduced, as discussed in Section 1.5.

This leads to an unusual recursion scheme: we are not aware of a similar scheme being used in the literature on computational linear algebra. The algorithm uses three recursive calls with $(m/2) \times (m/2)$ matrices whose determinant has degree at most $D/2$ for two of them and at most D for the third; our complexity analysis in Section 3.3 shows that such a recursion gives the cost in Theorem 1.2. Precisely, if $\deg(\det(\mathbf{A}_1)) \leq D/2$ then degree properties of minimal kernel bases imply that $\deg(\det(\mathbf{K}_2)) \leq D/2$, yielding the two calls in half the degree; otherwise the algorithm uses inexpensive row and column operations on \mathbf{A} to reduce to the case $\deg(\det(\mathbf{A}_1)) \leq D/2$.

Although this approach removes the obstacle of row basis computation which arises in [32], it adds a requirement: all recursive calls must take input matrices that are reduced. In the next section we discuss how to ensure the reducedness of \mathbf{A}_1 and \mathbf{B} thanks to a straightforward generalization of [42, Sec. 3], and we describe a new algorithm which handles the more involved case of \mathbf{K}_2 .

1.5. New tools, and ensuring reduced form in recursive calls

When outlining the approach of our determinant algorithm via the identity in Eq. (1), we implicitly assumed that the matrices used as input in recursive calls, i.e. \mathbf{A}_1 and \mathbf{K}_2 and \mathbf{B} , do satisfy the input requirement of row reducedness: this is not necessarily the case, even if starting from a reduced matrix \mathbf{A} .

Concerning \mathbf{A}_1 , one may locate such a reduced submatrix of \mathbf{A} and then permute rows and columns of \mathbf{A} (which only affects the sign of $\det(\mathbf{A})$) to make this submatrix become the leading principal submatrix \mathbf{A}_1 . This is a classical operation on reduced matrices which suggests using a form slightly stronger than reduced form called *weak Popov form* [36] (see Section 2.4). Assuming that \mathbf{A} has this form ensures that its leading principal submatrix \mathbf{A}_1 has it as well. This assumption is acceptable in terms of complexity since one can transform a reduced \mathbf{A} into a weak Popov \mathbf{P} by means of fast linear algebra in a cost negligible compared to our target [42, Sec. 3]; note that \mathbf{A} and \mathbf{P} have the same determinant up to an easily found constant (see Algorithm 1).

Next, the cases of \mathbf{K}_2 and \mathbf{B} are strongly linked. First, we will not discuss \mathbf{K}_2 but the whole kernel basis $[\mathbf{K}_1 \ \mathbf{K}_2]$. The fastest known algorithm for computing such a basis is that of [58], and for best efficiency it outputs a matrix in *shifted* reduced form, which is a generalization of reducedness involving degree weights given by a tuple $s \in \mathbb{Z}^m$ called a *shift* (see Sections 2.3

and 2.4 for definitions); the non-shifted case is for $s = \mathbf{0}$. As in the determinant algorithm of [32], here the shift for $[\mathbf{K}_1 \ \mathbf{K}_2]$ is taken as the list of row degrees of \mathbf{A} , denoted by $s = \text{rdeg}(\mathbf{A})$; for the characteristic matrix one has $s = (1, \dots, 1)$ but non-uniform shifts may arise in recursive calls. We want $[\mathbf{K}_1 \ \mathbf{K}_2]$ to be not only s -reduced, but in s -weak Popov form: a direct consequence is that \mathbf{B} is in weak Popov form, and is thus suitable input for a recursive call.

To obtain $[\mathbf{K}_1 \ \mathbf{K}_2]$ we use the kernel basis algorithm of [58] and transform its output into s -weak Popov form. A minor issue is that the fastest known algorithm for such transformations was written in [42, Sec. 3] for non-shifted forms; yet it easily extends to shifted forms as we show in Section 4, obtaining the next result.

Theorem 1.3. *There is an algorithm REDUCEDTOWEAKPOPOV which takes as input a matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ with $m \leq n$ and a shift $s \in \mathbb{Z}^n$ such that \mathbf{A} is in s -reduced form, and returns an s -weak Popov form of \mathbf{A} using $O(m^{\omega-2}nD + m^{\omega-1}n)$ operations in \mathbb{K} , where $D = |\text{rdeg}_s(\mathbf{A})| - m \cdot \min(s)$.*

Here, following usual notation recalled in Section 2.1, $|\text{rdeg}_s(\mathbf{A})|$ is the sum of the s -degrees of the rows of \mathbf{A} . This result extends [42, Thm. 13] since for $s = \mathbf{0}$ the quantity D is the sum of the row degrees of \mathbf{A} and in particular $D \leq m \deg(\mathbf{A})$, leading to the cost bound $O(m^{\omega-1}n \deg(\mathbf{A}))$.

To summarize, at this stage we have outlined how to ensure, without exceeding our target cost bound, that \mathbf{A}_1 and \mathbf{B} are valid input for recursive calls, i.e. are in weak Popov form. Having $\det(\mathbf{A}_1)$ and $\det(\mathbf{B})$, it remains to find $\det(\mathbf{K}_2)$ and then the sought $\det(\mathbf{A})$ follows. We noted that, to ensure the form of \mathbf{B} but also for efficiency reasons, the kernel basis $[\mathbf{K}_1 \ \mathbf{K}_2]$ is computed in s -weak Popov form for the shift $s = \text{rdeg}(\mathbf{A})$. This causes the main difficulty related to our modification of the determinant algorithm of [32]: \mathbf{K}_2 is not valid input for a recursive call since it is in ν -weak Popov form for some shift ν , a subtuple of s which is possibly nonzero.

A first idea is to extend our approach to the shifted case, allowing recursive calls with such a ν -reduced matrix: this is straightforward but gives an inefficient algorithm. Indeed, along the recursion the shift drifts away from its initial value and becomes arbitrarily large and unbalanced with respect to the degrees of the input matrices of recursive calls. For example, as mentioned above the sum of row degrees of the initial non-shifted $m \times m$ matrix \mathbf{A} is $D = \deg(\det(\mathbf{A}))$, whereas for the ν -shifted $(m/2) \times (m/2)$ matrix \mathbf{K}_2 we only have the same bound D instead of one related to $\deg(\det(\mathbf{K}_2))$ itself, which is known to be at most $D/2$ in our algorithm. This gap, here between D and $D/2$, will only grow as the algorithm goes down the tree of recursive calls, meaning that degrees in matrices handled recursively are not sufficiently well controlled.

Another idea is to compute a $\mathbf{0}$ -reduced matrix which has the same determinant as \mathbf{K}_2 . Finding a $\mathbf{0}$ -reduced form of \mathbf{K}_2 within our target cost seems to be a difficult problem. The best known algorithms for general $\mathbf{0}$ -reduction involve $\log(m)$ factors, either explicitly [23] or implicitly [38] (in the latter approach one starts by using the above-discussed triangularization procedure of [32] which we are modifying here to avoid $\log(m)$ factors). More specific algorithms exploit the form of \mathbf{K}_2 , interpreting the problem as a change of shift from ν to $\mathbf{0}$; yet at the time of writing efficient changes of shifts have only been achieved when the target shift is larger than the origin shift [27, Sec. 5], a fact that offers degree control for the transformation between the two matrices. Another possibility is to compute the so-called ν -Popov form \mathbf{P} of \mathbf{K}_2 , since its transpose \mathbf{P}^T is $\mathbf{0}$ -reduced by definition (see Section 2.4), and $\det(\mathbf{P}^T) = \det(\mathbf{P})$ is $\det(\mathbf{K}_2)$ up to a constant. However this suffers from the same issue, as computing \mathbf{P} is essentially the same as changing the shift ν into the nonpositive shift $-\delta$, where δ is the list of diagonal degrees of \mathbf{K}_2 [42, 26].

To circumvent these issues, we use the property that the transpose \mathbf{K}_2^T of a ν -reduced matrix is in $-\mathbf{d}$ -reduced form where $\mathbf{d} = \text{rdeg}_\nu(\mathbf{K}_2)$. This fact naturally comes up here since $\det(\mathbf{K}_2) =$

$\det(\mathbf{K}_2^\top)$), but seems otherwise rarely exploited in polynomial matrix algorithms: in fact we did not find a previous occurrence of it apart from related degree considerations in [56, Lem. 2.2].

Transposing the above two approaches using \mathbf{K}_2^\top instead of \mathbf{K}_2 , we observe that computing a $\mathbf{0}$ -reduced form of \mathbf{K}_2^\top is a change of shift from $-\mathbf{d}$ to $\mathbf{0}$, and computing the $-\mathbf{d}$ -Popov form \mathbf{P} of \mathbf{K}_2^\top is essentially a change of shift from $-\mathbf{d}$ to $-\delta$. In both cases the target shift is larger than the origin shift, implying that the kernel-based change of shift of [27, Sec. 5] involves matrices of well-controlled degrees. Still, this is not enough to make this change of shift efficient as such, the difficulty being now that the average row degree of \mathbf{K}_2^\top may not be small: only its average column degree, which corresponds to the average row degree of \mathbf{K}_2 , is controlled.

Our solution uses the second approach, computing the $-\mathbf{d}$ -Popov form \mathbf{P} , because it offers the a priori knowledge that the column degrees of \mathbf{P} are exactly δ . We exploit this degree knowledge to carry out partial linearization techniques, originally designed for approximant bases [46, 28], which we extend here to kernel bases. These techniques allow us to reduce our problem to a kernel basis computation where the matrix entries have uniformly small degrees, implying that it can be efficiently handled via the minimal approximant basis algorithm PM-BASIS from [21]. The next result summarizes the new algorithmic tool developed in Section 5 for finding \mathbf{P} .

Theorem 1.4. *Let $s \in \mathbb{Z}^m$, let $\mathbf{A} \in \mathbb{K}[x]^{m \times m}$ be in $-s$ -weak Popov form, let $\delta \in \mathbb{Z}_{\geq 0}^m$ be the $-s$ -pivot degree of \mathbf{A} , and assume that $s \geq \delta$. There is an algorithm WEAKPOPOVTOPOPOV which takes as input (\mathbf{A}, s) and computes the $-s$ -Popov form of \mathbf{A} by*

- performing PM-BASIS at order less than $|s|/m + 4$ on an input matrix of row dimension at most $6m$ and column dimension at most $3m$,
- multiplying the inverse of a matrix in $\mathbb{K}^{m \times m}$ by a matrix in $\mathbb{K}[x]^{m \times m}$ of column degree δ ,
- and performing $O(m^2)$ extra operations in \mathbb{K} .

Thus, computing the $-s$ -Popov form of \mathbf{A} can be done in $O(m^\omega M'(|s|/m))$ operations in \mathbb{K} .

This theorem is a generalization of [42, Sec. 4] to shifted forms, for shifts $-s$ that satisfy the assumption $s \geq \delta$. Indeed, if \mathbf{A} is $\mathbf{0}$ -weak Popov, then one recovers [42, Thm. 20] by taking $s = (\deg(\mathbf{A}), \dots, \deg(\mathbf{A}))$ in the above theorem. For comparison, the naive generalization of [42, Sec. 4] to shifted forms runs in $O(m^\omega M'(\max(s)))$, which exceeds our target complexity as soon as $\max(s) \gg |s|/m$. Hence the use of partial linearization techniques, which were not needed in the non-shifted case featuring $\max(s) = |s|/m = \deg(\mathbf{A})$.

As mentioned above, our Algorithm WEAKPOPOVTOPOPOV is based on the computation of a kernel basis with a priori knowledge of the degree profile of the output. This kernel problem is very close to the one handled in [56, Sec. 4], except that in this reference one only has upper bounds on the output degrees, implying a certain number—possibly logarithmic in m —of calls to PM-BASIS to recover the output and its actual degrees. In the same spirit but in the context of approximant bases, [28, Sec. 5] uses partial linearization techniques to reduce an arbitrary input with known output degrees to essentially one call to PM-BASIS, whereas [55, Algo. 2] assumes weaker output degree information and makes a potentially logarithmic number of calls to PM-BASIS.

1.6. Perspectives

We plan to implement our characteristic polynomial algorithm in the LinBox ecosystem [48, 49]. First prototype experiments suggest that, for large finite fields, it could be competitive with the existing fastest-known implementation, based on the randomized algorithm of [39]. The native support for small fields of our algorithm should outperform the algorithm of [39]

which requires expensive field extensions. Another perspective stems from the remark that our algorithm resorts to fast polynomial multiplication (see assumption \mathcal{H}_ω), while previous ones did not [31, 39]: we would like to understand whether the same cost can be achieved by a purely linear algebraic approach. Finally, perhaps the most challenging problem related to characteristic polynomial computation is to compute Frobenius forms deterministically in the time of matrix multiplication, the current best known complexity bound being $O(m^\omega \log(m) \log(\log(m)))$ [44]; and more generally computing Smith forms of polynomial matrices with a cost free of factors logarithmic in the matrix dimension.

2. Preliminaries on polynomial matrices

In this section we present the notation as well as basic definitions and properties that will be used throughout the paper.

2.1. Notation

Tuples of integers will often be manipulated entry-wise. In particular, for tuples $s, t \in \mathbb{Z}^n$ of the same length n , we write $s + t$ for their entry-wise sum, and the inequality $s \leq t$ means that each entry in s is less than or equal to the corresponding entry in t . The concatenation of tuples is denoted by (s, t) . We write $|t|$ for the sum of the entries of t . The tuple of zeros is denoted by $\mathbf{0} = (0, \dots, 0)$; its length is understood from the context.

For an $m \times n$ matrix \mathbf{A} over some ring, we write $\mathbf{A}_{i,j}$ for its entry at index (i, j) . We extend this to submatrices: given sets $I \subseteq \{1, \dots, m\}$ and $J \subseteq \{1, \dots, n\}$ of row and column indices, we write $\mathbf{A}_{I,J}$ for the submatrix of \mathbf{A} formed by its entries indexed by $I \times J$. Besides, $\mathbf{A}_{I,*}$ stands for the submatrix of \mathbf{A} formed by its rows with index in I , and we use the similar notation $\mathbf{A}_{*,J}$. The transpose of \mathbf{A} is denoted by \mathbf{A}^\top . The identity matrix of size n is denoted by \mathbf{I}_n , while the $n \times n$ matrix with 1 on the antidiagonal and 0 elsewhere is denoted by \mathbf{J}_n . In particular, when writing $s\mathbf{J}_n$ for a tuple $s = (s_1, \dots, s_n) \in \mathbb{Z}^n$, we mean the reversed tuple $s\mathbf{J}_n = (s_n, \dots, s_1)$.

Now consider \mathbf{A} with polynomial entries, i.e. $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$. The *degree* of \mathbf{A} is denoted by $\deg(\mathbf{A})$ and is the largest of the degrees of its entries, or $-\infty$ if $\mathbf{A} = \mathbf{0}$. The *row degree* of \mathbf{A} is the tuple $\text{rdeg}(\mathbf{A}) \in (\mathbb{Z}_{\geq 0} \cup \{-\infty\})^m$ whose i th entry is $\max_{1 \leq j \leq n} (\deg(\mathbf{A}_{i,j}))$. More generally, for a tuple $s = (s_1, \dots, s_n) \in \mathbb{Z}^n$, the *s-row degree* of \mathbf{A} is the tuple $\text{rdeg}_s(\mathbf{A}) \in (\mathbb{Z} \cup \{-\infty\})^m$ whose i th entry is $\max_{1 \leq j \leq n} (\deg(\mathbf{A}_{i,j}) + s_j)$. In this context, the tuple s is commonly called a (*degree*) *shift* [4]. The (*shifted*) *column degree* of \mathbf{A} is defined similarly.

We write \mathbf{X}^s for the $n \times n$ diagonal matrix $\text{diag}(x^{s_1}, \dots, x^{s_n})$ which is over the ring $\mathbb{K}[x, x^{-1}]$ of Laurent polynomials over \mathbb{K} . Note that, hereafter, Laurent polynomials will only arise in proofs and explanations, more specifically in considerations about shifted degrees: they never arise in algorithms, which for the sake of clarity only involve polynomials in $\mathbb{K}[x]$. The usefulness of this matrix \mathbf{X}^s will become clear in the definition of leading matrices in the next subsection.

The next lemma gives a link between shifted row degrees and shifted column degrees. We will mostly use the following particular case of it: the column degree of \mathbf{A} is at most $\mathbf{d} \in \mathbb{Z}_{\geq 0}^n$ (entry-wise) if and only if the $-\mathbf{d}$ -row degree of \mathbf{A} is nonpositive.

Lemma 2.1 ([56, Lemma 2.2]). *Let \mathbf{A} be a matrix in $\mathbb{K}[x]^{m \times n}$, \mathbf{d} be a tuple in \mathbb{Z}^n , and \mathbf{t} be a tuple in \mathbb{Z}^m . Then, $\text{cdeg}_t(\mathbf{A}) \leq \mathbf{d}$ if and only if $\text{rdeg}_{-\mathbf{d}}(\mathbf{A}) \leq -\mathbf{t}$.*

2.2. Bases of modules, kernel bases and approximant bases

We recall that any $\mathbb{K}[x]$ -submodule \mathcal{M} of $\mathbb{K}[x]^{1 \times n}$ is free, and admits a basis formed by r elements of $\mathbb{K}[x]^{1 \times n}$, where $r \leq n$ is called the rank of \mathcal{M} [see e.g. 15]. Such a basis can thus be represented as an $r \times n$ matrix \mathbf{B} over $\mathbb{K}[x]$ whose rows are the basis elements; this basis matrix \mathbf{B} has rank r .

For a matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$, its *row space* is the $\mathbb{K}[x]$ -submodule $\{\mathbf{pA}, \mathbf{p} \in \mathbb{K}[x]^{1 \times m}\}$ of $\mathbb{K}[x]^{1 \times n}$, that is, the set of all $\mathbb{K}[x]$ -linear combinations of its rows. If $\mathbf{B} \in \mathbb{K}[x]^{r \times n}$ is a basis of this row space, then \mathbf{B} is said to be a *row basis* of \mathbf{A} ; in particular, r is the rank of \mathbf{B} and of \mathbf{A} .

The *left kernel* of \mathbf{A} , denoted by $\mathcal{K}(\mathbf{A})$, is the $\mathbb{K}[x]$ -module $\{\mathbf{p} \in \mathbb{K}[x]^{1 \times m} \mid \mathbf{pA} = \mathbf{0}\}$. A matrix $\mathbf{K} \in \mathbb{K}[x]^{k \times m}$ is a *left kernel basis* of \mathbf{A} if its rows form a basis of $\mathcal{K}(\mathbf{A})$, in which case $k = m - r$. Similarly, a *right kernel basis* of \mathbf{A} is a matrix $\mathbf{K} \in \mathbb{K}[x]^{n \times (n-r)}$ whose *columns* form a basis of the right kernel of \mathbf{A} .

Given positive integers $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_n) \in \mathbb{Z}_{>0}^n$ and a matrix $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$, the set of *approximants for \mathbf{F} at order $\boldsymbol{\gamma}$* [see e.g. 51, 3] is the $\mathbb{K}[x]$ -submodule of $\mathbb{K}[x]^{1 \times m}$ defined as

$$\mathcal{A}_{\boldsymbol{\gamma}}(\mathbf{F}) = \{\mathbf{p} \in \mathbb{K}[x]^{1 \times m} \mid \mathbf{pF} = \mathbf{0} \bmod \mathbf{X}^{\boldsymbol{\gamma}}\}.$$

The identity $\mathbf{pF} = \mathbf{0} \bmod \mathbf{X}^{\boldsymbol{\gamma}}$ means that $\mathbf{pF}_{s,j} = 0 \bmod x^{\gamma_j}$ for $1 \leq j \leq n$. Since all m rows of the matrix $x^{\max(\boldsymbol{\gamma})} \mathbf{I}_m$ are in $\mathcal{A}_{\boldsymbol{\gamma}}(\mathbf{F})$, this module has rank m .

2.3. Leading matrices and reduced forms of polynomial matrices

We will often compute with polynomial matrices that have a special form, called the (*shifted*) *reduced form*. It corresponds to a type of minimality of the degrees of such matrices, and also provides good control of these degrees during computations as illustrated by the *predictable degree property* [17] [29, Thm. 6.3-13] which we recall below. In this section, we introduce the notion of *row reducedness*; to avoid confusion, we will not use the similar notion of column reducedness in this paper, and thus all further mentions of reducedness refer to row reducedness.

For shifted reduced forms, we follow the definitions in [4, 5]. Let $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ and $\mathbf{s} \in \mathbb{Z}^n$, and let $\mathbf{t} = (t_1, \dots, t_m) = \text{rdeg}_s(\mathbf{A})$. Then, the \mathbf{s} -leading matrix of \mathbf{A} is the matrix $\text{lm}_s(\mathbf{A}) \in \mathbb{K}^{m \times n}$ whose entry (i, j) is the coefficient of degree $t_i - s_j$ of the entry (i, j) of \mathbf{A} , or 0 if $t_i = -\infty$. Equivalently, $\text{lm}_s(\mathbf{A})$ is the coefficient of degree zero of $\mathbf{X}^{-\mathbf{t}} \mathbf{A} \mathbf{X}^{\mathbf{s}}$, whose entries are in $\mathbb{K}[x^{-1}]$. The matrix \mathbf{A} is said to be in \mathbf{s} -reduced form if its \mathbf{s} -leading matrix has full row rank. In particular, a matrix in \mathbf{s} -reduced form must have full row rank.

For a matrix $\mathbf{M} \in \mathbb{K}[x]^{k \times m}$, we have $\text{rdeg}_s(\mathbf{MA}) \leq \text{rdeg}_t(\mathbf{M})$ and this is an equality when no cancellation of leading terms occurs in this left-multiplication. The predictable degree property states that \mathbf{A} is \mathbf{s} -reduced if and only if $\text{rdeg}_s(\mathbf{MA}) = \text{rdeg}_t(\mathbf{M})$ holds for any $\mathbf{M} \in \mathbb{K}[x]^{k \times m}$. Here is a useful consequence of this characterization.

Lemma 2.2. *Let $\mathbf{s} \in \mathbb{Z}^n$, let $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$, and let $\mathbf{t} = \text{rdeg}_s(\mathbf{A})$. If \mathbf{A} is \mathbf{s} -reduced, then the identity $\text{lm}_s(\mathbf{MA}) = \text{lm}_t(\mathbf{M}) \text{lm}_s(\mathbf{A})$ holds for any $\mathbf{M} \in \mathbb{K}[x]^{k \times m}$.*

Proof. Let $\mathbf{d} = \text{rdeg}_s(\mathbf{MA})$. By definition, $\text{lm}_s(\mathbf{MA})$ is the coefficient of degree 0 of the matrix $\mathbf{X}^{-\mathbf{d}} \mathbf{M} \mathbf{A} \mathbf{X}^{\mathbf{s}} = \mathbf{X}^{-\mathbf{d}} \mathbf{M} \mathbf{X}^{\mathbf{t}} \mathbf{X}^{-\mathbf{t}} \mathbf{A} \mathbf{X}^{\mathbf{s}}$, whose entries are in $\mathbb{K}[x^{-1}]$. Besides, since $\text{rdeg}_s(\mathbf{A}) = \mathbf{t}$ and since the predictable degree property gives $\text{rdeg}_t(\mathbf{M}) = \mathbf{d}$, the matrices $\mathbf{X}^{-\mathbf{d}} \mathbf{M} \mathbf{X}^{\mathbf{t}}$ and $\mathbf{X}^{-\mathbf{t}} \mathbf{A} \mathbf{X}^{\mathbf{s}}$ are over $\mathbb{K}[x^{-1}]$ and their coefficients of degree 0 are $\text{lm}_t(\mathbf{M})$ and $\text{lm}_s(\mathbf{A})$, respectively. \square

Another characterization of matrices in \mathbf{s} -reduced form is that they have minimal \mathbf{s} -row degree among all matrices which represent the same $\mathbb{K}[x]$ -module [53, Def. 2.13]; in this paper, we will use the following consequence of this minimality.

Lemma 2.3. *Let \mathcal{M} be a submodule of $\mathbb{K}[x]^{1 \times n}$ of rank m , let $s \in \mathbb{Z}^n$, and let $\mathbf{t} \in \mathbb{Z}^m$ be the s -row degree of some s -reduced basis of \mathcal{M} . Without loss of generality, assume that \mathbf{t} is nondecreasing. Let $\mathbf{B} \in \mathbb{K}[x]^{m \times n}$ be a matrix of rank m whose rows are in \mathcal{M} , and let $\mathbf{d} \in \mathbb{Z}^m$ be its s -row degree sorted in nondecreasing order. If $\mathbf{d} \leq \mathbf{t}$, then \mathbf{B} is an s -reduced basis of \mathcal{M} , and $\mathbf{d} = \mathbf{t}$.*

Proof. Up to permuting the rows of \mathbf{B} , we assume that $\text{rdeg}_s(\mathbf{B}) = \mathbf{d}$ without loss of generality. Let $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ be an s -reduced basis of \mathcal{M} such that $\text{rdeg}_s(\mathbf{A}) = \mathbf{t}$. Since the rows of \mathbf{B} are in \mathcal{M} , there exists a matrix $\mathbf{U} \in \mathbb{K}[x]^{m \times m}$ such that $\mathbf{B} = \mathbf{U}\mathbf{A}$; and \mathbf{U} is nonsingular since \mathbf{B} and \mathbf{A} have rank m . Since \mathbf{A} is s -reduced, the predictable degree property applies, ensuring that

$$\mathbf{d} = \text{rdeg}_s(\mathbf{B}) = \text{rdeg}_s(\mathbf{U}\mathbf{A}) = \text{rdeg}_t(\mathbf{U}).$$

This means that $\deg(\mathbf{U}_{i,j}) \leq d_i - t_j$ for all $1 \leq i, j \leq m$.

Now, assume by contradiction that $\mathbf{d} = \mathbf{t}$ does not hold. Thus, $d_k < t_k$ for some $1 \leq k \leq m$. Then, for $i \leq k$ and $j \geq k$ we have $d_i \leq d_k < t_k \leq t_j$, hence $\deg(\mathbf{U}_{i,j}) < 0$. Thus, the submatrix $\mathbf{U}_{\{1, \dots, k\}, \{k, \dots, m\}}$ is zero, which implies that \mathbf{U} is singular; this is a contradiction, hence $\mathbf{d} = \mathbf{t}$.

Since \mathbf{t} is nondecreasing, the inequality $\deg(\mathbf{U}_{i,j}) \leq t_i - t_j$ implies that \mathbf{U} is a block lower triangular matrix whose diagonal blocks have degree 0; hence these blocks are invertible matrices over \mathbb{K} , and \mathbf{U} is unimodular [see 42, Lemma 6 for similar degree considerations, starting from stronger assumptions on \mathbf{A} and \mathbf{B}]. Thus, \mathbf{B} is a basis of \mathcal{M} .

Furthermore, it is easily observed that $\text{lm}_d(\mathbf{U}) \in \mathbb{K}^{m \times m}$ is block lower triangular with the same invertible diagonal blocks as \mathbf{U} ; hence $\text{lm}_d(\mathbf{U})$ is invertible. On the other hand, Lemma 2.2 states that $\text{lm}_s(\mathbf{B}) = \text{lm}_d(\mathbf{U})\text{lm}_s(\mathbf{A})$. Thus $\text{lm}_s(\mathbf{B})$ has rank $m = \text{rank}(\text{lm}_s(\mathbf{A}))$, and \mathbf{B} is s -reduced. \square

2.4. Pivots and weak Popov forms of polynomial matrices

For algorithmic purposes, it is often convenient to work with reduced forms that satisfy some additional requirements, called *weak Popov forms*. These are intrinsically related to the notion of *pivot* of a polynomial matrix.

For a nonzero vector $\mathbf{p} \in \mathbb{K}[x]^{1 \times n}$ and a shift $s \in \mathbb{Z}^n$, the s -pivot of \mathbf{p} is its rightmost entry p_j such that $\deg(p_j) + s_j = \text{rdeg}_s(\mathbf{p})$ [4, 36]; it corresponds to the rightmost nonzero entry of $\text{lm}_s(\mathbf{p})$. The index $j = \pi$ and the degree $\deg(p_\pi) = \delta$ of this entry are called the s -pivot index and s -pivot degree, respectively. For brevity, in this paper the pair (π, δ) is called the s -pivot profile of \mathbf{p} . By convention, the zero vector in $\mathbb{K}[x]^{1 \times n}$ has s -pivot index 0 and s -pivot degree $-\infty$. These notions are extended to matrices $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ by forming row-wise lists. For example, the s -pivot index of \mathbf{A} is $\boldsymbol{\pi} = (\pi_1, \dots, \pi_m) \in \mathbb{Z}_{>0}^m$ where π_i is the s -pivot index of the row $\mathbf{A}_{i,*}$. The s -pivot degree δ and the s -pivot profile $(\pi_i, \delta_i)_{1 \leq i \leq m}$ of \mathbf{A} are defined similarly.

Then, \mathbf{A} is said to be in s -weak Popov form if it has no zero row and $\boldsymbol{\pi}$ is strictly increasing; and \mathbf{A} is said to be in s -unordered weak Popov form if it is in s -weak Popov form up to row permutation, i.e. the entries of $\boldsymbol{\pi}$ are pairwise distinct. Furthermore, a matrix is in s -Popov form if it is in s -weak Popov form, its s -pivots are monic, and each of these s -pivots has degree strictly larger than the other entries in the same column. For a given $\mathbb{K}[x]$ -submodule \mathcal{M} of $\mathbb{K}[x]^{1 \times n}$, there is a unique basis of \mathcal{M} which is in s -Popov form [4].

For a given matrix \mathbf{B} , the matrix \mathbf{A} is said to be an s -reduced (resp. s -weak Popov, s -Popov) form of \mathbf{B} if \mathbf{A} is a row basis of \mathbf{B} and \mathbf{A} is in s -reduced (resp. s -weak Popov, s -Popov) form.

Like for s -reducedness, the property of a matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ to be in s -weak Popov form depends only on its s -leading matrix $\text{lm}_s(\mathbf{A}) \in \mathbb{K}^{m \times n}$, namely on the fact that it has a staircase shape. Indeed, \mathbf{A} is in s -weak (resp. s -unordered weak) Popov form if and only if $\text{lm}_s(\mathbf{A})$ has no

zero row and $\mathbf{J}_m \text{lm}_s(\mathbf{A}) \mathbf{J}_n$ is in row echelon form (resp. in row echelon form up to row permutation); this was used as a definition by Beckermann et al. [4, 5]. In particular, for any constant matrix $\mathbf{C} \in \mathbb{K}^{m \times n}$, we have $\text{lm}_0(\mathbf{C}) = \mathbf{C}$ and therefore \mathbf{C} is in $\mathbf{0}$ -weak (resp. $\mathbf{0}$ -unordered weak) Popov form if and only if it has no zero row and $\mathbf{J}_m \mathbf{C} \mathbf{J}_n$ is in row echelon form (resp. in row echelon form up to row permutation). Taking $\mathbf{C} = \text{lm}_s(\mathbf{A})$, the next lemma follows.

Lemma 2.4. *Let $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ and let $s \in \mathbb{Z}^n$. Then, \mathbf{A} is in s -weak (resp. s -unordered weak) Popov form if and only if $\text{lm}_s(\mathbf{A})$ is in $\mathbf{0}$ -weak (resp. $\mathbf{0}$ -unordered weak) Popov form.*

Furthermore, if \mathbf{A} is in s -weak Popov form and (j_1, \dots, j_m) is the list of indices of pivot columns in the row echelon form $\mathbf{J}_m \text{lm}_s(\mathbf{A}) \mathbf{J}_n$ (in other words, this list is the column rank profile of that matrix), then the s -pivot index of \mathbf{A} is equal to $(n + 1 - j_m, \dots, n + 1 - j_1)$. This leads to the following lemma which states that the s -pivot profile is an invariant of left-unimodularly equivalent s -weak Popov forms [29, 4, 5], generalizing the fact that for matrices over \mathbb{K} the set of indices of pivot columns is an invariant of left-equivalent row echelon forms.

Lemma 2.5. *Let $s \in \mathbb{Z}^n$ and let $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ be in s -unordered weak Popov form with s -pivot profile $(\pi_i, \delta_i)_{1 \leq i \leq m}$. Then, the s -pivot profile of the s -Popov form of \mathbf{A} is $(\pi_{\sigma(i)}, \delta_{\sigma(i)})_{1 \leq i \leq m}$, where $\sigma : \{1, \dots, m\} \rightarrow \{1, \dots, m\}$ is the permutation such that $(\pi_{\sigma(i)})_{1 \leq i \leq m}$ is strictly increasing.*

Proof. Without loss of generality we assume that \mathbf{A} is in s -weak Popov form, implying also $\sigma(i) = i$ for $1 \leq i \leq m$. Let $\mathbf{P} \in \mathbb{K}[x]^{m \times n}$ be the s -Popov form of \mathbf{A} : we want to prove that \mathbf{A} and \mathbf{P} have the same s -pivot index and the same s -pivot degree. Let \mathbf{U} be the unimodular matrix such that $\mathbf{P} = \mathbf{U}\mathbf{A}$; then Lemma 2.2 yields $\text{lm}_s(\mathbf{P}) = \text{lm}_t(\mathbf{U})\text{lm}_s(\mathbf{A})$, where $t = \text{rdeg}_s(\mathbf{A})$. Since both $\text{lm}_s(\mathbf{P})$ and $\text{lm}_s(\mathbf{A})$ have full row rank, $\text{lm}_t(\mathbf{U}) \in \mathbb{K}^{m \times m}$ is invertible. Then

$$\mathbf{J}_m \text{lm}_s(\mathbf{P}) \mathbf{J}_n = \mathbf{J}_m \text{lm}_t(\mathbf{U}) \mathbf{J}_m \mathbf{J}_m \text{lm}_s(\mathbf{A}) \mathbf{J}_n$$

holds, and thus the row echelon forms $\mathbf{J}_m \text{lm}_s(\mathbf{P}) \mathbf{J}_n$ and $\mathbf{J}_m \text{lm}_s(\mathbf{A}) \mathbf{J}_n$ have the same pivot columns since $\mathbf{J}_m \text{lm}_t(\mathbf{U}) \mathbf{J}_m \in \mathbb{K}^{m \times m}$ is invertible. It follows from the discussion preceding this lemma that \mathbf{P} has the same s -pivot index as \mathbf{A} .

As a consequence, \mathbf{P} has the same s -pivot degree as \mathbf{A} if and only if $\text{rdeg}_s(\mathbf{P}) = \text{rdeg}_s(\mathbf{A})$. Suppose by contradiction that there exists an index i such that $\text{rdeg}_s(\mathbf{P}_{i,*}) < \text{rdeg}_s(\mathbf{A}_{i,*})$. Then, build the matrix $\mathbf{B} \in \mathbb{K}[x]^{m \times n}$ which is equal to \mathbf{A} except for its i th row which is replaced by $\mathbf{P}_{i,*}$. By construction, \mathbf{B} has rank m (since it is in s -weak Popov form) and its rows are in the row space of \mathbf{A} . Writing \mathbf{d} for the tuple $\text{rdeg}_s(\mathbf{B})$ sorted in nondecreasing order, and \mathbf{u} for the tuple \mathbf{t} sorted in nondecreasing order, we have $\mathbf{d} \leq \mathbf{u}$ and $\mathbf{d} \neq \mathbf{u}$, which contradicts Lemma 2.3. Hence there is no such index i , and since this proof by contradiction is symmetric in \mathbf{A} and \mathbf{P} , there is no index i such that $\text{rdeg}_s(\mathbf{A}_{i,*}) < \text{rdeg}_s(\mathbf{P}_{i,*})$ either. Thus $\text{rdeg}_s(\mathbf{A}) = \text{rdeg}_s(\mathbf{P})$. \square

We will also use the following folklore fact, which is a corollary of Lemma 2.2, and has often been used in algorithms for approximant bases or kernel bases in order to preserve the reducedness of matrices during the computation.

Lemma 2.6. *Let $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ and $\mathbf{B} \in \mathbb{K}[x]^{k \times m}$, and let $s \in \mathbb{Z}^n$ and $t = \text{rdeg}_s(\mathbf{A}) \in \mathbb{Z}^m$. Then,*

- if \mathbf{A} is s -reduced and \mathbf{B} is t -reduced, then $\mathbf{B}\mathbf{A}$ is s -reduced;
- if \mathbf{A} is in s -weak Popov form and \mathbf{B} is in t -weak Popov form, then $\mathbf{B}\mathbf{A}$ is in s -weak Popov form.

Proof. Since \mathbf{A} is s -reduced, Lemma 2.2 states that $\text{Im}_s(\mathbf{BA}) = \mathbf{ML}$ where $\mathbf{L} = \text{Im}_s(\mathbf{A}) \in \mathbb{K}^{m \times n}$ and $\mathbf{M} = \text{Im}_t(\mathbf{B}) \in \mathbb{K}^{k \times m}$. The first item then follows from the fact that if \mathbf{M} has rank k and \mathbf{L} has rank m , then \mathbf{ML} has rank k . Similarly, the second item reduces to prove that, assuming \mathbf{M} and \mathbf{L} are in row echelon form with full row rank, then \mathbf{ML} is also in row echelon form. Let (a_1, \dots, a_k) (resp. (b_1, \dots, b_m)) be the pivot indices of \mathbf{M} (resp. \mathbf{L}). Then the i th row of \mathbf{ML} is a nonzero multiple of row a_i of \mathbf{L} combined with multiples of rows of \mathbf{L} of index greater than a_i . Consequently, the pivot indices of the rows of \mathbf{ML} are $b_{a_1} < \dots < b_{a_k}$, which proves that \mathbf{ML} is in row echelon form. \square

Finally, under assumptions that generalize the situation encountered in our determinant algorithm below, we show that the pivot entries of a kernel basis $[\mathbf{K}_1 \ \mathbf{K}_2]$ are located in its rightmost columns, that is, in \mathbf{K}_2 .

Lemma 2.7. *Let $t \in \mathbb{Z}^n$, let $\mathbf{F} \in \mathbb{K}[x]^{n \times n}$ be in t -weak Popov form, and let $\mathbf{u} = \text{rdeg}_t(\mathbf{F})$. Let $\mathbf{G} \in \mathbb{K}[x]^{m \times n}$ and $\mathbf{v} \in \mathbb{Z}^m$ be such that $\mathbf{v} \geq \text{rdeg}_t(\mathbf{G})$, and let $\mathbf{K} = [\mathbf{K}_1 \ \mathbf{K}_2] \in \mathbb{K}[x]^{m \times (m+n)}$ be a (\mathbf{u}, \mathbf{v}) -weak Popov basis of $\mathcal{K}(\begin{bmatrix} \mathbf{F} \\ \mathbf{G} \end{bmatrix})$, where \mathbf{K}_1 and \mathbf{K}_2 have m and n columns, respectively. Then, the (\mathbf{u}, \mathbf{v}) -pivot entries of \mathbf{K} are all located in \mathbf{K}_2 ; in particular, \mathbf{K}_2 is in \mathbf{v} -weak Popov form.*

Proof. Since the (\mathbf{u}, \mathbf{v}) -pivot entry of a row is the rightmost entry of that row which reaches the (\mathbf{u}, \mathbf{v}) -row degree, it is enough to prove that $\text{rdeg}_v(\mathbf{K}_2) \geq \text{rdeg}_u(\mathbf{K}_1)$. First, from $\mathbf{v} \geq \text{rdeg}_t(\mathbf{G})$, we obtain $\text{rdeg}_v(\mathbf{K}_2) \geq \text{rdeg}_{\text{rdeg}_t(\mathbf{G})}(\mathbf{K}_2)$. Now, by definition, $\text{rdeg}_{\text{rdeg}_t(\mathbf{G})}(\mathbf{K}_2) \geq \text{rdeg}_t(\mathbf{K}_2\mathbf{G})$. Since the rows of \mathbf{K} are in $\mathcal{K}(\begin{bmatrix} \mathbf{F} \\ \mathbf{G} \end{bmatrix})$, we have $\mathbf{K}_2\mathbf{G} = -\mathbf{K}_1\mathbf{F}$, hence $\text{rdeg}_t(\mathbf{K}_2\mathbf{G}) = \text{rdeg}_t(\mathbf{K}_1\mathbf{F})$. Since \mathbf{F} is t -reduced, we can apply the predictable degree property: $\text{rdeg}_t(\mathbf{K}_1\mathbf{F}) = \text{rdeg}_u(\mathbf{K}_1)$. This proves the sought inequality. For the last point, note that the (\mathbf{u}, \mathbf{v}) -pivot entries of $[\mathbf{K}_1 \ \mathbf{K}_2]$ located in \mathbf{K}_2 correspond to \mathbf{v} -pivot entries in \mathbf{K}_2 . Thus, since $[\mathbf{K}_1 \ \mathbf{K}_2]$ is in (\mathbf{u}, \mathbf{v}) -weak Popov form with all (\mathbf{u}, \mathbf{v}) -pivot entries in \mathbf{K}_2 , it follows that the \mathbf{v} -pivot index of \mathbf{K}_2 is increasing. \square

2.5. Basic subroutines and their complexity

To conclude these preliminaries, we recall known fast algorithms for three polynomial matrix subroutines used in our determinant algorithm: multiplication with unbalanced degrees, minimal approximant bases, and minimal kernel bases; we give the corresponding complexity estimates adapted to our context and in particular using our framework stated in Section 1.1.

Unbalanced multiplication. Polynomial matrix algorithms often involve multiplication with matrix operands whose entries have degrees that may be unbalanced but still satisfy properties that can be exploited to perform the multiplication efficiently. Here we will encounter products of reduced matrices with degree properties similar to those discussed in [58, Sec. 3.6], where an efficient approach for computing such products was given.

Lemma 2.8. *There is an algorithm UNBALANCEDMULTIPLICATION which takes as input a matrix $\mathbf{B} \in \mathbb{K}[x]^{k \times m}$ with $k \leq m$, a matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ with $n \leq m$, and an integer D greater than or equal to both the sum of the positive entries of $\text{rdeg}_0(\mathbf{A})$ and that of $\text{rdeg}_{\text{rdeg}_0(\mathbf{A})}(\mathbf{B})$, and returns the product \mathbf{BA} using $O(m^\omega \mathbb{M}(D/m))$ operations in \mathbb{K} , assuming \mathcal{H}_{sm} and \mathcal{H}_ω .*

Proof. Zhou et al. [58, Sec. 3.6] gave such an algorithm, yet with a cost analysis which hides logarithmic factors; because these factors are our main concern here we will rely on the version in [27, Sec. 4]. In this reference, Algorithm UNBALANCEDMULTIPLICATION was described for square matrices. One could adapt it to the case of rectangular \mathbf{A} and \mathbf{B} as in the statement above.

However, for the sake of conciseness and with no impact on the asymptotic cost bound, we consider the more basic approach of forming the square $m \times m$ matrices $\mathbf{D} = \begin{bmatrix} \mathbf{B} \\ \mathbf{0} \end{bmatrix}$ and $\mathbf{C} = [\mathbf{A} \ \mathbf{0}]$, computing \mathbf{DC} using the above-cited algorithm, and retrieving \mathbf{BA} from it. Now, by construction, both the sum of the positive entries of $\text{rdeg}_0(\mathbf{C})$ and that of $\text{rdeg}_{\text{rdeg}_0(\mathbf{A})}(\mathbf{D})$ are at most D , hence [27, Prop. 4.1] applies: defining \bar{m} and d as the smallest powers of 2 greater than or equal to m and D/m , it states that the computation of \mathbf{DC} costs $O(\sum_{0 \leq i \leq \log_2(\bar{m})} 2^i (2^{-i} \bar{m})^\omega M(2^i d))$ operations in \mathbb{K} . Using \mathcal{H}_{sm} and \mathcal{H}_ω , which ensure respectively that $M(2^i d) \leq M(2^i) M(d)$ and $M(2^i) \in O(2^{i(\omega-1-\varepsilon)})$ for some $\varepsilon > 0$, we obtain that this bound is in $O(\bar{m}^\omega M(d) \sum_{0 \leq i \leq \log_2(\bar{m})} 2^{-i\varepsilon}) \subseteq O(\bar{m}^\omega M(d))$. This is in $O(m^\omega M(D/m))$, since $\bar{m} \in \Theta(m)$ and $d \in \Theta(1 + D/m)$. \square

Minimal approximant basis. The second basic tool we will use is approximant bases computation; for this, we will use the algorithm PM-BASIS, originally described in [21]. Precisely, we rely on the slightly modified version presented in [28] which ensures that the computed basis is in shifted weak Popov form.

Lemma 2.9. *There is an algorithm PM-BASIS which takes as input a tuple $\boldsymbol{\gamma} \in \mathbb{Z}_{>0}^n$, a matrix $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$ with $\text{cdeg}(\mathbf{F}) < \boldsymbol{\gamma}$, and a shift $s \in \mathbb{Z}^m$, and returns a basis of $\mathcal{A}_{\boldsymbol{\gamma}}(\mathbf{F})$ in s -weak Popov form using $O((m+n)m^{\omega-1} M'(\max(\boldsymbol{\gamma})))$ operations in \mathbb{K} .*

Proof. The algorithm is [28, Algo. 2]; to accommodate non-uniform order $\boldsymbol{\gamma}$, it is called with input order $\Gamma = \max(\boldsymbol{\gamma})$ and input matrix $\mathbf{F}\mathbf{X}^{(\Gamma, \dots, \Gamma) - \boldsymbol{\gamma}}$ as explained in [28, Rmk. 3.3]. According to [28, Prop. 3.2], this costs $O((1 + \frac{n}{m}) \sum_{0 \leq i \leq \lceil \log_2(\Gamma) \rceil} 2^i m^\omega M(2^{-i}\Gamma))$ operations in \mathbb{K} , which is precisely the claimed bound by definition of $M'(\cdot)$. \square

Minimal kernel basis. We will make use of the algorithm of Zhou et al. [58], which itself relies on unbalanced products and approximant bases, and returns a kernel basis in shifted reduced form efficiently for input matrices with small average row degree.

Lemma 2.10. *There is an algorithm KERNELBASIS which takes as input a full column rank matrix $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$ with $m \geq n$ and $m \in O(n)$, and a shift $s \in \mathbb{Z}_{\geq 0}^m$ such that $s \geq \text{rdeg}_0(\mathbf{F})$, and returns a basis of $\mathcal{K}(\mathbf{F})$ in s -reduced form using $O(m^\omega M'(D/m))$ operations in \mathbb{K} , assuming $\mathcal{H}_{\text{sl}}, \mathcal{H}_{\text{sm}}, \mathcal{H}_\omega$. Here D is the sum of the entries of s , and the sum of the s -row degree of this basis is at most D .*

Proof. The algorithm of Zhou et al. [58] computes an s -reduced basis $\mathbf{K} \in \mathbb{K}[x]^{k \times m}$ of $\mathcal{K}(\mathbf{F})$; precisely, this reference is about computing a basis of the right kernel in column reduced form, yet this naturally translates into left kernels and row reduced forms by taking suitable transposes. Furthermore, the last claim in the lemma follows from [58, Thm. 3.4], which states that any such basis \mathbf{K} is such that $|\text{rdeg}_s(\mathbf{K})| \leq |s| = D$. For the complexity, we rely on the analysis in [27, Prop. B.1] which shows that, defining \bar{m} and d as the smallest powers of 2 greater than or equal to m and D/m , this computation costs

$$O\left(\sum_{j=0}^{\log_2(\bar{m})} 2^j \left(\sum_{i=0}^{\log_2(2^{-j}\bar{m})} 2^i (2^{-i-j}\bar{m})^\omega M(2^{i+j}d) + \sum_{i=0}^{\log_2(2^j d)} 2^i (2^{-j}\bar{m})^\omega M(2^{j-i}d)\right)\right)$$

operations in \mathbb{K} . Now the same analysis as in the proof of Lemma 2.8 shows that, assuming \mathcal{H}_{sm} and \mathcal{H}_ω , the first inner sum is in $O((2^{-j}\bar{m})^\omega M(2^j d))$, and by definition of $M'(\cdot)$ the second inner sum is in $O((2^{-j}\bar{m})^\omega M'(2^j d))$. Thus the total cost is in $O(\sum_{0 \leq j \leq \log_2(\bar{m})} 2^j (2^{-j}\bar{m})^\omega M'(2^j d))$, which is in $O(\bar{m}^\omega M'(d) \sum_{0 \leq j \leq \log_2(\bar{m})} 2^{j(1-\omega)} M'(2^j))$ since \mathcal{H}_{sl} and \mathcal{H}_{sm} ensure that $M'(\cdot)$ is submultiplicative. Similarly to the proof of Lemma 2.8, this bound is in $O(m^\omega M'(D/m))$ thanks to \mathcal{H}_ω . \square

3. Determinant algorithm for reduced matrices

In this section, we present the main algorithm in this paper, which computes the determinant of a matrix in reduced form using the subroutines listed in Section 2.5 as well as the algorithms REDUCEDTOWEAKPOPOV and WEAKPOPOVTOPOPOV from Theorems 1.3 and 1.4. Taking for granted the proof of these theorems in Sections 4 and 5, here we prove the correctness of our determinant algorithm in Section 3.2 and analyse its complexity in Section 3.3, thus proving Theorem 1.2.

3.1. Two properties of determinants of reduced matrices

Leading coefficient of the determinant. All bases of a given submodule of $\mathbb{K}[x]^{1 \times n}$ of rank n have the same determinant up to a constant factor, i.e. up to multiplication by an element of $\mathbb{K} \setminus \{0\}$. Many algorithms operating on polynomial matrices such as PM-BASIS and KERNELBASIS compute such bases, so that their use in a determinant algorithm typically leads to obtaining the sought determinant up to a constant factor; then finding the actual determinant requires to efficiently recover this constant [see e.g. 32, Sec. 4]. Since in this paper we seek determinants of matrices in reduced form, this issue is easily handled using the next result.

Lemma 3.1. *Let $s \in \mathbb{Z}^n$ and $\mathbf{A} \in \mathbb{K}[x]^{n \times n}$. If \mathbf{A} is in s -reduced form, the leading coefficient of $\det(\mathbf{A})$ is $\det(\text{Im}_s(\mathbf{A}))$. In particular, if \mathbf{A} is in s -weak Popov form, the leading coefficient of $\det(\mathbf{A})$ is the product of the leading coefficients of the diagonal entries of \mathbf{A} .*

Proof. The second claim is a direct consequence of the first, since for \mathbf{A} in s -weak Popov form, $\text{Im}_s(\mathbf{A})$ is lower triangular with diagonal entries equal to the leading coefficients of the diagonal entries of $\mathbf{A}\mathbf{X}^s$, which are the leading coefficients of the diagonal entries of \mathbf{A} . For the first claim in the case $s = \mathbf{0}$, we refer to [29, Sec. 6.3.2], and in particular Eq. (23) therein. Now, for an arbitrary s and \mathbf{A} in s -reduced form, we consider the nonnegative shift $t = s - (\min(s), \dots, \min(s))$ and observe that $\text{Im}_s(\mathbf{A}) = \text{Im}_0(\mathbf{A}\mathbf{X}^t)$, hence $\mathbf{A}\mathbf{X}^t$ is $\mathbf{0}$ -reduced, and thus the leading coefficient of $\det(\mathbf{A}\mathbf{X}^t) = \det(\mathbf{A}) \det(\mathbf{X}^t)$ (which is the same as that of $\det(\mathbf{A})$) is equal to $\det(\text{Im}_0(\mathbf{A}\mathbf{X}^t)) = \det(\text{Im}_s(\mathbf{A}))$. \square

From shifted to non-shifted. In Section 1.5, we explained that one step of our algorithm consists in finding $\det(\mathbf{K})$ for a matrix \mathbf{K} in ν -weak Popov form, and that it achieves this by computing the $-d$ -Popov form of \mathbf{K}^\top , which is already in $-d$ -weak Popov form. The next lemma substantiates this; note that in Section 1.5 we had left out the reversal matrix \mathbf{J}_n for the sake of exposition.

Lemma 3.2. *Let $\nu \in \mathbb{Z}^n$, let $\mathbf{K} \in \mathbb{K}[x]^{n \times n}$, and let $d = \text{rdeg}_\nu(\mathbf{K})$.*

- (a) *if $\text{Im}_\nu(\mathbf{K})$ has no zero column, then $\text{Im}_{-d}(\mathbf{K}^\top) = \text{Im}_\nu(\mathbf{K})^\top$ and $\text{rdeg}_{-d}(\mathbf{K}^\top) = -\nu$;*
- (b) *if \mathbf{K} is in ν -reduced form, then \mathbf{K}^\top is in $-d$ -reduced form;*
- (c) *if \mathbf{K} is in ν -weak Popov form, then $\mathbf{J}_n \mathbf{K}^\top \mathbf{J}_n$ is in $-d \mathbf{J}_n$ -weak Popov form;*
- (d) *if furthermore \mathbf{P} is the $-d \mathbf{J}_n$ -Popov form of $\mathbf{J}_n \mathbf{K}^\top \mathbf{J}_n$, then \mathbf{P}^\top is in $\mathbf{0}$ -weak Popov form and $\det(\mathbf{K}) = \det(\text{Im}_\nu(\mathbf{K})) \det(\mathbf{P}^\top)$.*

Proof. By definition, $\text{Im}_\nu(\mathbf{K})^\top$ is the coefficient of degree 0 of $(\mathbf{X}^{-d} \mathbf{K} \mathbf{X}^\nu)^\top = \mathbf{X}^\nu \mathbf{K}^\top \mathbf{X}^{-d}$, which is a matrix over $\mathbb{K}[x^{-1}]$. The assumption on $\text{Im}_\nu(\mathbf{K})$ implies that this coefficient of degree 0 has no zero row. It follows that $\text{rdeg}_{-d}(\mathbf{K}^\top) = -\nu$ and that this coefficient of degree 0 is $\text{Im}_{-d}(\mathbf{K}^\top)$. Item (b) follows from Item (a) by definition of shifted reduced forms.

From now on, we assume that \mathbf{K} is in ν -weak Popov form. Then $\text{Im}_\nu(\mathbf{K})$ is invertible and lower triangular, and in particular $\text{Im}_{-d}(\mathbf{K}^\top) = \text{Im}_\nu(\mathbf{K})^\top$. Since \mathbf{J}_n is a permutation matrix, we

obtain $\text{lm}_{-d}\mathbf{J}_n(\mathbf{J}_n\mathbf{K}^\top\mathbf{J}_n) = \mathbf{J}_n\text{lm}_{-d}(\mathbf{K}^\top)\mathbf{J}_n = \mathbf{J}_n\text{lm}_v(\mathbf{K})^\top\mathbf{J}_n$, which is invertible and lower triangular. Hence $\mathbf{J}_n\mathbf{K}^\top\mathbf{J}_n$ is in $-d\mathbf{J}_n$ -weak Popov form.

For Item (d), since \mathbf{P} is $n \times n$ and in $-d\mathbf{J}_n$ -Popov form, we have $\text{lm}_0(\mathbf{P}^\top) = \mathbf{I}_n$, hence \mathbf{P}^\top is in $\mathbf{0}$ -weak Popov form. Furthermore, since $\mathbf{J}_n\mathbf{K}^\top\mathbf{J}_n$ is unimodularly equivalent to \mathbf{P} , its determinant is $\det(\mathbf{K}) = \det(\mathbf{J}_n\mathbf{K}^\top\mathbf{J}_n) = \lambda \det(\mathbf{P})$ for some $\lambda \in \mathbb{K} \setminus \{0\}$. Applying Lemma 3.1 to \mathbf{P} shows that $\det(\mathbf{P})$ is monic, hence λ is the leading coefficient of $\det(\mathbf{K})$; applying the same lemma to \mathbf{K} yields $\lambda = \det(\text{lm}_v(\mathbf{K}))$. \square

3.2. Algorithm and correctness

Our main determinant algorithm is `DETERMINANTOFWEAKPOPOV` (Algorithm 2), which takes as input a matrix in $\mathbf{0}$ -weak Popov form and computes its determinant using recursive calls on matrices of smaller dimension. Then, we compute the determinant of a $\mathbf{0}$ -reduced matrix by first calling `REDUCEDTOWEAKPOPOV` to find a $\mathbf{0}$ -weak Popov matrix which has the same determinant up to a nonzero constant, and then calling the previous algorithm on that matrix. This is detailed in Algorithm 1.

Algorithm 1 `DETERMINANTOFREDUCED(A)`

Input: a matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times m}$ in $\mathbf{0}$ -reduced form.

Output: the determinant of \mathbf{A} .

- 1: $\mathbf{P} \in \mathbb{K}[x]^{m \times m} \leftarrow \text{REDUCEDTOWEAKPOPOV}(\mathbf{A}, \mathbf{0})$
 - 2: $\Delta \leftarrow \text{DETERMINANTOFWEAKPOPOV}(\mathbf{P}); \ell_\Delta \in \mathbb{K} \setminus \{0\} \leftarrow$ leading coefficient of Δ
 - 3: **return** $\det(\text{lm}_0(\mathbf{A})) \Delta / \ell_\Delta$
-

The correctness of Algorithm 1 is obvious: according to Theorem 1.3 and Proposition 3.3, \mathbf{P} is a $\mathbf{0}$ -weak Popov form of \mathbf{A} , and Δ is the determinant of \mathbf{P} up to multiplication by some element of $\mathbb{K} \setminus \{0\}$. Thus $\det(\mathbf{A}) = \ell \Delta$ for some $\ell \in \mathbb{K} \setminus \{0\}$, and Lemma 3.1 yields $\ell = \det(\text{lm}_0(\mathbf{A})) / \ell_\Delta$.

Concerning the cost bound, Theorem 1.3 states that the first step uses $O(m^\omega(1 + D/m))$ operations in \mathbb{K} , where $D = |\text{rdeg}_0(\mathbf{A})|$; since \mathbf{A} is $\mathbf{0}$ -reduced, this is $D = \deg(\det(\mathbf{A}))$ [29, Sec. 6.3.2]. In the last step, the determinant computation costs $O(m^\omega)$ operations, while scaling Δ by a constant costs $O(D)$ operations. The second step uses $O(m^\omega M'(D/m))$ operations according to Proposition 3.3, under the assumptions \mathcal{H}_{sl} , \mathcal{H}_{sm} , and \mathcal{H}_ω .

We now describe the main algorithm of this paper (Algorithm 2) and focus on its correctness. We also mention cost bounds for all steps of the algorithm that are not recursive calls, but we defer the core of the complexity analysis to Section 3.3.

Proposition 3.3. *Algorithm 2 is correct, and assuming that \mathcal{H}_{sl} , \mathcal{H}_{sm} , and \mathcal{H}_ω hold (hence in particular $\omega > 2$), it uses $O(m^\omega M'(D/m))$ operations in \mathbb{K} .*

Proof of correctness. The fact that \mathbf{A} is in $\mathbf{0}$ -weak Popov form has two consequences on the tuple s computed at Line 1: first, it is the $\mathbf{0}$ -pivot degree of \mathbf{A} (i.e. its diagonal degrees), and second, the sum $D = |s|$ is equal to the degree of the determinant of \mathbf{A} [29, Sec. 6.3.2].

The main base case of the recursion is when $m = 1$ and is handled at Line 2; it uses no operation in \mathbb{K} . We use a second base case at Line 3: if $D = 0$, then \mathbf{A} is an $m \times m$ matrix over \mathbb{K} . Since it is in $\mathbf{0}$ -weak Popov, it is invertible and lower triangular, hence $\det(\mathbf{A})$ is the product of its diagonal entries, which is computed in $O(m)$ multiplications in \mathbb{K} . This base case is not necessary for obtaining the correctness and the cost bound in Proposition 3.3; still, not using it would incur a cost of $O(m^\omega)$ operations in the case $D = 0$.

Algorithm 2 DETERMINANTOFWEAKPOPOV(**A**)

Input: a matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times m}$ in $\mathbf{0}$ -weak Popov form.

Output: the determinant of \mathbf{A} , up to multiplication by some element of $\mathbb{K} \setminus \{0\}$.

- 1: $s = (s_1, \dots, s_m) \in \mathbb{Z}_{\geq 0}^m \leftarrow \text{rdeg}_0(\mathbf{A})$; $D \leftarrow s_1 + \dots + s_m$ $\triangleright D = \text{degree of } \det(\mathbf{A})$
 - 2: **if** $m = 1$ **then return** the polynomial f such that $\mathbf{A} = [f]$ \triangleright base case: 1×1 matrix
 - 3: **if** $D = 0$ **then return** the product of diagonal entries of \mathbf{A} \triangleright base case: matrix over \mathbb{K}
 - 4: **if** $D < m$ **then** \triangleright handle constant rows to reduce to dimension $\leq D$
 - 5: $\mathbf{B} \leftarrow \mathbf{A} \text{lm}_0(\mathbf{A})^{-1}$ from which rows and columns with indices in $\{i \mid s_i = 0\}$ are removed
 - 6: **return** DETERMINANTOFWEAKPOPOV(\mathbf{B})
 - 7: **if** $s_1 + \dots + s_{\lfloor m/2 \rfloor} > D/2$ **then return** DETERMINANTOFWEAKPOPOV($\mathbf{J}_m \mathbf{A} \text{lm}_0(\mathbf{A})^{-1} \mathbf{J}_m$)
 - 8: Write $\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{A}_3 & \mathbf{A}_4 \end{bmatrix}$, with \mathbf{A}_1 of size $\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor$ and \mathbf{A}_4 of size $\lceil m/2 \rceil \times \lceil m/2 \rceil$
 - 9: $\mathbf{K} \in \mathbb{K}[x]^{\lceil m/2 \rceil \times m} \leftarrow \text{KERNELBASIS}(\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_3 \end{bmatrix}, s)$
 - 10: $[\mathbf{K}_1 \ \mathbf{K}_2] \in \mathbb{K}[x]^{\lceil m/2 \rceil \times m} \leftarrow \text{REDUCEDTOWEAKPOPOV}(\mathbf{K}, s)$, where \mathbf{K}_2 is $\lceil m/2 \rceil \times \lceil m/2 \rceil$
 - 11: $\mathbf{B} \leftarrow \text{UNBALANCEDMULTIPLICATION}([\mathbf{K}_1 \ \mathbf{K}_2], \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{A}_4 \end{bmatrix}, D)$ $\triangleright \mathbf{B} = \mathbf{K}_1 \mathbf{A}_2 + \mathbf{K}_2 \mathbf{A}_4$
 - 12: $\Delta_1 \leftarrow \text{DETERMINANTOFWEAKPOPOV}(\mathbf{B})$ \triangleright first recursive call
 - 13: $\Delta_2 \leftarrow \text{DETERMINANTOFWEAKPOPOV}(\mathbf{A}_1)$ \triangleright second recursive call
 - 14: $\mathbf{P} \leftarrow \text{WEAKPOPOVTOPOPOV}(\mathbf{J}_{\lfloor m/2 \rfloor} \mathbf{K}_2^T \mathbf{J}_{\lfloor m/2 \rfloor}, \text{rdeg}_{(s_{\lfloor m/2 \rfloor + 1}, \dots, s_m)}(\mathbf{K}_2) \mathbf{J}_{\lfloor m/2 \rfloor})$
 - 15: $\Delta_3 \leftarrow \text{DETERMINANTOFWEAKPOPOV}(\mathbf{P}^T)$ \triangleright third recursive call
 - 16: **return** $\Delta_1 \Delta_2 / \Delta_3$
-

For the recursion, we proceed inductively: we assume that the algorithm correctly computes the determinant for all $\mathbf{0}$ -weak Popov matrices of dimension less than m , and based on this we show that it is also correct for any $\mathbf{0}$ -weak Popov matrix \mathbf{A} of dimension m .

Case 1: $D < m$. Then \mathbf{A} has at least one constant row; using linear algebra we reduce to the case of a matrix \mathbf{B} of dimension at most D with all rows of degree at least 1. Since $s = \text{rdeg}_0(\mathbf{A})$, we can write $\mathbf{A} = \mathbf{X}^s \text{lm}_0(\mathbf{A}) + \mathbf{R}$ for a matrix $\mathbf{R} \in \mathbb{K}[x]^{m \times m}$ such that $\text{rdeg}_0(\mathbf{R}) < s$. Since \mathbf{A} is $\mathbf{0}$ -reduced, $\text{lm}_0(\mathbf{A})$ is invertible and $\mathbf{A} \text{lm}_0(\mathbf{A})^{-1} = \mathbf{X}^s + \mathbf{R} \text{lm}_0(\mathbf{A})^{-1}$ with $\text{rdeg}_0(\mathbf{R} \text{lm}_0(\mathbf{A})^{-1}) < s$, which implies $\text{lm}_0(\mathbf{A} \text{lm}_0(\mathbf{A})^{-1}) = \mathbf{I}_m$. In particular, $\mathbf{A} \text{lm}_0(\mathbf{A})^{-1}$ is in $\mathbf{0}$ -weak Popov form, and for each i such that the row $\mathbf{A}_{i,*}$ is constant, i.e. $s_i = 0$, the i th row of $\mathbf{A} \text{lm}_0(\mathbf{A})^{-1}$ is the i th row of the identity matrix. Therefore the matrix \mathbf{B} at Line 5 is in $\mathbf{0}$ -weak Popov form, has the same determinant as \mathbf{A} up to a constant, and has dimension $\#\{i \mid s_i \neq 0\} \leq D$. Hence the correctness in this case. In terms of complexity, computing \mathbf{B} essentially amounts to computing the product $\mathbf{A} \text{lm}_0(\mathbf{A})^{-1}$, which is done by row-wise expanding \mathbf{A} into a $(m + D) \times m$ matrix over \mathbb{K} , right-multiplying by $\text{lm}_0(\mathbf{A})^{-1}$, and compressing the result back into a polynomial matrix: this costs $O(m^\omega(1 + D/m)) \subseteq O(m^\omega)$ operations.

Case 2: $s_1 + \dots + s_{\lfloor m/2 \rfloor} > D/2$. Then we modify the input \mathbf{A} so as to reduce to *Case 3*. As we have seen above, $\text{lm}_0(\mathbf{A} \text{lm}_0(\mathbf{A})^{-1}) = \mathbf{I}_m$. We now reverse the diagonal entries by reversing the order of rows and columns: let $\mathbf{B} = \mathbf{J}_m \mathbf{A} \text{lm}_0(\mathbf{A})^{-1} \mathbf{J}_m$. Then $\text{lm}_0(\mathbf{B}) = \mathbf{J}_m \text{lm}_0(\mathbf{A} \text{lm}_0(\mathbf{A})^{-1}) \mathbf{J}_m = \mathbf{I}_m$, hence \mathbf{B} is in $\mathbf{0}$ -weak Popov form: Line 7 calls the algorithm on this matrix to obtain $\det(\mathbf{B})$ up to a constant, and this yields $\det(\mathbf{A})$ since it is equal to $\det(\text{lm}_0(\mathbf{A})) \det(\mathbf{B})$. To conclude the proof of correctness in that case (assuming correctness in *Case 3*), it remains to observe that \mathbf{B} has the same matrix dimension m as \mathbf{A} , and that the matrix \mathbf{B} has degrees such that calling the algorithm with input \mathbf{B} does not enter *Case 2* but *Case 3*. Indeed, we have $\text{rdeg}_0(\mathbf{B}) = s \mathbf{J}_m$, hence the sum of the first $\lfloor m/2 \rfloor$ entries of the tuple $\text{rdeg}_0(\mathbf{B})$ is $s_m + \dots + s_{\lfloor m/2 \rfloor + 1} = D - (s_1 + \dots + s_{\lfloor m/2 \rfloor})$, which is at most $D/2$ by assumption. In terms of complexity, the main step is to compute the

product $\mathbf{A} \operatorname{Im}_0(\mathbf{A})^{-1}$, which costs $O(m^\omega(1 + D/m))$ operations as we have seen above; this is in $O(m^\omega M'(D/m))$.

Case 3: $s_1 + \dots + s_{\lfloor m/2 \rfloor} \leq D/2$. Then, Line 7 performs no action, and Line 8 defines submatrices of \mathbf{A} . By construction, $\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_3 \end{bmatrix}$ has full column rank and $s \geq \operatorname{rdeg}_0(\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_3 \end{bmatrix})$ holds. Thus, according to Lemma 2.10, Line 9 uses $O(m^\omega M'(D/m))$ operations to compute an s -reduced basis \mathbf{K} of $\mathcal{K}(\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_3 \end{bmatrix})$, with $|\operatorname{rdeg}_s(\mathbf{K})| \leq D$. Then, Theorem 1.3 states that Line 10 transforms \mathbf{K} into an s -weak Popov basis $[\mathbf{K}_1 \ \mathbf{K}_2]$ of this kernel at a cost of $O(m^\omega(1 + D/m))$ operations, since $|\operatorname{rdeg}_s(\mathbf{K})| \leq D$ and $\min(s) \geq 0$. Since all s -reduced bases of $\mathcal{K}(\mathbf{F})$ have the same s -row degree up to permutation, $|\operatorname{rdeg}_s([\mathbf{K}_1 \ \mathbf{K}_2])| \leq D$ holds, hence the assumptions of Lemma 2.8 are satisfied and Line 11 uses $O(m^\omega M(D/m))$ operations to compute $\mathbf{B} = \mathbf{K}_1 \mathbf{A}_2 + \mathbf{K}_2 \mathbf{A}_4$.

The important observation at this stage is the identity

$$\begin{bmatrix} \mathbf{I}_{\lfloor m/2 \rfloor} & \mathbf{0} \\ \mathbf{K}_1 & \mathbf{K}_2 \end{bmatrix} \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{A}_3 & \mathbf{A}_4 \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{0} & \mathbf{B} \end{bmatrix} \quad (2)$$

which, provided that \mathbf{K}_2 is nonsingular, implies $\det(\mathbf{A}) = \det(\mathbf{B}) \det(\mathbf{A}_1) / \det(\mathbf{K}_2)$. We are going to show that this is the formula used in Line 16 to compute $\det(\mathbf{A})$.

First, \mathbf{A}_1 has dimension less than m and, being a principal submatrix of the $\mathbf{0}$ -weak Popov matrix \mathbf{A} , it is also in $\mathbf{0}$ -weak Popov form. Hence the recursive call at Line 13 is sound and Δ_2 is equal to $\det(\mathbf{A}_1)$ up to a constant.

Since \mathbf{A} is in $\mathbf{0}$ -weak Popov form and $\begin{bmatrix} \mathbf{I}_{\lfloor m/2 \rfloor} & \mathbf{0} \\ \mathbf{K}_1 & \mathbf{K}_2 \end{bmatrix}$ is in $\operatorname{rdeg}_0(\mathbf{A})$ -weak Popov form, their product $\begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{0} & \mathbf{B} \end{bmatrix}$ is in $\mathbf{0}$ -weak Popov form; see Lemma 2.6, or note that $\operatorname{Im}_0(\begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{0} & \mathbf{B} \end{bmatrix})$ is invertible and lower triangular according to Lemma 2.2. It follows that \mathbf{B} is in $\mathbf{0}$ -weak Popov form and has dimension less than m : Line 12 recursively computes Δ_1 , equal to $\det(\mathbf{B})$ up to a constant.

It remains to prove that Δ_3 computed at Lines 14 and 15 is equal to $\det(\mathbf{K}_2)$ up to a constant. Let $\nu = \operatorname{rdeg}_0(\mathbf{A}_4) = (s_{\lfloor m/2 \rfloor + 1}, \dots, s_m)$ be the shift used at Line 14, and let $\mathbf{d} = \operatorname{rdeg}_\nu(\mathbf{K}_2) = \operatorname{rdeg}_s([\mathbf{K}_1 \ \mathbf{K}_2])$. Applying Lemma 2.7 (with $\mathbf{F} = \mathbf{A}_1$, $\mathbf{G} = \mathbf{A}_3$, $\mathbf{t} = \mathbf{0}$, and ν as above) shows that $[\mathbf{K}_1 \ \mathbf{K}_2]$ has all its s -pivot entries in \mathbf{K}_2 , and in particular \mathbf{K}_2 is in ν -weak Popov form. Let $\delta \in \mathbb{Z}_{\geq 0}^n$ be the ν -pivot degree of \mathbf{K}_2 , where $n = \lfloor m/2 \rfloor$, and note that $\mathbf{d} = \delta + \nu \geq \delta$ since $\nu \geq \mathbf{0}$. Then, Lemma 3.2 states that $\mathbf{J}_n \mathbf{K}_2^\top \mathbf{J}_n$ is in $-\mathbf{d} \mathbf{J}_n$ -weak Popov form; its $-\mathbf{d} \mathbf{J}_n$ -pivot degree is the list of degrees of its diagonal entries, that is, $\delta \mathbf{J}_n$. Since $\mathbf{d} \mathbf{J}_n \geq \delta \mathbf{J}_n$, we can apply Theorem 1.4, which implies that Line 14 computes the $-\mathbf{d} \mathbf{J}_n$ -Popov form \mathbf{P} of $\mathbf{J}_n \mathbf{K}_2^\top \mathbf{J}_n$ using $O(m^\omega M'(|\mathbf{d}|/m))$ operations; as we have seen above, $|\mathbf{d}| = |\operatorname{rdeg}_s([\mathbf{K}_1 \ \mathbf{K}_2])| \leq D$. Then, from the last item of Lemma 3.2, \mathbf{P}^\top is in $\mathbf{0}$ -weak Popov form and $\det(\mathbf{K}) = \det(\operatorname{Im}_\nu(\mathbf{K})) \det(\mathbf{P}^\top)$, hence Line 15 correctly computes $\det(\mathbf{K})$ up to a constant. \square

To conclude this presentation of our determinant algorithm, we note that it would be beneficial, in a practical implementation, to add an early exit. Precisely, just after computing $\det(\mathbf{B})$ at Line 12, one could perform the following action before (possibly) proceeding to the next steps:

12b: **if** $\deg(\Delta_1) = D$ **then return** Δ_1 \triangleright *early exit*

Indeed, recall that Δ_1 is $\det(\mathbf{B})$ up to a constant; furthermore we claim that

- for a generic \mathbf{A} , we have $\deg(\Delta_1) = D$,
- if $\deg(\Delta_1) = D$ (i.e. $\deg(\det(\mathbf{B})) = D$), then $\det(\mathbf{B})$ is $\det(\mathbf{A})$ up to a constant.

It follows that for a generic matrix \mathbf{A} , then Δ_1 is $\det(\mathbf{A})$ up to a constant, hence the correctness of this early exit (see also [21, Sec. 4.2.2] for similar considerations). To prove the above claim,

first note that since $[\mathbf{K}_1 \ \mathbf{K}_2]$ is a kernel basis, it has unimodular column bases [22, Lem. 2.2], and thus it can be completed into a unimodular matrix $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_2 \\ \mathbf{K}_1 & \mathbf{K}_2 \end{bmatrix} \in \mathbb{K}[x]^{m \times m}$ [57, Lem. 2.10]. Therefore

$$\mathbf{U}\mathbf{A} = \begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_2 \\ \mathbf{K}_1 & \mathbf{K}_2 \end{bmatrix} \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{A}_3 & \mathbf{A}_4 \end{bmatrix} = \begin{bmatrix} \mathbf{B}_1 & \mathbf{B}_2 \\ \mathbf{0} & \mathbf{B} \end{bmatrix}$$

where $[\mathbf{B}_1 \ \mathbf{B}_2] = [\mathbf{U}_1 \ \mathbf{U}_2]\mathbf{A}$. Since $\det(\mathbf{U})$ is in $\mathbb{K} \setminus \{0\}$, $\det(\mathbf{A})$ is $\det(\mathbf{B}_1) \det(\mathbf{B})$ up to a constant. For the second item of the claim, $\deg(\Delta_1) = D$ implies $\deg(\det(\mathbf{B})) = D = \deg(\det(\mathbf{A}))$, hence $\det(\mathbf{B}_1)$ is in $\mathbb{K} \setminus \{0\}$. The first item follows from the fact that \mathbf{B}_1 is a row basis of $\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_3 \end{bmatrix}$ [56, Lem. 3.1]; since the latter matrix has more rows than columns, if \mathbf{A}_1 and \mathbf{A}_3 have generic entries, then such a row basis \mathbf{B}_1 is unimodular which means $\det(\mathbf{B}_1) \in \mathbb{K} \setminus \{0\}$ and thus $\deg(\Delta_1) = D$.

3.3. Complexity analysis

We have seen above that all computations in Algorithm 2 other than recursive calls have an arithmetic cost in $O(m^\omega M'(D/m))$ operations in \mathbb{K} ; here, we complete the proof of the cost bound in Proposition 5.14. In this section, we use the assumptions \mathcal{H}_{sl} , \mathcal{H}_{sm} , and \mathcal{H}_ω as well as their consequences stated in Section 1.1.

Let $C(m, D)$ denote the arithmetic cost of Algorithm 2; recall that D is the degree of the determinant of the input, which is also the sum of its row degrees. First consider the case $m \leq D$. If $s_1 + \dots + s_{\lfloor m/2 \rfloor} > D/2$, the reduction to the case $s_1 + \dots + s_{\lfloor m/2 \rfloor} \leq D/2$ with the same m and D performed at Line 7 costs $O(m^\omega(1 + D/m))$. Once we are in the latter case, there are three recursive calls with input matrices having the following dimensions and degrees:

- At Line 12, the matrix \mathbf{B} is $\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor$, and applying the predictable degree property on Eq. (2) gives in particular $\text{rdeg}_0(\mathbf{B}) = \text{rdeg}_s([\mathbf{K}_1 \ \mathbf{K}_2])$, hence $|\text{rdeg}_0(\mathbf{B})| \leq D$.
- At Line 13, the matrix \mathbf{A}_1 is $\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor$ and the sum of its row degrees is $s_1 + \dots + s_{\lfloor m/2 \rfloor}$, which is at most $D/2$ by assumption.
- At Line 15, the matrix \mathbf{P}^\top is $\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor$ and its $\mathbf{0}$ -pivot degree is $\text{rdeg}_0(\mathbf{P}^\top) = \delta \mathbf{J}_m$. Recall indeed that this is the list of diagonal degrees of \mathbf{P}^\top , which is the same as that of \mathbf{P} , and thus the same as that of $\mathbf{J}_n \mathbf{K}_2^\top \mathbf{J}_n$ according to Lemma 2.5. Now, from $|\delta + \nu| = |\delta| + |\nu| \leq D$ and the assumption $|\nu| = s_{\lfloor m/2 \rfloor + 1} + \dots + s_m > D/2$, we obtain $|\text{rdeg}_0(\mathbf{P}^\top)| = |\delta| \leq D/2$.

We assume without loss of generality that m is a power of 2. If it is not, a given input matrix can be padded with zeros, and ones on the main diagonal, so as to form a square matrix with dimension the next power of two and the same determinant. According to the three items above, the cost bound then satisfies:

$$C(m, D) \leq 2C(m/2, \lfloor D/2 \rfloor) + C(m/2, D) + O(m^\omega M'(D/m)).$$

Letting the $O(\cdot)$ term aside, we illustrate this recurrence relation in Fig. 1.

Let $\mu = \log_2(m)$ and let K be the constant of the $O(\cdot)$ term above. Recalling that $m \leq D$ and $\lfloor \lfloor D/2^j \rfloor / 2 \rfloor = \lfloor D/2^{j+1} \rfloor$, unrolling this recurrence to the i th recursion level for $0 \leq i \leq \mu$ yields

$$\begin{aligned} C(m, D) &\leq \sum_{j=0}^i a_{i,j} C\left(\frac{m}{2^i}, \left\lfloor \frac{D}{2^j} \right\rfloor\right) + K \sum_{k=0}^{i-1} \sum_{j=0}^k a_{k,j} \left(\frac{m}{2^k}\right)^\omega M'\left(\frac{D/2^j}{m/2^k}\right) \\ &\leq \sum_{j=0}^i a_{i,j} C\left(\frac{m}{2^i}, \left\lfloor \frac{D}{2^j} \right\rfloor\right) + Km^\omega M'\left(\frac{D}{m}\right) \left(\sum_{k=0}^{i-1} \sum_{j=0}^k a_{k,j} 2^{-k\omega} M'(2^{k-j}) \right), \end{aligned}$$

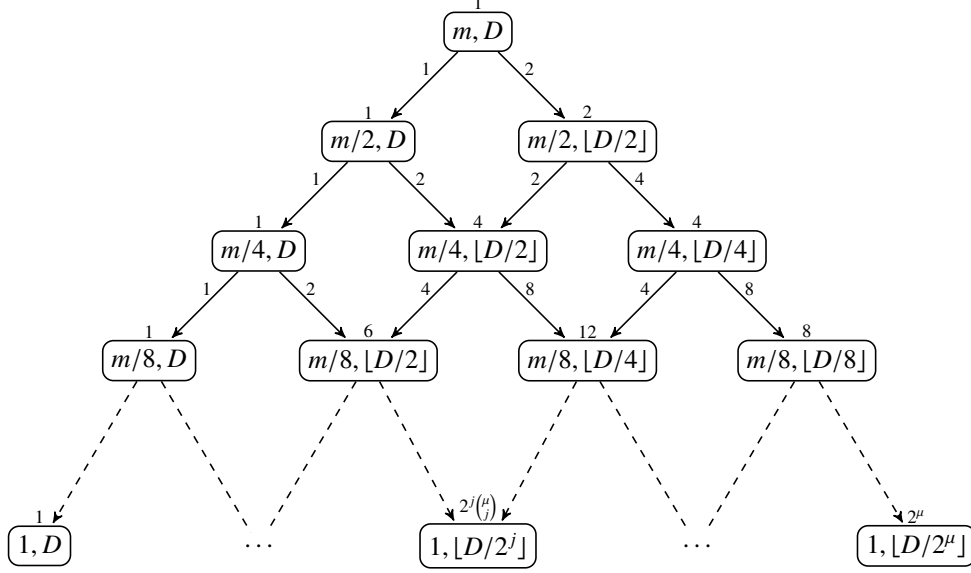


Figure 1: Directed acyclic graph of recursive calls, of depth $\mu = \log_2(m)$. Each boxed node shows the matrix dimensions and the determinantal degree of a recursive call. Beginning with one call in dimension and determinantal degree (m, D) , for a given node the number above it indicates the number of times a recursive call corresponding to this node is made, and the numbers of recursive sub-calls this node generates are indicated on both arrows starting from this node.

where the last inequality comes from the submultiplicativity of $M'(\cdot)$ and the coefficients $a_{i,j}$ satisfy

$$\begin{cases} a_{i,0} = 1, \\ a_{i,i} = 2^i, \\ a_{i,j} = a_{i-1,j} + 2a_{i-1,j-1} \text{ for } 0 < j < i. \end{cases}$$

In Fig. 1, one can observe the similarity between Pascal's triangle and the number of calls with parameters $(m/2^i, D/2^j)$. This translates as a connection between the $a_{i,j}$'s and the binomial coefficients: one can prove by induction that $a_{i,j} = 2^j \binom{i}{j}$.

Now, by assumption $M'(d) \in O(d^{\omega-1-\varepsilon})$ for some $\varepsilon > 0$; for such an $\varepsilon > 0$, let \tilde{K} be a constant such that $M'(d) \leq \tilde{K}d^{\omega-1-\varepsilon}$ for all $d \geq 0$. Then

$$\sum_{k=0}^{i-1} \sum_{j=0}^k a_{k,j} 2^{-k\omega} M'(2^{k-j}) \leq \tilde{K} \sum_{k=0}^{i-1} 2^{-(1+\varepsilon)k} \sum_{j=0}^k \frac{a_{k,j}}{2^{j(\omega-1-\varepsilon)}} \leq \tilde{K} \sum_{k=0}^{i-1} 2^{-(1+\varepsilon)k} \sum_{j=0}^k \binom{k}{j} = \tilde{K} \sum_{k=0}^{i-1} 2^{-\varepsilon k}$$

and, defining the constant $\hat{K} = \tilde{K} \sum_{k=0}^{+\infty} 2^{-\varepsilon k}$, for $i = \mu$ we obtain

$$C(m, D) \leq \sum_{j=0}^{\mu} a_{\mu,j} C\left(1, \left\lfloor \frac{D}{2^j} \right\rfloor\right) + \hat{K} m^{\omega} M'\left(\frac{D}{m}\right).$$

As we have seen above, parameters $(1, d)$ for any $d \in \mathbb{Z}_{\geq 0}$ correspond to base cases with a 1×1 matrix, and they incur no arithmetic cost (one might want to consider them to use $O(1)$ operations

each; then the total cost of these base cases is bounded asymptotically by $\sum_{j=0}^{\mu} a_{\mu,j} \leq m^{\log_2(3)}$. Thus we obtain $C(m, D) = O(m^\omega M'(D/m))$, under the assumption $m \leq D$.

For the case $D < m$ handled at Line 4, Section 3.2 showed that $C(m, D) = C(\hat{m}, D) + O(m^\omega)$ where \hat{m} is the number of non-constant rows of the input matrix. Since $\hat{m} \leq D$, our proof above shows that $C(\hat{m}, D)$ is in $O(\hat{m}^\omega M'(D/\hat{m}))$. Now our assumptions on $M(\cdot)$ imply in particular that $M'(\cdot)$ is subquadratic, hence this bound is in $O(\hat{m}^\omega(D/\hat{m})^2) = O(\hat{m}^{\omega-2}D^2) \subseteq O(m^\omega)$. This concludes the complexity analysis.

4. Shifted forms: from reduced to weak Popov

This section proves Theorem 1.3 by generalizing the approach of [42, Sec. 2 and 3], which focuses on the non-shifted case $s = \mathbf{0}$. It first uses Gaussian elimination on the $\mathbf{0}$ -leading matrix of \mathbf{A} to find a unimodular matrix \mathbf{U} such that \mathbf{UA} is in $\mathbf{0}$ -weak Popov form, and then exploits the specific form of \mathbf{U} to compute \mathbf{UA} efficiently. Here we extend this approach to arbitrary shifts and show how to take into account the possible unbalancedness of the row degree of \mathbf{A} .

First, we generalize [42, Lem. 8] from $s = \mathbf{0}$ to an arbitrary s , by describing how \mathbf{U} can be obtained by computing a $\mathbf{0}$ -weak Popov form of the s -leading matrix of \mathbf{A} .

Lemma 4.1. *Let $s \in \mathbb{Z}^n$ and let $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$ be s -reduced with $\text{rdeg}_s(\mathbf{A})$ nondecreasing. There exists an invertible lower triangular $\mathbf{T} \in \mathbb{K}^{m \times m}$ which can be computed in $O(m^{\omega-1}n)$ operations in \mathbb{K} and is such that $\mathbf{T} \text{lm}_s(\mathbf{A})$ is in $\mathbf{0}$ -unordered weak Popov form. For any such matrix \mathbf{T} ,*

- $\mathbf{U} = \mathbf{X}^t \mathbf{T} \mathbf{X}^{-t}$ has polynomial entries and is unimodular, where $\mathbf{t} = \text{rdeg}_s(\mathbf{A})$,
- $\text{rdeg}_s(\mathbf{UA}) = \text{rdeg}_s(\mathbf{A})$ and $\text{lm}_s(\mathbf{UA}) = \mathbf{T} \text{lm}_s(\mathbf{A})$,
- \mathbf{UA} is in s -unordered weak Popov form.

Proof. Consider the matrix $\text{lm}_s(\mathbf{A})\mathbf{J}_n$ and its generalized Bruhat decomposition $\text{lm}_s(\mathbf{A})\mathbf{J}_n = \mathbf{CPR}$ as defined in [35]: $\mathbf{C} \in \mathbb{K}^{m \times m}$ is in column echelon form, $\mathbf{R} \in \mathbb{K}^{m \times n}$ is in row echelon form, and $\mathbf{P} \in \mathbb{K}^{m \times m}$ is a permutation matrix. Therefore $\mathbf{J}_m \mathbf{R} \mathbf{J}_n$ is in $\mathbf{0}$ -weak Popov form (see the paragraph before Lemma 2.4 in Section 2.4), and \mathbf{PRJ}_n is in $\mathbf{0}$ -unordered weak Popov form. Since $\text{lm}_s(\mathbf{A})$ has full row rank, \mathbf{C} is lower triangular and invertible, hence $\mathbf{PRJ}_n = \mathbf{C}^{-1} \text{lm}_s(\mathbf{A})$ which proves the existence of $\mathbf{T} = \mathbf{C}^{-1}$. Computing the decomposition costs $O(m^{\omega-1}n)$ operations [13, Cor. 25] while inverting \mathbf{C} costs $O(m^\omega)$ operations. Alternatively, [42, Sec. 3] shows how to compute \mathbf{T} within the same cost bound using an LUP decomposition with a modified pivoting strategy.

For any such matrix \mathbf{T} , write $\mathbf{T} = (T_{ij})_{ij}$ and $\mathbf{t} = (t_i)_i$. Then the entry (i, j) of \mathbf{U} is $T_{ij}x^{t_i-t_j}$. Thus, \mathbf{U} is lower triangular in $\mathbb{K}(x)^{m \times m}$ with diagonal entries in $\mathbb{K} \setminus \{0\}$, and since $\mathbf{t} = \text{rdeg}_s(\mathbf{A})$ is nondecreasing, \mathbf{U} is a unimodular matrix in $\mathbb{K}[x]^{m \times m}$.

Now, consider the nonnegative shift $\mathbf{u} = \mathbf{t} - (\min(\mathbf{t}), \dots, \min(\mathbf{t})) \in \mathbb{Z}_{\geq 0}^m$, and note that $\mathbf{U} = \mathbf{X}^{\mathbf{u}} \mathbf{T} \mathbf{X}^{-\mathbf{u}}$. (The introduction of \mathbf{u} is to circumvent the fact that we have not defined the row degree of a matrix over the Laurent polynomials, a notion which would be needed if we used \mathbf{X}^t rather than $\mathbf{X}^{\mathbf{u}}$ in Eq. (4).) Since \mathbf{A} is in s -reduced form, the predictable degree property yields

$$\text{rdeg}_s(\mathbf{UA}) = \text{rdeg}_t(\mathbf{U}) = \text{rdeg}_{\mathbf{u}}(\mathbf{U}) + (\min(\mathbf{t}), \dots, \min(\mathbf{t})). \quad (3)$$

On the other hand,

$$\text{rdeg}_{\mathbf{u}}(\mathbf{U}) = \text{rdeg}_0(\mathbf{UX}^{\mathbf{u}}) = \text{rdeg}_0(\mathbf{X}^{\mathbf{u}} \mathbf{T}) = \mathbf{u} = \text{rdeg}_s(\mathbf{A}) - (\min(\mathbf{t}), \dots, \min(\mathbf{t})), \quad (4)$$

where the third equality follows from the fact that \mathbf{T} is a constant matrix with no zero row. Then, from Eqs. (3) and (4), we obtain $\text{rdeg}_s(\mathbf{UA}) = \text{rdeg}_s(\mathbf{A}) = \mathbf{t}$.

Then, $\text{Im}_s(\mathbf{UA})$ is formed by the coefficients of nonnegative degree of

$$\mathbf{X}^{-\text{rdeg}_s(\mathbf{UA})} \mathbf{UAX}^s = \mathbf{X}^{-t} \mathbf{UAX}^s = \mathbf{TX}^{-t} \mathbf{AX}^s.$$

Since \mathbf{T} is constant and $\text{Im}_s(\mathbf{A})$ is formed by the coefficients of nonnegative degree of $\mathbf{X}^{-t} \mathbf{AX}^s$, we obtain that $\text{Im}_s(\mathbf{UA}) = \mathbf{T} \text{Im}_s(\mathbf{A})$. The third item then directly follows from Lemma 2.4. \square

Knowing \mathbf{T} , and therefore \mathbf{U} , the remaining difficulty is to efficiently compute \mathbf{UA} . For this, in the case $s = \mathbf{0}$, the approach in [42] has a cost bound which involves the maximum degree $d = \deg(\mathbf{A}) = \max(\text{rdeg}_0(\mathbf{A}))$ [42, Thm. 13], and uses the following steps:

- first compute $x^d \mathbf{X}^{-\text{rdeg}_0(\mathbf{A})} \mathbf{UA} = \mathbf{T}(\mathbf{X}^{(d, \dots, d) - \text{rdeg}_0(\mathbf{A})} \mathbf{A})$;
- then scale the rows via the left multiplication by $x^{-d} \mathbf{X}^{\text{rdeg}_0(\mathbf{A})}$.

While the scaling does not use arithmetic operations, the first step asks to multiply the constant matrix \mathbf{T} by an $m \times n$ polynomial matrix of degree d : this costs $O(m^{\omega-1}nd)$ operations in \mathbb{K} . Such a cost would not allow us to reach our target bound for the computation of characteristic polynomials, since d may be too large, namely when \mathbf{A} has unbalanced row degrees.

Below, we refine the approach to better conform with the row degrees of \mathbf{A} . This leads to an improvement of the above cost bound to $O(m^{\omega-1}n(1 + D/m))$ where $D = |\text{rdeg}_0(\mathbf{A})|$, thus involving the average row degree of \mathbf{A} instead of its maximum degree d . For this, we follow the strategy of splitting the rows of \mathbf{A} into subsets having degrees in prescribed intervals; when these intervals get further away from the average row degree of \mathbf{A} , the corresponding subset contains a smaller number of rows. Earlier works using a similar strategy such as [54], [58, Sec. 2.6], and [22, Sec. 4], are not directly applicable to the problem here. Furthermore, we exploit the fact that \mathbf{A} has nondecreasing row degree and that \mathbf{T} has a lower triangular shape to avoid logarithmic factors in the cost bound. This results in Algorithm 3.

Algorithm 3 REDUCEDTOWEAKPOPOV(\mathbf{A}, s)

Input: a matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$, a shift $s \in \mathbb{Z}^n$ such that \mathbf{A} is in s -reduced form.

Output: an s -weak Popov form of \mathbf{A} .

- 1: \triangleright Step 1: Ensure nonnegative shift and nondecreasing s -row degree \triangleleft
 - 2: $\hat{s} \in \mathbb{Z}_{\geq 0}^n \leftarrow s - (\min(s), \dots, \min(s))$
 - 3: $(\mathbf{B}, \mathbf{t}) \in \mathbb{K}[x]^{m \times n} \times \mathbb{Z}_{\geq 0}^m \leftarrow$ matrix \mathbf{B} obtained from \mathbf{A} by row permutation such that the tuple $\mathbf{t} = \text{rdeg}_{\hat{s}}(\mathbf{B})$ is nondecreasing
 - 4: \triangleright Step 2: Compute the factor \mathbf{T} in the unimodular transformation $\mathbf{U} = \mathbf{X}^t \mathbf{TX}^{-t}$ \triangleleft
 - 5: $\mathbf{L} \in \mathbb{K}^{m \times n} \leftarrow \text{Im}_{\hat{s}}(\mathbf{B})$, that is, the entry (i, j) of \mathbf{L} is the coefficient of degree $t_i - s_j$ of the entry (i, j) of \mathbf{B} , where $\hat{s} = (s_1, \dots, s_n)$ and $\mathbf{t} = (t_1, \dots, t_m)$
 - 6: $\mathbf{T} \in \mathbb{K}^{m \times m} \leftarrow$ invertible lower triangular matrix such that \mathbf{TL} is in $\mathbf{0}$ -unordered weak Popov form \triangleright can be computed via a generalized Bruhat decomposition, see Lemma 4.1
 - 7: \triangleright Step 3: Compute the product $\mathbf{P} = \mathbf{UB}$ \triangleleft
 - 8: $D \leftarrow t_1 + \dots + t_m$; $K \leftarrow \lfloor \log_2(mt_m/D) \rfloor + 1$ $\triangleright K = \min\{k \in \mathbb{Z}_{>0} \mid t_m < 2^k D/m\}$
 - 9: $i_0 \leftarrow 0$; $i_k \leftarrow \max\{i \mid t_i < 2^k D/m\}$ for $1 \leq k \leq K$ $\triangleright 0 = i_0 < i_1 \leq \dots \leq i_{K-1} < i_K = m$
 - 10: $\mathbf{P} \leftarrow$ zero matrix in $\mathbb{K}[x]^{m \times n}$
 - 11: **for** k from 1 to K **do**
 - 12: $\mathcal{R} \leftarrow \{i_{k-1} + 1, \dots, m\}$; $\mathcal{C} \leftarrow \{i_{k-1} + 1, \dots, i_k\}$; $\theta \leftarrow t_{i_k} = \max(\mathbf{t}_{\mathcal{C}})$
 - 13: $\mathbf{P}_{\mathcal{R},*} \leftarrow \mathbf{P}_{\mathcal{R},*} + \mathbf{X}^{t_{\mathcal{R}} - (\theta, \dots, \theta)} \mathbf{T}_{\mathcal{R},\mathcal{C}} \mathbf{X}^{(\theta, \dots, \theta) - t_{\mathcal{C}}} \mathbf{B}_{\mathcal{C},*}$
 - 14: **return** the row permutation of \mathbf{P} which has increasing \hat{s} -pivot index
-

$\mathbf{X}^{t_{\mathcal{R}_k} - (\theta, \dots, \theta)}$ ($\mathbf{T}_{\mathcal{R}_k, C_k}(\mathbf{X}^{(\theta, \dots, \theta) - t_{C_k}} \mathbf{B}_{C_k, *})$) uses $O((m/2^{k-1})^{\omega-1} n(\theta+1)) \subseteq O((m/2^{k-1})^{\omega-1} n(2^k D/m+1))$ operations in \mathbb{K} . Thus, since $\omega > 2$, the cost of Step 3 is

$$\begin{aligned} O\left(\sum_{1 \leq k \leq K} \left(\frac{m}{2^{k-1}}\right)^{\omega-1} n\left(\frac{2^k D}{m} + 1\right)\right) &\subseteq O\left(m^{\omega-2} n D \left(\sum_{1 \leq k \leq K} 2^{k(2-\omega)}\right) + m^{\omega-1} n \left(\sum_{1 \leq k \leq K} 2^{k(1-\omega)}\right)\right) \\ &\subseteq O\left(m^{\omega-2} n D + m^{\omega-1} n\right). \quad \square \end{aligned}$$

5. Shifted forms: from weak Popov to Popov

This section culminates in Section 5.5 with the description of Algorithm WEAKPOPOVTOPOPOV and a proof of Theorem 1.4. Based on [42, Lem. 14] (which extends to the shifted case), the result in Theorem 1.4 can easily be used to solve the same problem in the rectangular case with an $m \times n$ matrix \mathbf{A} ; while this is carried out in Algorithm WEAKPOPOVTOPOPOV, it is out of the main scope of this paper and thus for conciseness we only give a cost analysis in the case $m = n$.

Our approach is to obtain the $-s$ -Popov form of \mathbf{A} from a shifted reduced kernel basis of some matrix \mathbf{F} built from \mathbf{A} . This fact is substantiated in Section 5.1, which also proves that we have precise a priori information on the degrees of this sought kernel basis.

A folklore method for kernel basis computation is to find an approximant basis at an order sufficiently large so that it contains a kernel basis as a submatrix. More precisely, assuming we know a list of bounds s such that there exists a basis of $\mathcal{K}(\mathbf{F})$ with column degree bounded by s , the following algorithm computes such a kernel basis which is furthermore $-s$ -reduced:

- $\gamma \leftarrow \text{cdeg}_s(\mathbf{F}) + 1$;
- $\mathbf{M} \leftarrow$ basis of $\mathcal{A}_\gamma(\mathbf{F})$ in $-s$ -reduced form;
- **return** the submatrix of \mathbf{M} formed by its rows with nonpositive $-s$ -degree.

The idea is that any row \mathbf{p} of \mathbf{M} is such that $\mathbf{pF} = \mathbf{0} \bmod \mathbf{X}^\gamma$, and if it further satisfies $\text{cdeg}(\mathbf{p}) \leq s$ then $\text{cdeg}(\mathbf{pF}) \leq \text{cdeg}_s(\mathbf{F}) < \gamma$, so that $\mathbf{pF} = \mathbf{0}$ holds. Here the complexity mainly depends on $|s|$ and $|\text{cdeg}_s(\mathbf{F})|$, quantities that are often large in which case the algorithm of Zhou et al. [58] is more efficient. Nevertheless there are cases, such as the one arising in this section, where both sums are controlled, and elaborating over this approach leads to an efficient algorithm.

In order to propose Algorithm KNOWNDEGREEKERNELBASIS in Section 5.4, efficiently computing the kernel basis using essentially a single call to PM-BASIS, we transform the input into one with a balanced shift and a balanced order. Here and in what follows, for a nonnegative tuple $\mathbf{t} \in \mathbb{Z}_{\geq 0}^n$, we say that \mathbf{t} is *balanced* if $\max(\mathbf{t}) \in O(|\mathbf{t}|/n)$, meaning that the maximum entry in \mathbf{t} is not much larger than the average of all entries of \mathbf{t} . Section 5.2 deals with the shifts by describing a transformation of the input inspired from [46, Sec. 3], allowing us to reduce to the case of a shift s whose entries are balanced. Section 5.3 deals with balancing the order γ by performing the overlapping partial linearization of [46, Sec. 2].

For the latter transformation, as noted above, we assume that there exists a basis of the considered kernel whose column degree is bounded by s , or equivalently that $-s$ -reduced kernel bases have nonpositive $-s$ -row degree. On the other hand, for the first transformation we must ensure that $-s$ -reduced kernel bases have nonnegative $-s$ -row degree. Thus our algorithm works under the requirement that $-s$ -reduced bases of $\mathcal{K}(\mathbf{F})$ have $-s$ -row degree exactly $\mathbf{0}$, hence its name. This is a restriction compared to [56, Algo. 1] which only assumes that $-s$ -reduced bases of $\mathcal{K}(\mathbf{F})$ have nonpositive $-s$ -row degree and has to perform several approximant basis computations to

recover the whole kernel basis, as outlined in the introduction. In short, we have managed to exploit the fact that we have better a priori knowledge of degrees in kernel bases than in the context of [56], leading to a faster kernel computation which, when used in our determinant algorithm, brings no logarithmic factor in the complexity.

5.1. Normalization via kernel basis computation

We normalize the matrix \mathbf{A} into its $-s$ -Popov form \mathbf{P} using a kernel basis computation, an approach already used in a context similar to ours in the non-shifted case in [42, Lem. 19]. Roughly, this stems from the fact that the identity $\mathbf{UA} = \mathbf{P}$ with a unimodular \mathbf{U} can be rewritten as

$$\begin{bmatrix} \mathbf{U} & \mathbf{P} \end{bmatrix} \begin{bmatrix} \mathbf{A} \\ -\mathbf{I}_m \end{bmatrix} = \mathbf{0};$$

and that, for a well-chosen shift, one retrieves $[\mathbf{U} \ \mathbf{P}]$ as a shifted reduced kernel basis. The next statement gives a choice of shift suited to our situation, and describes the degree profile of such kernel bases. The focus on the shift $-\delta$ comes from the fact that any $-\delta$ -reduced form \mathbf{R} of \mathbf{A} is only a constant transformation away from being the $-s$ -Popov form \mathbf{P} of \mathbf{A} [see 26, Lem. 4.1].

Lemma 5.1. *Let $s \in \mathbb{Z}^m$, let $\mathbf{A} \in \mathbb{K}[x]^{m \times m}$ be in $-s$ -weak Popov form, let $\delta \in \mathbb{Z}_{\geq 0}^m$ be the $-s$ -pivot degree of \mathbf{A} , and assume that $s \geq \delta$. Let \mathbf{R} be a $-\delta$ -weak Popov form of \mathbf{A} and let \mathbf{U} be the unimodular matrix such that $\mathbf{UA} = \mathbf{R}$. Let further $\mathbf{d} = (s - \delta, \delta) \in \mathbb{Z}^{2m}$ and $\mathbf{F} = \begin{bmatrix} \mathbf{A} \\ -\mathbf{I}_m \end{bmatrix} \in \mathbb{K}[x]^{2m \times m}$. Then,*

- *the $-\mathbf{d}$ -pivot profile of $[\mathbf{U} \ \mathbf{R}]$ is $(m + j, \delta_j)_{1 \leq j \leq m}$,*
- *$[\mathbf{U} \ \mathbf{R}]$ is in $-\mathbf{d}$ -weak Popov form with $\text{rdeg}_{-\mathbf{d}}([\mathbf{U} \ \mathbf{R}]) = \mathbf{0}$ and $\text{cdeg}([\mathbf{U} \ \mathbf{R}]) \leq \mathbf{d}$,*
- *$[\mathbf{U} \ \mathbf{R}]$ is a basis of $\mathcal{K}(\mathbf{F})$.*

Proof. First, we prove that \mathbf{R} has $-\delta$ -pivot degree δ ; note that this implies $\text{rdeg}_{-\delta}(\mathbf{R}) = \mathbf{0}$ and $\text{cdeg}(\mathbf{R}) = \delta$ since \mathbf{R} is in $-\delta$ -weak Popov form. Lemma 2.5 shows that the $-s$ -Popov form \mathbf{P} of \mathbf{A} has the same $-s$ -pivot degree as \mathbf{A} , that is, δ . Hence, by definition of Popov forms, $\text{cdeg}(\mathbf{P}) = \delta$. Then, [26, Lem. 4.1] states that \mathbf{P} is also in $-\delta$ -Popov form. Since \mathbf{R} is a $-\delta$ -weak Popov form of \mathbf{P} , by Lemma 2.5 it has the same $-\delta$ -pivot degree as \mathbf{P} , that is, δ .

Now, by the predictable degree property and since $\text{rdeg}_{-s}(\mathbf{A}) = -s + \delta$,

$$\text{rdeg}_{-s+\delta}(\mathbf{U}) = \text{rdeg}_{-s}(\mathbf{UA}) = \text{rdeg}_{-s}(\mathbf{R}) \leq \text{rdeg}_{-\delta}(\mathbf{R}) = \mathbf{0},$$

where the inequality holds because $-s \leq -\delta$. Thus, by choice of \mathbf{d} , the $-\mathbf{d}$ -pivot entries of $[\mathbf{U} \ \mathbf{R}]$ are the $-\delta$ -pivot entries of its submatrix \mathbf{R} ; this proves the first item.

Then, the second item follows: the matrix $[\mathbf{U} \ \mathbf{R}]$ is in $-\mathbf{d}$ -weak Popov form since its $-\mathbf{d}$ -pivot index is increasing; its $-\mathbf{d}$ -row degree is equal to the $-\delta$ -row degree of \mathbf{R} which is $\mathbf{0}$; and $\text{rdeg}_{-\mathbf{d}}([\mathbf{U} \ \mathbf{R}]) = \mathbf{0}$ implies $\text{cdeg}([\mathbf{U} \ \mathbf{R}]) \leq \mathbf{d}$ by Lemma 2.1.

Let $[\mathbf{K}_1 \ \mathbf{K}_2] \in \mathbb{K}[x]^{m \times 2m}$ be a basis $\mathcal{K}(\mathbf{F})$ (it has $m = 2m - m$ rows since \mathbf{F} is $2m \times m$ and has rank m). Since $\mathbf{UA} = \mathbf{R}$, the rows of $[\mathbf{U} \ \mathbf{R}]$ are in this kernel. As a result, there exists a matrix $\mathbf{V} \in \mathbb{K}[x]^{m \times m}$ such that $[\mathbf{U} \ \mathbf{R}] = \mathbf{V}[\mathbf{K}_1 \ \mathbf{K}_2]$. In particular, $\mathbf{U} = \mathbf{V}\mathbf{K}_1$, and since \mathbf{U} is unimodular, this implies that \mathbf{V} is unimodular as well. Thus, the basis $[\mathbf{K}_1 \ \mathbf{K}_2]$ is unimodularly equivalent to $[\mathbf{U} \ \mathbf{R}]$, and the latter matrix is also a basis of $\mathcal{K}(\mathbf{F})$. \square

One may note similarities with [27, Lem. 5.1], which is about changing the shift of reduced forms via kernel basis computations. Here, we consider \mathbf{A} in $-s$ -reduced form and are interested

We use the notation in this definition in all of Section 5.2. More explicitly, $\bar{\mathbf{K}} = [\bar{\mathbf{K}}_1 \ \cdots \ \bar{\mathbf{K}}_m]$ where $\bar{\mathbf{K}}_j \in \mathbb{K}[x]^{k \times \alpha_j}$ is the unique matrix such that

$$\mathbf{K}_{*,j} = \bar{\mathbf{K}}_j \begin{bmatrix} 1 \\ x^\delta \\ \vdots \\ x^{(\alpha_j-1)\delta} \end{bmatrix}$$

and the first $\alpha_j - 1$ columns of $\bar{\mathbf{K}}_j$ have degree less than δ .

This construction originates from [46, Sec. 3], where it was designed for approximant basis computations, in order to make the shift more balanced at the cost of a small increase of the dimension; this is stated in Lemma 5.4. As mentioned above, several slightly different versions of this column partial linearization have been given in the literature: each version requires some minor adaptations of the original construction to match the context. Here, in order to benefit from properties proved in [28, Sec. 5.1], we follow the construction in [28, Lem. 5.2]: one can check that Definition 5.3 is a specialization of the construction in that reference, for the shift $-s \in \mathbb{Z}_{\leq 0}^m$ and taking the second parameter to be $t = \max(-s)$ so that $\mathbf{t} = -s - \max(-s) + t = -s$.

Lemma 5.4. *The entries of \bar{s} are in $\{0, 1, \dots, \delta\}$. Furthermore, if $\delta \geq |s|/m$, then $m \leq \bar{m} \leq 2m$.*

Proof. The first remark is obvious. For the second one, note that $1 \leq \alpha_j \leq 1 + s_j/\delta$ holds by construction, for $1 \leq j \leq m$. Hence $m \leq \bar{m} \leq m + |s|/\delta \leq 2m$. \square

Importantly, this column partial linearization behaves well with respect to shifted row degrees and shifted reduced forms.

Lemma 5.5. *Let $\mathbf{K} \in \mathbb{K}[x]^{k \times m}$ with $\text{rdeg}_{-s}(\mathbf{K}) \geq \mathbf{0}$, and let $\bar{\mathbf{K}} \in \mathbb{K}[x]^{k \times \bar{m}}$ be its column partial linearization. Then, row degrees are preserved: $\text{rdeg}_{-\bar{s}}(\bar{\mathbf{K}}) = \text{rdeg}_{-s}(\mathbf{K})$. Furthermore, if \mathbf{K} is in $-s$ -weak Popov form, then $\bar{\mathbf{K}}$ is in $-\bar{s}$ -weak Popov form.*

Proof. We show that this follows from the second item in [28, Lem. 5.2]; as noted above, the shift $\mathbf{t} = (t_1, \dots, t_m)$ in that reference corresponds to $-s = (-s_1, \dots, -s_m)$ here. Let $(\pi_i, \delta_i)_{1 \leq i \leq k}$ denote the $-s$ -pivot profile of \mathbf{K} . By definition of the $-s$ -pivot index and degree, $\text{rdeg}_{-s}(\mathbf{K}) = (\delta_i - s_{\pi_i})_{1 \leq i \leq k}$, hence our assumption $\text{rdeg}_{-s}(\mathbf{K}) \geq \mathbf{0}$ means that $\delta_i \geq s_{\pi_i}$ for $1 \leq i \leq k$. Thus, we can apply [28, Lem. 5.2] to each row of \mathbf{K} , from which we conclude that $\bar{\mathbf{K}}$ has $-\bar{s}$ -pivot profile $(\alpha_1 + \dots + \alpha_{\pi_i}, \delta_i - s_{\pi_i} + \beta_{\pi_i})_{1 \leq i \leq k}$.

First, since the entry of $-\bar{s}$ at index $\alpha_1 + \dots + \alpha_{\pi_i}$ is $-\beta_{\pi_i}$, this implies that

$$\text{rdeg}_{-\bar{s}}(\bar{\mathbf{K}}) = (\delta_i - s_{\pi_i} + \beta_{\pi_i} - \beta_{\pi_i})_{1 \leq i \leq k} = (\delta_i - s_{\pi_i})_{1 \leq i \leq k} = \text{rdeg}_{-s}(\mathbf{K}).$$

Furthermore, if \mathbf{K} is in $-s$ -weak Popov form, then $(\pi_i)_{1 \leq i \leq k}$ is increasing, hence $(\alpha_1 + \dots + \alpha_{\pi_i})_{1 \leq i \leq k}$ is increasing as well and $\bar{\mathbf{K}}$ is in $-\bar{s}$ -weak Popov form. \square

Remark 5.6. *One may further note the following properties:*

- Writing $\text{lm}_{-s}(\mathbf{K}) = [\mathbf{L}_{*,1} \ \cdots \ \mathbf{L}_{*,m}] \in \mathbb{K}^{k \times m}$, we have

$$\text{lm}_{-\bar{s}}(\bar{\mathbf{K}}) = \underbrace{[\mathbf{0} \ \cdots \ \mathbf{0} \ \mathbf{L}_{*,1}]}_{\alpha_1} \ \cdots \ \underbrace{[\mathbf{0} \ \cdots \ \mathbf{0} \ \mathbf{L}_{*,m}]}_{\alpha_m} \in \mathbb{K}^{k \times \bar{m}}.$$

- If \mathbf{K} is in $-s$ -Popov form, then $\overline{\mathbf{K}}$ is in $-\overline{s}$ -Popov form.

These properties are not used here, but for reference we provide a proof in [Appendix A](#).

We will also need properties for the converse operation, going from some matrix $\mathbf{P} \in \mathbb{K}[x]^{m \times \overline{m}}$ to its compression \mathbf{PE} .

Lemma 5.7. *Let $\mathbf{P} \in \mathbb{K}[x]^{k \times \overline{m}}$ have $-\overline{s}$ -pivot profile $(\pi_i, \delta_i)_{1 \leq i \leq k}$ and assume $\pi_i = \alpha_1 + \dots + \alpha_{j_i}$ for some $j_i \in \mathbb{Z}_{>0}$, for $1 \leq i \leq k$. Then, $\text{rdeg}_{-s}(\mathbf{PE}) = \text{rdeg}_{-\overline{s}}(\mathbf{P})$, and \mathbf{PE} has $-s$ -pivot profile*

$$(j_i, \delta_i + (\alpha_{j_i} - 1)\delta)_{1 \leq i \leq k} = (j_i, \delta_i + s_{j_i} - \beta_{j_i})_{1 \leq i \leq k}.$$

If \mathbf{P} is in $-\overline{s}$ -weak Popov form, then \mathbf{PE} is in $-s$ -weak Popov form.

Proof. The $-s$ -pivot profile of \mathbf{PE} is directly obtained by applying the first item in [\[28, Lem. 5.2\]](#) to each row of \mathbf{P} . From this $-s$ -pivot profile, we get

$$\text{rdeg}_{-s}(\mathbf{PE}) = (\delta_i + s_{j_i} - \beta_{j_i} - s_{j_i})_{1 \leq i \leq k} = (\delta_i - \beta_{j_i})_{1 \leq i \leq k}.$$

On the other hand, since the entry of $-\overline{s}$ at index $\alpha_1 + \dots + \alpha_{j_i}$ is $-\beta_{j_i}$, the $-\overline{s}$ -pivot profile of \mathbf{P} yields $\text{rdeg}_{-\overline{s}}(\mathbf{P}) = (\delta_i - \beta_{j_i})_{1 \leq i \leq k}$. Thus, $\text{rdeg}_{-s}(\mathbf{PE}) = \text{rdeg}_{-\overline{s}}(\mathbf{P})$. Now, if \mathbf{P} is in $-\overline{s}$ -weak Popov form, then $(\alpha_1 + \dots + \alpha_{j_i})_{1 \leq i \leq k}$ is increasing, which implies that $(j_i)_{1 \leq i \leq k}$ is increasing. As a result, \mathbf{PE} is in $-s$ -weak Popov form, since $(j_i)_{1 \leq i \leq k}$ is the $-s$ -pivot index of \mathbf{PE} . \square

Our approach for computing the kernel of \mathbf{F} is based on the fact that if \mathbf{K} is a basis of $\mathcal{K}(\mathbf{F})$ and $\overline{\mathbf{K}}$ is its column partial linearization, then $\mathbf{K}\mathbf{F} = \overline{\mathbf{K}}\mathbf{E}\mathbf{F} = \mathbf{0}$. This identity shows that the kernel $\mathcal{K}(\mathbf{E}\mathbf{F})$ contains the rows of $\overline{\mathbf{K}}$, so we may hope to recover $\overline{\mathbf{K}}$, and thus $\mathbf{K} = \overline{\mathbf{K}}\mathbf{E}$, from a basis of $\mathcal{K}(\mathbf{E}\mathbf{F})$; the main advantage is that the latter basis is computed with the balanced shift $-\overline{s}$. Note that a basis of $\mathcal{K}(\mathbf{E}\mathbf{F})$ does not straightforwardly yield $\overline{\mathbf{K}}$, at least because this kernel also contains $\mathcal{K}(\mathbf{E})$. In [Lemma 5.8](#) we exhibit a basis \mathbf{S} for $\mathcal{K}(\mathbf{E})$, and then in [Lemma 5.9](#) we show that $\mathcal{K}(\mathbf{E}\mathbf{F})$ is generated by $\overline{\mathbf{K}}$ and \mathbf{S} . We also give properties which allow us, from a basis of $\mathcal{K}(\mathbf{E}\mathbf{F})$, to easily recover a basis \mathbf{K} of $\mathcal{K}(\mathbf{F})$ which has the sought form (see [Lemma 5.11](#)).

Lemma 5.8. *The matrix $\mathbf{S} = \text{diag}(\mathbf{S}_1, \dots, \mathbf{S}_m) \in \mathbb{K}[x]^{(\overline{m}-m) \times \overline{m}}$, where*

$$\mathbf{S}_j = \begin{bmatrix} x^\delta & -1 & & \\ & \ddots & \ddots & \\ & & x^\delta & -1 \end{bmatrix} \in \mathbb{K}[x]^{(\alpha_{j-1}) \times \alpha_j}$$

for $1 \leq j \leq m$, is the $\mathbf{0}$ -Popov basis of the kernel $\mathcal{K}(\mathbf{E})$. Furthermore, \mathbf{S} is also in $-\overline{s}$ -Popov form, it has $-\overline{s}$ -row degree $\mathbf{0}$, and its $-\overline{s}$ -pivot profile is $(\alpha_1 + \dots + \alpha_{j-1} + i, \delta)_{1 \leq i < \alpha_j, 1 \leq j \leq m}$.

Proof. By construction, $\mathbf{S}\mathbf{E} = \mathbf{0}$ and \mathbf{S} is in $\mathbf{0}$ -Popov form. Besides, \mathbf{S} has rank $\overline{m} - m$, which is the rank of $\mathcal{K}(\mathbf{E})$ since \mathbf{E} has rank m . Now, observe that there is no nonzero vector of degree less than δ in the left kernel of the vector $[1 \ x^\delta \ \dots \ x^{(\alpha_{j-1})\delta}]^T$, and thus there is no nonzero vector of degree less than δ in $\mathcal{K}(\mathbf{E})$. It follows that \mathbf{S} is a basis of $\mathcal{K}(\mathbf{E})$. Indeed, if $\mathbf{K} \in \mathbb{K}[x]^{(\overline{m}-m) \times \overline{m}}$ is a basis of $\mathcal{K}(\mathbf{E})$ in $\mathbf{0}$ -reduced form, then $\text{rdeg}(\mathbf{K}) \geq (\delta, \dots, \delta)$. Since $\mathbf{S}\mathbf{E} = \mathbf{0}$, we have $\mathbf{S} = \mathbf{U}\mathbf{K}$ for some nonsingular $\mathbf{U} \in \mathbb{K}[x]^{(\overline{m}-m) \times (\overline{m}-m)}$. By the predictable degree property,

$$(\delta, \dots, \delta) = \text{rdeg}(\mathbf{S}) = \text{rdeg}(\mathbf{U}\mathbf{K}) = \text{rdeg}_{\text{rdeg}(\mathbf{K})}(\mathbf{U}) \geq \text{rdeg}(\mathbf{U}) + (\delta, \dots, \delta),$$

Finally, we combine the above results to show that one can compute a basis of $\mathcal{K}(\mathbf{F})$ by computing a $-\bar{s}$ -weak Popov basis of $\mathcal{K}(\mathbf{EF})$ and taking a submatrix of it.

Lemma 5.11. *Let $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$ and let $r = \bar{m} - \text{rank}(\mathbf{F})$ be the rank of $\mathcal{K}(\mathbf{EF})$. Assume that $-s$ -reduced bases of $\mathcal{K}(\mathbf{F})$ have nonnegative $-s$ -row degree. Let $\mathbf{Q} \in \mathbb{K}[x]^{r \times \bar{m}}$ be a $-\bar{s}$ -weak Popov basis of $\mathcal{K}(\mathbf{EF})$. Let $\mathbf{P} \in \mathbb{K}[x]^{k \times \bar{m}}$ be the submatrix of the rows of \mathbf{Q} whose $-\bar{s}$ -pivot index is in $\{\alpha_1 + \dots + \alpha_j, 1 \leq j \leq m\}$. Then, \mathbf{PE} is a $-s$ -weak Popov basis of $\mathcal{K}(\mathbf{F})$.*

Proof. We first prove that the number of rows of \mathbf{PE} is the rank of $\mathcal{K}(\mathbf{F})$, that is, $k = m - \text{rank}(\mathbf{F})$. Indeed, by Item (iv) of Lemma 5.9, the $-\bar{s}$ -pivot index of the $-\bar{s}$ -Popov basis of $\mathcal{K}(\mathbf{EF})$ contains the $-\bar{s}$ -pivot index of \mathbf{S} . By Lemma 5.8, the latter is the tuple formed by the integers in the set $\{1, \dots, \bar{m}\} \setminus \{\alpha_1 + \dots + \alpha_j, 1 \leq j \leq m\}$ sorted in increasing order. Since \mathbf{P} is the submatrix of the rows of \mathbf{Q} whose $-\bar{s}$ -pivot index is not in this set, \mathbf{P} has $k = r - (\bar{m} - m) = m - \text{rank}(\mathbf{F})$ rows.

Now, by construction, \mathbf{P} is in $-\bar{s}$ -weak Popov form and its $-\bar{s}$ -pivot index has entries in $\{\alpha_1 + \dots + \alpha_i, 1 \leq i \leq m\}$. Thus we can apply Lemma 5.7; it ensures that \mathbf{PE} is in $-s$ -weak Popov form and that $\text{rdeg}_{-s}(\mathbf{PE}) = \text{rdeg}_{-\bar{s}}(\mathbf{P})$.

It remains to prove that \mathbf{PE} is a basis of $\mathcal{K}(\mathbf{F})$. Let $\mathbf{K} \in \mathbb{K}^{k \times m}$ be the $-s$ -Popov basis of $\mathcal{K}(\mathbf{F})$, and let $\bar{\mathbf{K}} \in \mathbb{K}^{k \times \bar{m}}$ be its column partial linearization. Let $\mathbf{d} = \text{rdeg}_{-s}(\bar{\mathbf{K}})$, which has nonnegative entries by assumption. Then, according to Lemma 5.5, $\bar{\mathbf{K}}$ is in $-\bar{s}$ -weak Popov form, and $\text{rdeg}_{-\bar{s}}(\bar{\mathbf{K}}) = \mathbf{d}$. Then, we define $\mathbf{B} \in \mathbb{K}^{(k+\bar{m}-m) \times \bar{m}}$ as in Eq. (5); by Item (iv) of Lemma 5.9, the matrix \mathbf{B} is a $-\bar{s}$ -unordered weak Popov basis of $\mathcal{K}(\mathbf{EF})$.

Then, Lemma 2.5 shows that \mathbf{Q} has the same $-\bar{s}$ -pivot profile as the row permutation of \mathbf{B} which is in $-\bar{s}$ -weak Popov form. Since \mathbf{P} (resp. $\bar{\mathbf{K}}$) is the submatrix of the rows of \mathbf{Q} (resp. \mathbf{B}) whose $-\bar{s}$ -pivot index is in $\{\alpha_1 + \dots + \alpha_j, 1 \leq j \leq m\}$, and since both \mathbf{P} and $\bar{\mathbf{K}}$ are in $-\bar{s}$ -weak Popov form, it follows that \mathbf{P} and $\bar{\mathbf{K}}$ have the same $-\bar{s}$ -pivot profile. In particular, we have $\text{rdeg}_{-\bar{s}}(\mathbf{P}) = \text{rdeg}_{-\bar{s}}(\bar{\mathbf{K}}) = \mathbf{d}$, from which we get $\text{rdeg}_{-s}(\mathbf{PE}) = \mathbf{d}$. According to Lemma 2.3, since the rows of \mathbf{PE} are in $\mathcal{K}(\mathbf{F})$, this implies that \mathbf{PE} is a basis of $\mathcal{K}(\mathbf{F})$. \square

5.3. Reducing to the case of a balanced order

Now we apply the overlapping partial linearization of [46, Sec. 2], more precisely the version in [28, Sec. 5.2] which supports arbitrary γ as showed in the definition below that we recall for completeness. In the next lemma, we will show that the sought kernel basis can be retrieved as a submatrix of an approximant basis for the transformed problem.

Definition 5.12 ([46, 28]). *Let $\gamma = (\gamma_1, \dots, \gamma_n) \in \mathbb{Z}_{>0}^n$, let $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$ with $\text{cdeg}(\mathbf{F}) < \gamma$, and let $\mu \in \mathbb{Z}_{>0}$. Then, for $1 \leq i \leq n$, let $\gamma_i = \alpha_i \mu + \beta_i$ with $\alpha_i = \left\lceil \frac{\gamma_i}{\mu} - 1 \right\rceil$ and $1 \leq \beta_i \leq \mu$. Let also $\bar{n} = \max(\alpha_1 - 1, 0) + \dots + \max(\alpha_n - 1, 0)$, and define*

$$\mathcal{L}_\mu(\gamma) = (\bar{\gamma}_1, \dots, \bar{\gamma}_n) \in \mathbb{Z}_{>0}^{n+\bar{n}},$$

where $\bar{\gamma}_i = (2\mu, \dots, 2\mu, \mu + \beta_i) \in \mathbb{Z}_{>0}^{\alpha_i}$ if $\alpha_i > 1$ and $\bar{\gamma}_i = \gamma_i$ otherwise. Considering the i th column of \mathbf{F} , we write its x^μ -adic representation as

$$\mathbf{F}_{*,i} = \mathbf{F}_{*,i}^{(0)} + \mathbf{F}_{*,i}^{(1)} x^\mu + \dots + \mathbf{F}_{*,i}^{(\alpha_i)} x^{\alpha_i \mu}$$

where $\text{cdeg}([\mathbf{F}_{*,i}^{(0)} \ \mathbf{F}_{*,i}^{(1)} \ \dots \ \mathbf{F}_{*,i}^{(\alpha_i)}]) < (\mu, \dots, \mu, \beta_i)$.

If $\alpha_i > 1$, we define

$$\bar{\mathbf{F}}_{*,i} = \left[\mathbf{F}_{*,i}^{(0)} + \mathbf{F}_{*,i}^{(1)} x^\mu \quad \mathbf{F}_{*,i}^{(1)} + \mathbf{F}_{*,i}^{(2)} x^\mu \quad \dots \quad \mathbf{F}_{*,i}^{(\alpha_i-1)} + \mathbf{F}_{*,i}^{(\alpha_i)} x^\mu \right] \in \mathbb{K}[x]^{m \times \alpha_i}$$

and $\mathbf{J}_i = [\mathbf{0} \ \mathbf{I}_{\alpha_i-1}] \in \mathbb{K}[x]^{(\alpha_i-1) \times \alpha_i}$, and otherwise we let $\bar{\mathbf{F}}_{*,i} = \mathbf{F}_{*,i}$ and $\mathbf{J}_i \in \mathbb{K}[x]^{0 \times 1}$. Then,

$$\mathcal{L}_{\gamma,\mu}(\mathbf{F}) = \begin{bmatrix} \bar{\mathbf{F}}_{*,1} & \bar{\mathbf{F}}_{*,2} & \cdots & \bar{\mathbf{F}}_{*,n} \\ \mathbf{J}_1 & & & \\ & \mathbf{J}_2 & & \\ & & \ddots & \\ & & & \mathbf{J}_n \end{bmatrix} \in \mathbb{K}[x]^{(m+\bar{n}) \times (n+\bar{n})}$$

is called the overlapping linearization of \mathbf{F} with respect to γ and μ .

Lemma 5.13. *Let $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$, and let $s \in \mathbb{Z}_{\geq 0}^m$ be such that there exists a basis of $\mathcal{K}(\mathbf{F})$ with nonpositive $-s$ -row degree. Let $\mu \in \mathbb{Z}_{>0}$ with $\mu > \max(s)$, $\mathbf{t} = (s, \mu - 1, \dots, \mu - 1) \in \mathbb{Z}_{\geq 0}^{m+\bar{n}}$, and $\gamma \in \mathbb{Z}_{>0}^n$ with $\gamma \geq \text{cdeg}_s(\mathbf{F}) + 1$. Let $\mathbf{M} \in \mathbb{K}[x]^{(m+\bar{n}) \times (n+\bar{n})}$ be a $-\mathbf{t}$ -reduced basis of $\mathcal{A}_{\mathcal{L}_{\mu}(\gamma)}(\mathcal{L}_{\gamma,\mu}(\mathbf{F}))$. Then, exactly $k = m - \text{rank}(\mathbf{F})$ rows of \mathbf{M} have nonpositive $-\mathbf{t}$ -degree, and the first m columns of these rows form a matrix $\mathbf{K} \in \mathbb{K}[x]^{k \times m}$ which is a $-s$ -reduced basis of $\mathcal{K}(\mathbf{F})$. Besides, if \mathbf{M} is in $-\mathbf{t}$ -weak Popov form, then \mathbf{K} is in $-s$ -weak Popov form.*

Proof. Let $[\mathbf{K} \ \mathbf{Q}]$ be the submatrix of \mathbf{M} formed by its rows of nonpositive $-\mathbf{t}$ -degree, where $\mathbf{K} \in \mathbb{K}[x]^{k \times m}$ and $\mathbf{Q} \in \mathbb{K}[x]^{k \times \bar{n}}$; we have $0 \leq k \leq m + \bar{n}$. By choice of \mathbf{t} , from $\text{rdeg}_{-\mathbf{t}}([\mathbf{K} \ \mathbf{Q}]) \leq \mathbf{0}$ we get $\text{deg}(\mathbf{Q}) < \mu$ and $\text{rdeg}_{-s}(\mathbf{K}) \leq \mathbf{0}$.

In particular, $\text{deg}(\mathbf{K}) \leq \max(s) < \mu$: applying the second item in [28, Lem. 5.6] to each row of $[\mathbf{K} \ \mathbf{Q}]$ shows that $\text{rdeg}(\mathbf{Q}) < \text{rdeg}(\mathbf{K})$ and that the rows of \mathbf{K} are in $\mathcal{A}_{\gamma}(\mathbf{F})$, that is, $\mathbf{K}\mathbf{F} = \mathbf{0} \pmod{\mathbf{X}^{\gamma}}$. On the other hand, $\text{cdeg}(\mathbf{K}) \leq s$ implies that $\text{cdeg}(\mathbf{K}\mathbf{F}) \leq \text{cdeg}_s(\mathbf{F}) < \gamma$, hence $\mathbf{K}\mathbf{F} = \mathbf{0}$, i.e. the rows of \mathbf{K} are in $\mathcal{K}(\mathbf{F})$. This implies that the rank of \mathbf{K} is at most the rank of the module $\mathcal{K}(\mathbf{F})$, i.e. $\text{rank}(\mathbf{K}) \leq m - r$ where $r = \text{rank}(\mathbf{F})$.

Furthermore, from $\text{rdeg}(\mathbf{Q}) < \text{rdeg}(\mathbf{K})$ and $\max(s) < \mu$ we obtain

$$\text{rdeg}_{-s}(\mathbf{K}) \geq \text{rdeg}(\mathbf{K}) - \max(s) > \text{rdeg}(\mathbf{Q}) - \mu + 1 = \text{rdeg}_{(-\mu+1, \dots, -\mu+1)}(\mathbf{Q}),$$

hence by choice of \mathbf{t} we have $\text{lm}_{-\mathbf{t}}([\mathbf{K} \ \mathbf{Q}]) = [\text{lm}_{-s}(\mathbf{K}) \ \mathbf{0}]$. Since this matrix is a subset of the rows of the nonsingular matrix $\text{lm}_{-\mathbf{t}}(\mathbf{M})$, it has full row rank, and thus $\text{lm}_{-s}(\mathbf{K})$ has full row rank. This shows that \mathbf{K} is in $-s$ -reduced form, and that $k = \text{rank}(\mathbf{K})$. If \mathbf{M} is furthermore in $-\mathbf{t}$ -weak Popov form, then $[\mathbf{K} \ \mathbf{Q}]$ is in $-\mathbf{t}$ -weak Popov form as well; then, the identity $\text{lm}_{-\mathbf{t}}([\mathbf{K} \ \mathbf{Q}]) = [\text{lm}_{-s}(\mathbf{K}) \ \mathbf{0}]$ shows that the $-\mathbf{t}$ -pivot entries of $[\mathbf{K} \ \mathbf{Q}]$ are all located in \mathbf{K} , hence \mathbf{K} is in $-s$ -weak Popov form.

It remains to prove that $k = m - r$ and that the rows of \mathbf{K} generate $\mathcal{K}(\mathbf{F})$.

By assumption, there exists a basis $\mathbf{P} \in \mathbb{K}[x]^{(m-r) \times m}$ of $\mathcal{K}(\mathbf{F})$ such that $\text{cdeg}(\mathbf{P}) \leq s$. In particular, the rows of \mathbf{P} are in $\mathcal{A}_{\gamma}(\mathbf{F})$, and applying the first item of [28, Lem. 5.6] to each of these rows shows that there exists a matrix $\mathbf{R} \in \mathbb{K}[x]^{(m-r) \times \bar{n}}$ with $\text{rdeg}(\mathbf{R}) < \text{rdeg}(\mathbf{P})$ and such that the rows of $[\mathbf{P} \ \mathbf{R}]$ are in $\mathcal{A}_{\mathcal{L}_{\mu}(\gamma)}(\mathcal{L}_{\gamma,\mu}(\mathbf{F}))$. Thus, $[\mathbf{P} \ \mathbf{R}]$ is a left multiple of \mathbf{M} .

A key remark now is that $\text{rdeg}_{-\mathbf{t}}([\mathbf{P} \ \mathbf{R}]) \leq \mathbf{0}$. Indeed, by Lemma 2.1 $\text{rdeg}_{-s}(\mathbf{P}) \leq \mathbf{0}$ follows from $\text{cdeg}(\mathbf{P}) \leq s$, and we have

$$\text{deg}(\mathbf{R}) < \text{deg}(\mathbf{P}) = \max(\text{cdeg}(\mathbf{P})) \leq \max(s) < \mu.$$

Thus, since \mathbf{M} is $-\mathbf{t}$ -reduced, the predictable degree property ensures that $[\mathbf{P} \ \mathbf{R}]$ is a left multiple of \mathbf{M} which does not involve the rows of \mathbf{M} of positive $-\mathbf{t}$ -degree, i.e. a left multiple of $[\mathbf{K} \ \mathbf{Q}]$. In particular, \mathbf{P} is a left multiple of \mathbf{K} : there exists a matrix $\mathbf{V} \in \mathbb{K}[x]^{(m-r) \times k}$ such that $\mathbf{V}\mathbf{K} = \mathbf{P}$.

Since \mathbf{P} has rank $m - r$, we obtain $k \geq m - r$, hence $k = m - r$. On the other hand, since the rows of \mathbf{K} are in $\mathcal{K}(\mathbf{F})$, there exists a matrix $\mathbf{W} \in \mathbb{K}[x]^{k \times k}$ such that $\mathbf{W}\mathbf{P} = \mathbf{K}$. It follows that $\mathbf{P} = \mathbf{V}\mathbf{W}\mathbf{P}$, and since \mathbf{P} has full row rank this implies $\mathbf{V}\mathbf{W} = \mathbf{I}_k$. This means that \mathbf{P} and \mathbf{K} are left unimodularly equivalent, hence \mathbf{K} is a basis of $\mathcal{K}(\mathbf{F})$, which concludes the proof. \square

5.4. Computing kernel bases with known pivot degree

After applying the transformations presented in Sections 5.2 and 5.3, we are left with the computation of an approximant basis for a balanced order $\mathcal{L}_\mu(\gamma)$ and a balanced shift $-\bar{s}$: this is done efficiently by PM-Basis, designed in [21] as an improvement of [3, Algo. SPHPS]. Here, we use the version in [28, Algo. 2] which ensures that the output basis is in $-\bar{s}$ -weak Popov form.

Algorithm 4 KNOWNDEGREEKERNELBASIS(\mathbf{F}, s)

Input: a matrix $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$, and a nonnegative shift $s \in \mathbb{Z}_{\geq 0}^m$.

Requirement: $-s$ -reduced bases of $\mathcal{K}(\mathbf{F})$ have $-s$ -row degree $\mathbf{0}$.

Output: a $-s$ -weak Popov basis of $\mathcal{K}(\mathbf{F})$.

- 1: \triangleright Step 1: Output column partial linearization \triangleleft
 - 2: $\delta \leftarrow \lceil D/m \rceil \in \mathbb{Z}_{>0}$, where $D = \max(|s|, |\text{cdeg}_s(\mathbf{F})|, 1)$
 - 3: Apply Definition 5.3 to (s, δ) to obtain the parameters and expansion-compression matrix:
 $(\alpha_1, \dots, \alpha_m) \in \mathbb{Z}_{>0}^m, \bar{m} \in \mathbb{Z}_{>0}, \bar{s} \in \mathbb{Z}_{\geq 0}^m, \mathbf{E} \in \mathbb{K}[x]^{\bar{m} \times m}$
 - 4: \triangleright Step 2: Overlapping partial linearization \triangleleft
 - 5: $\gamma \leftarrow \text{cdeg}_s(\mathbf{F}) + 1 \in \mathbb{Z}_{>0}^n$ \triangleright order for approximation, equal to $\text{cdeg}_{\bar{s}}(\mathbf{E}\mathbf{F}) + 1$
 - 6: Apply Definition 5.12 to $(\gamma, \mathbf{E}\mathbf{F}, \delta + 1)$ to obtain the order $\mathcal{L}_{\delta+1}(\gamma) \in \mathbb{Z}_{>0}^{n+\bar{n}}$ and the matrix
 $\mathcal{L}_{\gamma, \delta+1}(\mathbf{E}\mathbf{F}) \in \mathbb{K}[x]^{(\bar{m}+\bar{n}) \times (n+\bar{n})}$
 - 7: \triangleright Step 3: Compute $-t$ -weak Popov basis of $\mathcal{A}_{\mathcal{L}_{\delta+1}(\gamma)}(\mathcal{L}_{\gamma, \delta+1}(\mathbf{E}\mathbf{F}))$ \triangleleft
 - 8: $\mathbf{t} \leftarrow (\bar{s}, \delta, \dots, \delta) \in \mathbb{Z}_{>0}^{\bar{m}+\bar{n}}$
 - 9: $\Gamma \leftarrow \max(\mathcal{L}_{\delta+1}(\gamma)); \mathbf{G} = \mathcal{L}_{\gamma, \delta+1}(\mathbf{E}\mathbf{F})\mathbf{X}^{(\Gamma, \dots, \Gamma) - \mathcal{L}_{\delta+1}(\gamma)}$ \triangleright use uniform order (Γ, \dots, Γ)
 - 10: $\mathbf{M} \in \mathbb{K}[x]^{(\bar{m}+\bar{n}) \times (\bar{m}+\bar{n})} \leftarrow \text{PM-BASIS}(\Gamma, \mathbf{G}, -\mathbf{t})$
 - 11: \triangleright Step 4: Deduce first $-\bar{s}$ -weak Popov basis of $\mathcal{K}(\mathbf{E}\mathbf{F})$, then $-s$ -weak Popov basis of $\mathcal{K}(\mathbf{F})$ \triangleleft
 - 12: $\mathbf{Q} \in \mathbb{K}[x]^{\bar{k} \times \bar{m}} \leftarrow$ first \bar{m} columns of the rows of \mathbf{M} which have nonpositive $-\mathbf{t}$ -degree
 - 13: $\mathbf{P} \in \mathbb{K}[x]^{k \times \bar{m}} \leftarrow$ the rows of \mathbf{Q} whose $-\bar{s}$ -pivot index is in $\{\alpha_1 + \dots + \alpha_j, 1 \leq j \leq m\}$
 - 14: **return** $\mathbf{P}\mathbf{E}$
-

Proposition 5.14. *Algorithm 4 is correct. Let $D = \max(|s|, |\text{cdeg}_s(\mathbf{F})|, 1)$. Then, assuming $m \geq n$ and using notation from the algorithm, its cost is bounded by the sum of:*

- the cost of performing PM-BASIS at order at most $2\lceil D/m \rceil + 2$ on an input matrix of row dimension $\bar{m} + \bar{n} \leq 3m$ and column dimension $n + \bar{n} \leq 2m$;
- $O(m^2)$ extra operations in \mathbb{K} .

Thus, Algorithm 4 uses $O(m^\omega M'(D/m))$ operations in \mathbb{K} .

Proof. Using the assumption that $-s$ -reduced bases of $\mathcal{K}(\mathbf{F})$ have $-s$ -row degree $\mathbf{0}$ along with Item (iv) of Lemma 5.9 shows that the $-\bar{s}$ -reduced bases of $\mathcal{K}(\mathbf{E}\mathbf{F})$ have $-\bar{s}$ -row degree $\mathbf{0}$. Thus, we can apply Lemma 5.13 to $(\mathbf{E}\mathbf{F}, \bar{s}, \mu, \gamma)$ with $\mu = \delta + 1 > \max(\bar{s})$ and $\gamma = \text{cdeg}_{\bar{s}}(\mathbf{E}\mathbf{F}) + 1$, which is $\gamma = \text{cdeg}_s(\mathbf{F}) + 1$ according to Item (i) of Lemma 5.9. Note that \mathbf{M} is a $-t$ -weak Popov basis of $\mathcal{A}_{(\Gamma, \dots, \Gamma)}(\mathbf{G}) = \mathcal{A}_{\mathcal{L}_{\delta+1}(\gamma)}(\mathcal{L}_{\gamma, \delta+1}(\mathbf{E}\mathbf{F}))$ (see e.g. [28, Rmk. 3.3] for this approach to make the order uniform). Then, Lemma 5.13 states that the matrix \mathbf{Q} at Line 12 has $\bar{k} = \bar{m} - \text{rank}(\mathbf{E}\mathbf{F}) =$

$\bar{m} - \text{rank}(\mathbf{F})$ rows and is a $-\bar{s}$ -weak Popov basis of $\mathcal{K}(\mathbf{EF})$. Then, by Lemma 5.11, \mathbf{PE} is a $-\bar{s}$ -weak Popov basis of $\mathcal{K}(\mathbf{F})$, hence the correctness.

For the cost analysis, we assume $m \geq n$, and we start by summarizing bounds on the dimensions and degrees at play. Lemma 5.4 yields $m \leq \bar{m} \leq 2m$, while Item (i) of Lemma 5.9 ensures $\text{cdeg}_{\bar{s}}(\mathbf{EF}) = \text{cdeg}_{\bar{s}}(\mathbf{F})$. Each entry of $\mathcal{L}_{\delta+1}(\boldsymbol{\gamma})$ is at most $2(\delta+1) = 2\lceil D/m \rceil + 2$, by construction. Writing $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_n)$, by Definition 5.12 we have

$$\bar{n} = \max\left(\left\lceil \frac{\gamma_1}{\delta+1} - 1 \right\rceil, 0\right) + \dots + \max\left(\left\lceil \frac{\gamma_n}{\delta+1} - 1 \right\rceil, 0\right) \leq \frac{\gamma_1}{\delta+1} + \dots + \frac{\gamma_n}{\delta+1} = \frac{|\boldsymbol{\gamma}|}{\delta+1}.$$

Since $|\boldsymbol{\gamma}| = |\text{cdeg}_{\bar{s}}(\mathbf{F})| + n \leq D + m$, and since $\delta + 1 \geq (D + m)/m$, it follows that $\bar{n} \leq m$. Besides, $\text{rdeg}_{-\bar{s}}(\mathbf{P}) \leq \mathbf{0}$ by construction, so that $\text{cdeg}(\mathbf{P}) \leq \bar{s}$ by Lemma 2.1.

The only steps that involve operations in \mathbb{K} are the call to PM-BASIS at Line 10 and the multiplication \mathbf{PE} at Line 14. The construction of \mathbf{E} and the inequality $\text{cdeg}(\mathbf{P}) \leq \bar{s}$ imply that the product \mathbf{PE} mainly involves concatenating vectors of coefficients; concerning operations in \mathbb{K} , there are $\bar{m} - m$ columns of \mathbf{P} for which we may add the constant term of that column to the term of degree δ of the previous column. Therefore Line 14 has cost bound $O(\bar{m}k)$; since $\bar{m} \leq 2m$ and $k \leq m$ this is in $O(m^2)$. At Line 10, we call the approximant basis subroutine PM-BASIS discussed in Section 2.5; since $\Gamma = \max(\mathcal{L}_{\delta+1}(\boldsymbol{\gamma}))$ is at most $2\lceil D/m \rceil + 2 \in O(1 + D/m)$, Lemma 2.9 states that this call uses

$$O\left((\bar{m} + \bar{n} + n + \bar{n})(\bar{m} + \bar{n})^{\omega-1} M'(D/m)\right)$$

operations in \mathbb{K} . Since $n + \bar{n} \leq \bar{m} + \bar{n} \leq 3m$, this yields the claimed cost bound. \square

5.5. Proof of Theorem 1.4

Algorithm 5 WEAKPOPOVTOPOPOV(\mathbf{A}, s)

Input: a matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times n}$, a shift $s \in \mathbb{Z}^n$ such that \mathbf{A} is in $-s$ -weak Popov form.

Requirement: $s_\pi \geq \delta$, where $(\boldsymbol{\pi}, \delta)$ is the $-s$ -pivot profile of \mathbf{A} .

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

1: \triangleright Step 1: Find unimodular transformation and $-\delta$ -reduced form of $\mathbf{A}_{*,\boldsymbol{\pi}}$ \blacktriangleleft

2: $(\boldsymbol{\pi}, \delta) \leftarrow$ the $-s$ -pivot profile of \mathbf{A}

3: $[\mathbf{U} \ \mathbf{R}] \in \mathbb{K}[x]^{m \times 2m} \leftarrow \text{KNOWNDEGREEKERNELBASIS}\left(\begin{bmatrix} \mathbf{A}_{*,\boldsymbol{\pi}} \\ -\mathbf{I}_m \end{bmatrix}, (s_\pi - \delta, \delta)\right)$

4: \triangleright Step 2: Deduce $-s$ -Popov form of \mathbf{A} \blacktriangleleft

5: $\mathbf{P} \leftarrow$ zero matrix in $\mathbb{K}[x]^{m \times n}$

6: $\mathbf{P}_{*,\boldsymbol{\pi}} \leftarrow \text{Im}_{-\delta}(\mathbf{R})^{-1} \mathbf{R}$

7: $\mathbf{P}_{*,\{1,\dots,n\} \setminus \boldsymbol{\pi}} \leftarrow \text{Im}_{-\delta}(\mathbf{R})^{-1} \mathbf{U} \mathbf{A}_{*,\{1,\dots,n\} \setminus \boldsymbol{\pi}}$

8: **return** \mathbf{P}

For proving Theorem 1.4, we describe Algorithm WEAKPOPOVTOPOPOV (Algorithm 5) and we focus on the square case, $m = n$. Then, by definition of the $-s$ -weak Popov form, δ is the tuple of degrees of the diagonal entries of \mathbf{A} , and $\boldsymbol{\pi} = (1, \dots, m)$. Furthermore, in this case, $s_\pi = s$, $\mathbf{A}_{*,\boldsymbol{\pi}} = \mathbf{A}$, and $\mathbf{P}_{*,\boldsymbol{\pi}} = \mathbf{P}$; in particular, we can discard the step at Line 7 since the submatrices it involves are empty.

First note that the shift $\mathbf{d} = (s_\pi - \delta, \delta) = (s - \delta, \delta)$ used at Line 3 is nonnegative. Besides, Lemma 5.1 shows that $-\mathbf{d}$ -reduced bases of $\mathcal{K}(\mathbf{F})$ have $-\mathbf{d}$ -row degree $\mathbf{0}$. Thus the requirements

of Algorithm `KNOWNDEGREEKERNELBASIS` are met, and Proposition 5.14 shows that the matrix $[\mathbf{U} \ \mathbf{R}]$ computed at Line 3 is a $-d$ -weak Popov basis of $\mathcal{K}(\mathbf{F})$. Then, Corollary 5.2 shows that the matrix $\mathbf{P} = \mathbf{P}_{*,\pi} = \text{Im}_{-\delta}(\mathbf{R})^{-1}\mathbf{R}$ computed at Line 6 is the $-s$ -Popov form of \mathbf{A} . This proves that Algorithm 5 is correct.

Concerning the cost bound, we focus on the case $|s| > 0$. Indeed, since $s \geq \delta \geq \mathbf{0}$, if $|s| = 0$, then $s = \delta = \mathbf{0}$. In this case, no computation needs to be done: the $-s$ -Popov form of \mathbf{A} is \mathbf{I}_m , the unique matrix in Popov form whose pivot degree is $\mathbf{0}$.

The cost of Line 6 is that of multiplying $\text{Im}_{-\delta}(\mathbf{R})^{-1} \in \mathbb{K}^{m \times m}$ by $\mathbf{R} \in \mathbb{K}[x]^{m \times m}$, which has column degree δ as explained in Section 5.1; this computation corresponds to the second item in Theorem 1.4. This multiplication can be done by first performing a column linearization of \mathbf{R} into a $m \times (m + |\delta|)$ matrix $\overline{\mathbf{R}}$ over \mathbb{K} , computing $\text{Im}_{-\delta}(\mathbf{R})^{-1}\overline{\mathbf{R}}$, and finally compressing the result back into a polynomial matrix. This uses $O(m^\omega(1 + |\delta|/m))$ operations in \mathbb{K} .

Concerning Line 3, we rely on Proposition 5.14. Here, the matrix we give as input to Algorithm `KNOWNDEGREEKERNELBASIS` has dimensions $2m \times m$, hence the dimensions in Proposition 5.14 satisfy $\overline{m} \leq 4m$ and $\overline{n} \leq 2m$ (the latter bound comes from the proof of that proposition). Then, Proposition 5.14 states that Line 3 costs:

- $O(m^2)$ operations in \mathbb{K} , which is the third item in Theorem 1.4,
- one call to `PM-BASIS` on a matrix of row dimension $\overline{m} + \overline{n} \leq 6m$, column dimension $m + \overline{n} \leq 3m$, and at order at most $2\lceil D/(2m) \rceil + 2$, where $D = \max(|d|, |\text{cdeg}_d(\mathbf{F})|, 1)$ and \mathbf{F} is the input matrix $[\mathbf{A}^\top \ -\mathbf{I}_m]^\top$.

Besides, Proposition 5.14 also implies that Line 3 uses $O(m^\omega M'(D/m))$ operations in \mathbb{K} .

We are going to prove that $D = |s|$, which concludes the proof. Indeed, the previous paragraph then directly gives the overall cost bound $O(m^\omega M'(|s|/m))$ in Theorem 1.4, and using

$$2 \left\lceil \frac{D}{2m} \right\rceil + 2 < 2 \left(\frac{D}{2m} + 1 \right) + 2 = |s|/m + 4,$$

the previous paragraph also gives the first item in that theorem.

To observe that $D = |s|$, we first use the definition of d to obtain $|d| = |s - \delta| + |\delta| = |s|$. Since $|s| \geq 1$, this gives $D = \max(|s|, |\text{cdeg}_d(\mathbf{F})|)$. Now, $\text{cdeg}_d(\mathbf{F})$ is the entry-wise maximum of $\text{cdeg}_{s-\delta}(\mathbf{A})$ and $\text{cdeg}_\delta(-\mathbf{I}_m) = \delta$. Since \mathbf{A} is in $-s$ -row reduced form with $\text{rdeg}_{-s}(\mathbf{A}) = -s + \delta$, Lemma 3.2 ensures that \mathbf{A}^\top is in $s - \delta$ -reduced form with $\text{cdeg}_{s-\delta}(\mathbf{A}) = \text{rdeg}_{s-\delta}(\mathbf{A}^\top) = s$. Then, since $s \geq \delta$ we obtain $\text{cdeg}_d(\mathbf{F}) = s$, and thus $D = |s|$. This concludes the proof of Theorem 1.4.

As for the rectangular case $m < n$, the correctness of Algorithm 5 follows from the above proof in the square case, which shows that the algorithm correctly computes the $-s_\pi$ -Popov form $\mathbf{P}_{*,\pi}$ of $\mathbf{A}_{*,\pi}$, and from [37, Lem. 5.1] concerning the computation of $\mathbf{P}_{*,\{1,\dots,n\}\setminus\pi}$ and the fact that $\text{Im}_{-\delta}(\mathbf{R})^{-1}\mathbf{U}$ is the unimodular matrix which transforms \mathbf{A} into \mathbf{P} . The cost bound can be derived using the degree bounds on rectangular shifted Popov forms given in [5, 37] (we do not detail this here since this would add technical material unrelated to the main results of this article).

Appendix A.

In this appendix, we give proofs for Remarks 5.6 and 5.10. We use notation from Section 5.2: a shift $s \in \mathbb{Z}_{\geq 0}^m$, an integer $\delta \in \mathbb{Z}_{>0}$, and a matrix $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$ are given; $\mathbf{K} \in \mathbb{K}[x]^{k \times m}$ is a basis of $\mathcal{K}(\mathbf{F})$; $(\alpha_j, \beta_j)_{1 \leq j \leq m}$ and $\overline{s} \in \mathbb{Z}_{\geq 0}^m$ and $\overline{\mathbf{K}} \in \mathbb{K}[x]^{k \times \overline{m}}$ are as described in Definition 5.3. Following the context of Remarks 5.6 and 5.10, we assume $\text{rdeg}_{-s}(\mathbf{K}) \geq \mathbf{0}$, hence in particular $\text{rdeg}_{-\overline{s}}(\overline{\mathbf{K}}) = \text{rdeg}_{-s}(\mathbf{K})$ (see Lemma 5.5). We start with the claims in Remark 5.6.

Lemma A.1. Writing $\text{lm}_{-s}(\mathbf{K}) = [\mathbf{L}_{*,1} \ \cdots \ \mathbf{L}_{*,m}] \in \mathbb{K}^{k \times m}$, we have

$$\text{lm}_{-\bar{s}}(\bar{\mathbf{K}}) = [\underbrace{\mathbf{0} \ \cdots \ \mathbf{0} \ \mathbf{L}_{*,1}}_{\alpha_1} \ \cdots \ \underbrace{\mathbf{0} \ \cdots \ \mathbf{0} \ \mathbf{L}_{*,m}}_{\alpha_m}] \in \mathbb{K}^{k \times \bar{m}}.$$

If \mathbf{K} is in $-s$ -Popov form, then $\bar{\mathbf{K}}$ is in $-\bar{s}$ -Popov form.

Proof. For given $i \in \{1, \dots, k\}$ and $j \in \{1, \dots, m\}$, we rely on the definition of a leading matrix (see Section 2.3): the entry (i, j) of $\text{lm}_{-\bar{s}}(\bar{\mathbf{K}})$ is the coefficient of degree $d_i + \bar{s}_j$ of $\bar{\mathbf{K}}_{i,j}$, where $d_i = \text{rdeg}_{-\bar{s}}(\bar{\mathbf{K}}_{i,*}) = \text{rdeg}_{-s}(\mathbf{K}_{i,*}) \geq 0$. First consider the case $j \notin \{\alpha_1 + \dots + \alpha_\pi, 1 \leq \pi \leq m\}$. Then the j th entry of \bar{s} is $\bar{s}_j = \delta$, and by definition of the output column partial linearization $\bar{\mathbf{K}}_{i,j}$ has degree less than δ , hence its coefficient of degree $d_i + \delta \geq \delta$ must be zero. This proves that all columns of $\text{lm}_{-\bar{s}}(\bar{\mathbf{K}})$ with index not in $\{\alpha_1 + \dots + \alpha_\pi, 1 \leq \pi \leq m\}$ are indeed zero. It remains to prove that in the case $j = \alpha_1 + \dots + \alpha_\pi$ for some $1 \leq \pi \leq m$, then the entry (i, j) of $\text{lm}_{-\bar{s}}(\bar{\mathbf{K}})$ is equal to $\mathbf{L}_{i,\pi}$. This holds, since in this case we have $\bar{s}_j = \beta_\pi$, and by construction of $\bar{\mathbf{K}}$ the coefficient of degree $d_i + \beta_\pi$ of $\bar{\mathbf{K}}_{i,j}$ is equal to the coefficient of degree $d_i + (\alpha_\pi - 1)\delta + \beta_\pi = d_i + s_\pi$ of $\mathbf{K}_{i,\pi}$, which itself is equal to $\mathbf{L}_{i,\pi}$ by definition of a leading matrix.

Now assume \mathbf{K} is in $-s$ -Popov form. We have showed in Lemma 5.5 that $\bar{\mathbf{K}}$ is in $-\bar{s}$ -weak Popov form and that the row $\bar{\mathbf{K}}_{i,*}$ has its $-\bar{s}$ -pivot at index $\alpha_1 + \dots + \alpha_{\pi_i}$ where π_i is the $-s$ -pivot index of $\mathbf{K}_{i,*}$. By construction, the column $\bar{\mathbf{K}}_{*,\alpha_1 + \dots + \alpha_{\pi_i}}$ is the part of nonnegative degree of the column $x^{-(\alpha_{\pi_i} - 1)\delta} \mathbf{K}_{*,\pi_i} = x^{-s_{\pi_i} + \beta_i} \mathbf{K}_{*,\pi_i}$. It follows first that the $-\bar{s}$ -pivot entry $\bar{\mathbf{K}}_{i,\alpha_1 + \dots + \alpha_{\pi_i}}$ is monic since it is a high degree part of the (monic) $-s$ -pivot entry \mathbf{K}_{i,π_i} ; and second that $\bar{\mathbf{K}}_{i,\alpha_1 + \dots + \alpha_{\pi_i}}$ has degree strictly larger than all other entries in the column $\bar{\mathbf{K}}_{*,\alpha_1 + \dots + \alpha_{\pi_i}}$ since \mathbf{K}_{i,π_i} has degree strictly larger than all other entries in the column \mathbf{K}_{*,π_i} . Hence $\bar{\mathbf{K}}$ is in $-\bar{s}$ -Popov form. \square

Now we prove the claims in Remark 5.10 concerning variants of Item (iv) of Lemma 5.9, using further notation from Section 5.2: \mathbf{S} is the basis of $\mathcal{K}(\mathbf{E})$ described in Lemma 5.8 and $\mathbf{B} = [\frac{\bar{\mathbf{K}}}{\mathbf{S}}] \in \mathbb{K}[x]^{(k+\bar{m}-m) \times \bar{m}}$ is the basis of $\mathcal{K}(\mathbf{E}\mathbf{F})$ given in Eq. (5). As above we assume $\text{rdeg}_{-s}(\mathbf{K}) \geq \mathbf{0}$, and therefore $\text{rdeg}_{-\bar{s}}(\mathbf{B}) = (\text{rdeg}_{-s}(\mathbf{K}), \mathbf{0})$ as stated in Item (iv) of Lemma 5.9.

Lemma A.2. With the above notation and assumptions,

- If \mathbf{K} is in $-s$ -reduced form, then \mathbf{B} is in $-\bar{s}$ -reduced form.
- If \mathbf{K} is in $-s$ -Popov form, then \mathbf{B} is in $-\bar{s}$ -Popov form up to row permutation.

Proof. Suppose first that \mathbf{K} is in $-s$ -reduced form. We apply Lemma A.1:

$$\text{lm}_{-\bar{s}}(\bar{\mathbf{K}}) = [\mathbf{0} \ \cdots \ \mathbf{0} \ \mathbf{L}_{*,1} \ \cdots \ \mathbf{0} \ \cdots \ \mathbf{0} \ \mathbf{L}_{*,m}]$$

where $\text{lm}_{-s}(\mathbf{K}) = [\mathbf{L}_{*,1} \ \cdots \ \mathbf{L}_{*,m}] \in \mathbb{K}^{k \times m}$. The latter matrix has full rank, since \mathbf{K} is in $-s$ -reduced form. On the other hand, for matrices such as the diagonal blocks of \mathbf{S} we have

$$\text{lm}_{(-\delta, \dots, -\delta, -\beta)} \left(\begin{bmatrix} x^\delta & -1 & & \\ & \ddots & \ddots & \\ & & x^\delta & -1 \end{bmatrix} \right) = \begin{bmatrix} 1 & 0 & & \\ & \ddots & \ddots & \\ & & & 1 & 0 \end{bmatrix}$$

- [13] Dumas, J.G., Pernet, C., Sultan, Z., 2017. Fast computation of the rank profile matrix and the generalized Bruhat decomposition. *J. Symbolic Comput.* 83, 187–210. doi:[10.1016/j.jsc.2016.11.011](https://doi.org/10.1016/j.jsc.2016.11.011).
- [14] Dumas, J.G., Pernet, C., Wan, Z., 2005. Efficient computation of the characteristic polynomial, in: ISSAC'05, ACM, pp. 140–147. doi:[10.1145/1073884.1073905](https://doi.org/10.1145/1073884.1073905).
- [15] Dummit, D.S., Foote, R.M., 2004. *Abstract Algebra*. John Wiley & Sons.
- [16] Faddeev, D., Sominskii, I., 1949. Collected Problems in Higher Algebra, Problem n°979.
- [17] Forney, Jr., G.D., 1975. Minimal Bases of Rational Vector Spaces, with Applications to Multivariable Linear Systems. *SIAM Journal on Control* 13, 493–520. doi:[10.1137/0313029](https://doi.org/10.1137/0313029).
- [18] Frame, J., 1949. A simple recurrent formula for inverting a matrix (abstract). *Bull. of Amer. Math. Soc.* 55, 1045.
- [19] Gathen, J.v.z., Gerhard, J., 2013. *Modern Computer Algebra* (third edition). Cambridge University Press. doi:[10.1017/CB09781139856065](https://doi.org/10.1017/CB09781139856065).
- [20] Giesbrecht, M., 1995. Nearly optimal algorithms for canonical matrix forms. *SIAM Journal on Computing* 24, 948–969. doi:[10.1137/S0097539793252687](https://doi.org/10.1137/S0097539793252687).
- [21] Giorgi, P., Jeannerod, C.P., Villard, G., 2003. On the complexity of polynomial matrix computations, in: ISSAC'03, ACM, pp. 135–142. doi:[10.1145/860854.860889](https://doi.org/10.1145/860854.860889).
- [22] Giorgi, P., Neiger, V., 2018. Certification of minimal approximant bases, in: ISSAC'18, ACM, pp. 167–174. doi:[10.1145/3208976.3208991](https://doi.org/10.1145/3208976.3208991).
- [23] Gupta, S., Sarkar, S., Storjohann, A., Valeriotte, J., 2012. Triangular x -basis decompositions and derandomization of linear algebra algorithms over $K[x]$. *J. Symbolic Comput.* 47, 422–453. doi:[10.1016/j.jsc.2011.09.006](https://doi.org/10.1016/j.jsc.2011.09.006).
- [24] Harvey, D., Van Der Hoeven, J., Lecerf, G., 2017. Faster polynomial multiplication over finite fields. *J. ACM* 63. doi:[10.1145/3005344](https://doi.org/10.1145/3005344).
- [25] Ibarra, O.H., Moran, S., Hui, R., 1982. A generalization of the fast LUP matrix decomposition algorithm and applications. *Journal of Algorithms* 3, 45–56. doi:[10.1016/0196-6774\(82\)90007-4](https://doi.org/10.1016/0196-6774(82)90007-4).
- [26] Jeannerod, C.P., Neiger, V., Schost, E., Villard, G., 2016. Fast computation of minimal interpolation bases in Popov form for arbitrary shifts, in: ISSAC'16, ACM, pp. 295–302. doi:[10.1145/2930889.2930928](https://doi.org/10.1145/2930889.2930928).
- [27] Jeannerod, C.P., Neiger, V., Schost, E., Villard, G., 2017. Computing minimal interpolation bases. *J. Symbolic Comput.* 83, 272–314. doi:[10.1016/j.jsc.2016.11.015](https://doi.org/10.1016/j.jsc.2016.11.015).
- [28] Jeannerod, C.P., Neiger, V., Villard, G., 2020. Fast computation of approximant bases in canonical form. *J. Symbolic Comput.* 98, 192–224. doi:[10.1016/j.jsc.2019.07.011](https://doi.org/10.1016/j.jsc.2019.07.011).
- [29] Kailath, T., 1980. *Linear Systems*. Prentice-Hall.
- [30] Kaltofen, E., Villard, G., 2005. On the complexity of computing determinants. *Computational Complexity* 13, 91–130. doi:[10.1007/s00037-004-0185-3](https://doi.org/10.1007/s00037-004-0185-3).
- [31] Keller-Gehrig, W., 1985. Fast algorithms for the characteristic polynomial. *Theoretical Computer Science* 36, 309–317. doi:[10.1016/0304-3975\(85\)90049-0](https://doi.org/10.1016/0304-3975(85)90049-0).
- [32] Labahn, G., Neiger, V., Zhou, W., 2017. Fast, deterministic computation of the Hermite normal form and determinant of a polynomial matrix. *J. Complexity* 42, 44–71. doi:[10.1016/j.jco.2017.03.003](https://doi.org/10.1016/j.jco.2017.03.003).
- [33] Le Gall, F., 2014. Powers of tensors and fast matrix multiplication, in: ISSAC'14, ACM, pp. 296–303. doi:[10.1145/2608628.2608664](https://doi.org/10.1145/2608628.2608664).
- [34] Le Verrier, U., 1840. Sur les variations séculaires des éléments elliptiques des sept planètes principales. *Journal des Mathématiques Pures et Appliquées* 5, 220–254.
- [35] Manthey, W., Helmke, U., 2007. Bruhat canonical form for linear systems. *Linear Algebra Appl.* 425, 261–282. doi:[10.1016/j.laa.2007.01.022](https://doi.org/10.1016/j.laa.2007.01.022).
- [36] Mulders, T., Storjohann, A., 2003. On lattice reduction for polynomial matrices. *J. Symbolic Comput.* 35, 377–401. doi:[10.1016/S0747-7171\(02\)00139-6](https://doi.org/10.1016/S0747-7171(02)00139-6).
- [37] Neiger, V., Rosenkilde, J., Solomatov, G., 2018. Computing Popov and Hermite Forms of Rectangular Polynomial Matrices, in: ISSAC'18, ACM, pp. 295–302. doi:[10.1145/3208976.3208988](https://doi.org/10.1145/3208976.3208988).
- [38] Neiger, V., Vu, T.X., 2017. Computing canonical bases of modules of univariate relations, in: ISSAC'17, ACM, pp. 357–364. doi:[10.1145/3087604.3087656](https://doi.org/10.1145/3087604.3087656).
- [39] Pernet, C., Storjohann, A., 2007. Faster Algorithms for the Characteristic Polynomial, in: ISSAC'07, ACM, pp. 307–314. doi:[10.1145/1277548.1277590](https://doi.org/10.1145/1277548.1277590).
- [40] Popov, V.M., 1972. Invariant description of linear, time-invariant controllable systems. *SIAM Journal on Control* 10, 252–264. doi:[10.1137/0310020](https://doi.org/10.1137/0310020).
- [41] Samuelson, P.A., 1942. A method of determining explicitly the coefficients of the characteristic equation. *Annals of Mathematical Statistics* 13, 424–429.
- [42] Sarkar, S., Storjohann, A., 2011. Normalization of row reduced matrices, in: ISSAC'11, ACM, pp. 297–304. doi:[10.1145/1993886.1993931](https://doi.org/10.1145/1993886.1993931).
- [43] Souriau, J.M., 1948. Une méthode pour la décomposition spectrale et l'inversion des matrices. *Comptes-Rendus de l'Académie des Sciences* 227, 1010–1011.
- [44] Storjohann, A., 2001. Deterministic computation of the frobenius form, in: *Proceedings 42nd IEEE Symposium*

- on Foundations of Computer Science, pp. 368–377. doi:[10.1109/SFCS.2001.959911](https://doi.org/10.1109/SFCS.2001.959911).
- [45] Storjohann, A., 2003. High-order lifting and integrality certification. *J. Symbolic Comput.* 36, 613–648. doi:[10.1016/S0747-7171\(03\)00097-X](https://doi.org/10.1016/S0747-7171(03)00097-X).
 - [46] Storjohann, A., 2006. Notes on computing minimal approximant bases, in: *Challenges in Symbolic Computation Software*. URL: <http://drops.dagstuhl.de/opus/volltexte/2006/776>.
 - [47] Strassen, V., 1969. Gaussian elimination is not optimal. *Numer. Math.* 13, 354–356. doi:[10.1007/BF02165411](https://doi.org/10.1007/BF02165411).
 - [48] The FFLAS-FFPACK Group, 2019. FFLAS-FFPACK: Finite Field Linear Algebra Subroutines / Package, version 2.4.3. <http://github.com/linbox-team/fflas-ffpack>.
 - [49] The LinBox Group, 2019. Linbox: Linear algebra over black-box matrices, version 1.6.3. <https://github.com/linbox-team/linbox/>.
 - [50] Toom, A.L., 1963. The complexity of a scheme of functional elements realizing the multiplication of integers. *Soviet Mathematics Doklady* 3, 714–716.
 - [51] Van Barel, M., Bultheel, A., 1992. A general module theoretic framework for vector M-Padé and matrix rational interpolation. *Numer. Algorithms* 3, 451–462. doi:[10.1007/BF02141952](https://doi.org/10.1007/BF02141952).
 - [52] Wolovich, W.A., 1974. *Linear Multivariable Systems*. volume 11 of *Applied Mathematical Sciences*. Springer-Verlag New-York. doi:[10.1007/978-1-4612-6392-0](https://doi.org/10.1007/978-1-4612-6392-0).
 - [53] Zhou, W., 2012. Fast Order Basis and Kernel Basis Computation and Related Problems. Ph.D. thesis. University of Waterloo. URL: <http://hdl.handle.net/10012/7326>.
 - [54] Zhou, W., Labahn, G., 2009. Efficient computation of order bases, in: *ISSAC'09*, ACM. pp. 375–382. doi:[10.1145/1576702.1576753](https://doi.org/10.1145/1576702.1576753).
 - [55] Zhou, W., Labahn, G., 2012. Efficient algorithms for order basis computation. *J. Symbolic Comput.* 47, 793–819. doi:[10.1016/j.jsc.2011.12.009](https://doi.org/10.1016/j.jsc.2011.12.009).
 - [56] Zhou, W., Labahn, G., 2013. Computing column bases of polynomial matrices, in: *ISSAC'13*, ACM. pp. 379–386. doi:[10.1145/2465506.2465947](https://doi.org/10.1145/2465506.2465947).
 - [57] Zhou, W., Labahn, G., 2014. Unimodular completion of polynomial matrices, in: *ISSAC'14*, ACM. pp. 413–420. doi:[10.1145/2608628.2608640](https://doi.org/10.1145/2608628.2608640).
 - [58] Zhou, W., Labahn, G., Storjohann, A., 2012. Computing minimal nullspace bases, in: *ISSAC'12*, ACM. pp. 366–373. doi:[10.1145/2442829.2442881](https://doi.org/10.1145/2442829.2442881).