

Rank-Sensitive Computation of the Rank Profile of a Polynomial Matrix

George Labahn

Cheriton School of Computer Science
University of Waterloo, Ontario, Canada

Thi Xuan Vu

Department of Mathematics and Statistics
UiT, The Arctic University of Norway, Tromsø, Norway

Vincent Neiger

Sorbonne Université, CNRS, LIP6
F-75005 Paris, France

Wei Zhou

Cheriton School of Computer Science
University of Waterloo, Ontario, Canada

ABSTRACT

Consider a matrix $F \in \mathbb{K}[x]^{m \times n}$ of univariate polynomials over a field \mathbb{K} . We study the problem of computing the column rank profile of F . To this end we first give an algorithm which improves the minimal kernel basis algorithm of Zhou, Labahn, and Storjohann (Proceedings ISSAC 2012). We then provide a second algorithm which computes the column rank profile of F with a rank-sensitive complexity of $O^*(r^{\omega-2}n(m+D))$ operations in \mathbb{K} . Here, D is the sum of row degrees of F , ω is the exponent of matrix multiplication, and $O^*(\cdot)$ hides logarithmic factors.

CCS CONCEPTS

• Computing methodologies → Algebraic algorithms; • Theory of computation → Design and analysis of algorithms.

KEYWORDS

Polynomial matrix; kernel basis; rank profile; complexity.

ACM Reference Format:

George Labahn, Vincent Neiger, Thi Xuan Vu, and Wei Zhou. 2022. Rank-Sensitive Computation of the Rank Profile of a Polynomial Matrix. In *Proceedings of the 2022 International Symposium on Symbolic and Algebraic Computation (ISSAC '22)*, July 4–7, 2022, Villeneuve-d'Ascq, France. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3476446.3535495>

1 INTRODUCTION

In this paper, we consider the computation of rank properties of a univariate polynomial matrix $F \in \mathbb{K}[x]^{m \times n}$ over some base field \mathbb{K} . The rank of F can be determined by computing a basis for the left (or for the right) kernel of F . Under the assumption $m \geq n$ (which implicitly requires the input matrix to have full rank, see Section 4), an algorithm due to Zhou, Labahn, and Storjohann [39] computes a minimal basis for the left kernel of F using $O^*(m^\omega \lceil \rho/n \rceil)$ operations in \mathbb{K} , where ρ is the sum of the n largest row degrees of F . In this cost bound, ω is the exponent of matrix multiplication, and $O^*(\cdot)$ is $O(\cdot)$ but ignoring logarithmic factors. A natural alternative is

to compute a basis for the row space of F , called a row basis (or, similarly, a column basis). However, the fastest known row basis algorithm [37] starts by computing a basis of the left kernel of F , so one may as well get the rank directly from the latter.

Currently the best known cost bound for computing the rank of F only depends on the matrix dimensions m and n , and is not influenced by the rank r . More generally, the fastest known algorithms for basic computations with univariate polynomial matrices are not rank-sensitive. This is a significant drawback for the manipulation of matrices whose rank is unknown, and possibly low a priori.

Furthermore, there are specific situations where rank deficiency is actually expected by design, and one would like to take advantage of this in algorithms. Recently, in a context of computing generators of linearly recurrent sequences, minimal approximant bases of rank-deficient, structured matrices F have been encountered [13, Sec. 5]. It has also been observed that, for the computation of the Hermite normal form of F , finding the (column) rank profile of F provides a direct reduction to the case of a square, nonsingular matrix [26], for which fast methods are known [19]. A third situation occurs in verification protocols for polynomial matrices: most protocols proposed in [20] rely, directly or indirectly, on one for certifying “rank(F) $\geq \gamma$ ” [20, Prot. 3], itself asking the Prover to locate a square, nonsingular submatrix of F which has rank at least γ .

In fast \mathbb{K} -linear algebra, rank-sensitive algorithms and complexity bounds have proved highly valuable. For example, rank-sensitive Gaussian elimination costs $O(r^{\omega-2}mn)$ operations in \mathbb{K} [14, 29, 31], for an input (constant) matrix $A \in \mathbb{K}^{m \times n}$ of rank r . More recently, research on this topic has led to improvements of both complexity bounds and software implementations, and has also provided deep insight into the rank-related properties that are revealed depending on the chosen elimination strategies [6, 17]. For finding the rank or rank profile and for solving linear systems, [5, 33] report on running times as low as $(r^\omega + m + n + |A|)^{1+o(1)}$, with $|A|$ the number of nonzero entries of A . Now, for univariate polynomial matrices, despite the impact this would have on many computations, there is still an important lack of efficient rank-sensitive methods which would incorporate both fast linear algebra techniques and fast univariate polynomial multiplication.

One possibility is to make use of classical algorithms such as fraction-free Gaussian elimination (see [8]) while also keeping track of row or column operations to obtain a kernel basis and rank profile. The cost of such algorithms depends on m , n , r and the degree of matrices but does not involve the exponent of matrix multiplication ω . This is also the case for the algorithm of Mulders and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ISSAC '22, July 4–7, 2022, Villeneuve-d'Ascq, France

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8688-3/22/07...\$15.00

<https://doi.org/10.1145/3476446.3535495>

Storjohann [22] which transforms \mathbf{F} to weak Popov form and computes the rank profile with a cost of $O(rmn \deg(\mathbf{F})^2)$. Storjohann [29, Chap. 2] gives a recursive version of fraction-free Gaussian elimination which computes a kernel basis and rank profile of \mathbf{F} having complexity of $O(r^{\omega-1}mn \deg(\mathbf{F})^2)$ operations in \mathbb{K} . Storjohann and Villard [32] later gave a Las Vegas randomized algorithm which computed the rank and kernel basis of a polynomial matrix with complexity $O(r^{\omega-2}mn \deg(\mathbf{F}))$.

The main contribution of this paper is a column rank profile algorithm with a rank-sensitive cost of $O(r^{\omega-2}n(m+D))$ operations in \mathbb{K} . Here D is the sum of the row degrees of \mathbf{F} , with $D \leq m \deg(\mathbf{F})$. This is a follow-up and improvement to the algorithm given in the PhD thesis of Zhou [35, Sec. 11].

We first revisit [35, Algo. 11.1], to augment the minimal kernel basis algorithm of [39] so that it also determines the column rank profile of the input matrix. How the variant here improves upon those in the last two references is explained at the beginning of Section 4. In particular, within the same complexity bound, the new version supports arbitrary dimensions m, n and rank r of \mathbf{F} , which is essential for our purpose. This algorithm is not rank-sensitive: it has a cost of $O(m^{\omega-2}(m+n)(m+D))$ operations in \mathbb{K} .

We then give a rank-sensitive column rank profile algorithm, which uses the above kernel basis algorithm as its main subroutine. A sketch of a similar result has been given before in [35, Sec. 11.2], where the approach is to incorporate the *columns* gradually, always maintaining a number of columns which is bounded by the rank. At each step the above kernel basis procedure is called to obtain a partial column rank profile and discard rows that are $\mathbb{K}[x]$ -linearly dependent. At each step as well, to keep control of the cost of this kernel computation, a row basis computation is applied beforehand to reduce to a matrix having full row rank.

Here, we follow another path, by incorporating *rows* gradually. This allows us to benefit from the fact that the kernel procedure has quasi-linear cost with respect to the column dimension n , without having to resort to row basis computations. To enable proceeding row-wise, we exploit a property of kernel bases in so-called weak Popov form, showing that they give direct access to a set of linearly independent rows of the input in addition to its column rank profile. Once all rows of \mathbf{F} have been processed and a set of r linearly independent rows of \mathbf{F} has been found, the column rank profile of \mathbf{F} can be extracted efficiently again through the kernel algorithm.

Outline. In Section 2, we give the basic definitions and properties of our building blocks for polynomial matrix arithmetic including kernel bases, pivot profiles and rank profiles, and weak Popov forms. Section 3 introduces specific rank profile and kernel properties used in our algorithms. Section 4 describes our algorithm for computing the rank profile and kernel basis, while Section 5 presents our algorithm for the rank-sensitive computation of the rank profile. The paper ends with topics for future research, and contains as an appendix an illustration of our examples through SageMath code.

2 PRELIMINARIES

In this section we describe the notations used in this paper, and then give the basic definitions and a number of properties of polynomial matrices including *shifted degrees* and *pivot profiles*, *relation bases* and *kernel bases*, *reduced forms* and *weak Popov forms*.

2.1 Notation

We let $\mathbb{K}[x]$ denote a univariate polynomial ring over a field \mathbb{K} with $\mathbb{K}[x]^{m \times n}$ being the set of $m \times n$ univariate polynomial matrices. For $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$ and subsets I of $(1, \dots, m)$ and J of $(1, \dots, n)$, we write $\mathbf{F}_{I,J}$ for the submatrix of \mathbf{F} obtained by selecting rows indexed by I and columns indexed by J . We let $\mathbf{F}_{I,*} = \mathbf{F}_{I,\{1..n\}}$ denote the submatrix of \mathbf{F} obtained by selecting the rows indexed by I and keeping all columns and $\mathbf{F}_{*,J} = \mathbf{F}_{\{1..m\},J}$ for the submatrix of \mathbf{F} obtained by keeping all rows and selecting columns indexed by J .

For a tuple of integers $\mathbf{s} = (s_1, \dots, s_m) \in \mathbb{Z}^m$, the sum of its entries is denoted by $|\mathbf{s}| = s_1 + \dots + s_m$. When this concerns an input shift \mathbf{s} , we will often write D for this quantity, i.e. $D = |\mathbf{s}|$.

2.2 Kernel, row space, modules of relations

For a matrix \mathbf{F} in $\mathbb{K}[x]^{m \times n}$ of rank r , the set

$$\mathcal{K}(\mathbf{F}) := \{\mathbf{p} \in \mathbb{K}[x]^{1 \times m} \mid \mathbf{p}\mathbf{F} = \mathbf{0}\}$$

is a $\mathbb{K}[x]$ -module of rank $m - r$ and is called the (left) *kernel* of \mathbf{F} . The *row space* of \mathbf{F} is the module

$$\{\mathbf{p}\mathbf{F} \mid \mathbf{p} \in \mathbb{K}[x]^{1 \times m}\} \subseteq \mathbb{K}[x]^{1 \times n}.$$

A basis for one of these modules (a kernel basis or a row basis) is typically organized into a single polynomial matrix, for example, a basis of $\mathcal{K}(\mathbf{F})$ being represented by a full rank matrix $\mathbf{K} \in \mathbb{K}[x]^{(m-r) \times m}$. Also, for a nonsingular matrix \mathbf{M} in $\mathbb{K}[x]^{n \times n}$,

$$\mathcal{R}_{\mathbf{M}}(\mathbf{F}) := \{\mathbf{p} \in \mathbb{K}[x]^{1 \times m} \mid \mathbf{p}\mathbf{F} = \mathbf{0} \bmod \mathbf{M}\}$$

is a $\mathbb{K}[x]$ -module of rank m , called the (left) *relation module* of \mathbf{F} modulo \mathbf{M} . Here, the notation $\mathbf{A} = \mathbf{0} \bmod \mathbf{M}$ means that $\mathbf{A} = \mathbf{Q}\mathbf{M}$ for some matrix \mathbf{Q} .

Important particular cases are the relations of *approximation* and those of *interpolation* [1, 2, 34]. For the latter, $\mathbf{M} = \text{diag}(M_1, \dots, M_n)$ with $M_k = (x - \alpha_{k,1}) \cdots (x - \alpha_{k,\tau_k})$ for some $\tau_k \in \mathbb{Z}_{>0}$ and $1 \leq k \leq n$, where the $\alpha_{k,j}$'s are known elements from \mathbb{K} . Approximation is when these elements are zero: $\mathbf{M} = \text{diag}(x^{\tau_1}, \dots, x^{\tau_n})$, so that working mod \mathbf{M} amounts to truncating the column j modulo x^{τ_j} .

The notions of right kernel, column space, column bases, and right relations are of course defined similarly.

2.3 Shifted degrees, leading matrix

For a row vector $\mathbf{p} = [p_1 \cdots p_n]$ in $\mathbb{K}[x]^{1 \times n}$, its *degree* is

$$\text{rdeg}(\mathbf{p}) = \max_{1 \leq i \leq n} \deg(p_i)$$

that is, the largest degree of all its entries. Here we take the convention that the degree of a zero polynomial or zero row is $-\infty$. In many cases it is useful to shift (or re-weigh) the degrees. Given a shift $\mathbf{s} = (s_1, \dots, s_n) \in \mathbb{Z}^n$, the *s-degree* of \mathbf{p} is defined as

$$\text{rdeg}_{\mathbf{s}}(\mathbf{p}) = \max_{1 \leq i \leq n} (s_i + \deg(p_i)).$$

Note that this is equal to $\text{rdeg}(\mathbf{p}\mathbf{x}^{\mathbf{s}})$, where $\mathbf{x}^{\mathbf{s}}$ is the diagonal matrix with diagonal entries x^{s_1}, \dots, x^{s_n} .

For a matrix $\mathbf{P} \in \mathbb{K}[x]^{m \times n}$ and a shift $\mathbf{s} \in \mathbb{Z}^n$, the row degree $\text{rdeg}(\mathbf{P})$ of \mathbf{P} is the list of the degrees of its rows, and similarly the *s-row degree* $\text{rdeg}_{\mathbf{s}}(\mathbf{P})$ is the list of the *s*-degrees of its rows. Then, the *s-leading matrix* $\text{lm}_{\mathbf{s}}(\mathbf{P})$ of \mathbf{P} is the matrix in $\mathbb{K}^{m \times n}$ formed by the coefficients of degree zero of $\mathbf{x}^{-\mathbf{t}}\mathbf{P}\mathbf{x}^{\mathbf{s}}$, where $\mathbf{t} = \text{rdeg}_{\mathbf{s}}(\mathbf{P})$. By convention, a zero row in \mathbf{P} yields a zero row in $\text{lm}_{\mathbf{s}}(\mathbf{P})$.

2.4 Pivot and rank profiles

If the row vector \mathbf{p} is nonzero, the s -pivot index of \mathbf{p} is the largest index π such that $\deg(p_\pi) + s_\pi = \text{rdeg}_s(\mathbf{p})$. In this case p_π and $\deg(p_\pi)$ are the s -pivot entry and the s -pivot degree of \mathbf{p} . Note that π is also the index of the rightmost nonzero entry in $\text{lm}_s(\mathbf{p})$.

The pair $(\boldsymbol{\pi}, \boldsymbol{\delta}) = (\pi_i, \delta_i)_{1 \leq i \leq m}$ where $\boldsymbol{\pi} = (\pi_i)_{1 \leq i \leq m}$ and $\boldsymbol{\delta} = (\delta_i)_{1 \leq i \leq m}$ are the s -pivot index and degree for each row of the matrix \mathbf{P} , is called the s -pivot profile of \mathbf{P} . Observe that $\text{rdeg}_s(\mathbf{P})$ is equal to $(\delta_i + s_{\pi_i})_{1 \leq i \leq m}$.

The (column) rank profile of \mathbf{P} is the lexicographically minimal list of integers $J = (j_1, \dots, j_r)$ such that $\mathbf{P}_{*,J}$ has rank $r = \text{rank}(\mathbf{P})$. In what follows, rank profile always means column rank profile; otherwise, we will write explicitly row rank profile.

A matrix $\mathbf{H} = [h_{i,j}] \in \mathbb{K}[x]^{r \times n}$ with $r \leq n$ is in *Hermite normal form* [11, 21, 28] if there are indices $1 \leq j_1 < \dots < j_r \leq n$ such that

- (i) for $1 \leq i \leq r$, $h_{i,j_i} \neq 0$ is monic and $h_{i,j} = 0$ for $1 \leq j < j_i$;
- (ii) for $1 \leq k < i \leq r$, $\deg(h_{k,j_i}) < \deg(h_{i,j_i})$.

In this case, (j_1, \dots, j_r) is the rank profile of \mathbf{H} .

The Hermite normal form of $\mathbf{P} \in \mathbb{K}[x]^{m \times n}$ is its unique row basis $\mathbf{H} \in \mathbb{K}[x]^{r \times n}$, with $r = \text{rank}(\mathbf{P})$, which is in Hermite normal form. Then, the rank profile of \mathbf{P} is equal to that of \mathbf{H} , since $\mathbf{U}\mathbf{P} = \begin{bmatrix} \mathbf{H} \\ \mathbf{0} \end{bmatrix}$ for some unimodular matrix $\mathbf{U} \in \mathbb{K}[x]^{m \times m}$.

2.5 Reduced forms, predictable degree

With the above definitions, \mathbf{P} is said to be s -row reduced if $\text{lm}_s(\mathbf{P})$ has full row rank. A core feature of these matrices is the *predictable degree property*, which says that there cannot be any cancellation of high-degree terms of the matrix via $\mathbb{K}[x]$ -linear combinations of the rows (see [7, 18] for the case $s = \mathbf{0}$; [3, Lem. 3.6] for any s ; and [23, Thm. 1.11] for a proof of the equivalence in the next lemma).

LEMMA 2.1 (PREDICTABLE DEGREE). *Let $\mathbf{P} \in \mathbb{K}[x]^{m \times n}$ have no zero row, let $s \in \mathbb{Z}^n$ and $t = \text{rdeg}_s(\mathbf{P})$. Then, \mathbf{P} is in s -reduced form if and only if $\text{rdeg}_s(\mathbf{QP}) = \text{rdeg}_t(\mathbf{Q})$ for all $\mathbf{Q} \in \mathbb{K}[x]^{k \times m}$.*

COROLLARY 2.2. *Let $\mathbf{P} \in \mathbb{K}[x]^{m \times n}$ and let $s \in \mathbb{Z}^n$ be such that \mathbf{P} is s -reduced. Let $t = \text{rdeg}_s(\mathbf{P})$. Then, $\text{lm}_s(\mathbf{QP}) = \text{lm}_t(\mathbf{Q})\text{lm}_s(\mathbf{P})$ for any $\mathbf{Q} \in \mathbb{K}[x]^{k \times m}$.*

PROOF. Let $\mathbf{d} = \text{rdeg}_t(\mathbf{Q}) \in \mathbb{Z}^k$. By the predictable degree property, $\mathbf{d} = \text{rdeg}_s(\mathbf{QP})$. The conclusion then follows from the identity

$$\mathbf{x}^{-\mathbf{d}} \mathbf{QP} \mathbf{x}^{\mathbf{s}} = (\mathbf{x}^{-\mathbf{d}} \mathbf{Q} \mathbf{x}^{\mathbf{t}})(\mathbf{x}^{-\mathbf{t}} \mathbf{P} \mathbf{x}^{\mathbf{s}}). \quad \square$$

As a consequence, shifted reduced forms are preserved by multiplication, provided the shifts are appropriately chosen. This result is at the core of divide and conquer algorithms for bases of relation modules [2, 9] and kernel bases [39, Thm. 3.9].

LEMMA 2.3. *Let $\mathbf{P} \in \mathbb{K}[x]^{m \times n}$ and $s \in \mathbb{Z}^n$ such that \mathbf{P} is in s -reduced form. Let $t = \text{rdeg}_s(\mathbf{P}) \in \mathbb{Z}^m$, and let $\mathbf{Q} \in \mathbb{K}[x]^{k \times m}$ be in t -reduced form. Then, \mathbf{QP} is in s -reduced form.*

PROOF. The assumptions imply that both $\text{lm}_t(\mathbf{Q})$ and $\text{lm}_s(\mathbf{P})$ have full row rank. Thus their product has full row rank as well, and according to Corollary 2.2 this product is $\text{lm}_s(\mathbf{QP})$. \square

2.6 Weak Popov forms, predictable pivot

A matrix $\mathbf{P} = [p_{i,j}] \in \mathbb{K}[x]^{r \times n}$ with no zero row is s -Popov if

- (i) its s -pivot index (π_1, \dots, π_r) is strictly increasing;
- (ii) for $1 \leq i \leq r$, p_{i,π_i} is monic;
- (iii) for each $k, i \in \{1, \dots, r\}$ with $k \neq i$, $\deg(p_{k,\pi_i}) < \deg(p_{i,\pi_i})$.

If \mathbf{P} only satisfies the first condition, it is said to be s -weak Popov. Any s -weak Popov matrix is s -reduced. Furthermore, each matrix has a unique row basis in s -Popov form.

We remark that, for weak Popov forms, it is sometimes only required (see e.g. [22]) that the pivot indices be pairwise distinct, instead of increasing. Then, the forms with the added requirement of increasing indices were called ordered weak Popov forms. Here, we will only manipulate ordered weak Popov forms, and therefore we call them weak Popov forms for ease of presentation.

The shifted weak Popov form satisfies the following refinement of the predictable degree property and is also compatible with multiplication under well-chosen shifts (see [4, Sec. 5] for related considerations and [25, Lem. 2.6] for a proof of the next lemmas).

LEMMA 2.4 (PREDICTABLE PIVOT). *Let $\mathbf{P} \in \mathbb{K}[x]^{m \times n}$ have no zero row, let $s \in \mathbb{Z}^n$ and $t = \text{rdeg}_s(\mathbf{P})$, and let $(\pi_i, \delta_i)_{1 \leq i \leq m}$ be the s -pivot profile of \mathbf{P} . If \mathbf{P} is in s -weak Popov form, then the s -pivot profile of \mathbf{QP} is $(\pi_{j_i}, \delta_{j_i} + d_i)_{1 \leq i \leq k}$ for all $\mathbf{Q} \in \mathbb{K}[x]^{k \times m}$, where $(j_i, d_i)_{1 \leq i \leq k}$ is the t -pivot profile of \mathbf{Q} .*

LEMMA 2.5. *Let $\mathbf{P} \in \mathbb{K}[x]^{m \times n}$ and $s \in \mathbb{Z}^n$ such that \mathbf{P} is in s -weak Popov form. Let $t = \text{rdeg}_s(\mathbf{P}) \in \mathbb{Z}^m$, and let $\mathbf{Q} \in \mathbb{K}[x]^{k \times m}$ be in t -weak Popov form. Then, \mathbf{QP} is in s -weak Popov form.*

PROOF. By assumption, the s -pivot index $(\pi_i)_{1 \leq i \leq m}$ of \mathbf{P} and the t -pivot index $(j_i)_{1 \leq i \leq k}$ of \mathbf{Q} are both strictly increasing. Then, by Lemma 2.4, the s -pivot index of \mathbf{QP} is the subtuple $(\pi_{j_i})_{1 \leq i \leq k}$, which is strictly increasing. Hence \mathbf{QP} is in s -weak Popov form. \square

Note however that a similar product of shifted Popov forms does not yield a shifted Popov form, but only a shifted weak Popov form.

2.7 Example

We will use the following as a running example in this paper.

Example 2.6. Working over $\mathbb{K} = \mathbb{F}_2$, let $\mathbf{F} \in \mathbb{K}[x]^{5 \times 5}$ be given by

$$\begin{bmatrix} x^2 & x^3 + 1 & x^8 + x^6 + x^4 + x^3 + x^2 + x & x^4 + 1 & x^3 + 1 \\ 0 & x^4 + 1 & x^5 + x^4 + x^3 + x^2 & x + 1 & x^2 + 1 \\ 0 & x^2 + 1 & x + 1 & 0 & 1 \\ 0 & 0 & x^8 + 1 & x^4 + 1 & 0 \\ 0 & 0 & x^4 + 1 & 1 & 0 \end{bmatrix}$$

Then the matrix

$$\begin{bmatrix} 0 & 1 & x^2 + 1 & 0 & x + 1 \\ 0 & 1 & x^2 + 1 & 1 & x^4 + x \end{bmatrix} \in \mathbb{K}[x]^{2 \times 5}$$

is a weak Popov basis of $\mathcal{K}(\mathbf{F})$ (which is not in Popov form). It has pivot index $\boldsymbol{\pi} = (3, 5)$ and pivot degree $\boldsymbol{\delta} = (2, 4)$. Here is now an s -weak Popov basis of $\mathcal{K}(\mathbf{F})$ for the shift $\mathbf{s} = \text{rdeg}(\mathbf{F}) = (8, 5, 2, 8, 4)$:

$$\mathbf{K} = \begin{bmatrix} 0 & x^3 & x^5 + x^3 & 1 & x^3 + 1 \\ 0 & 1 & x^2 + 1 & 0 & x + 1 \end{bmatrix} \in \mathbb{K}[x]^{2 \times 5}.$$

Its s -pivot index is $(4, 5)$ and its s -pivot degree is $(0, 1)$. \square

The above kernel bases were computed using the SageMath software, as described in Fig. 1 on Page 9.

3 RANK AND DEGREE PROPERTIES RELATED TO KERNEL BASES

In this section we discuss rank and degree properties related to kernel bases and which are central for the correctness and complexity of the algorithms in Sections 4 and 5.

LEMMA 3.1. *Let $F \in \mathbb{K}[x]^{m \times n}$ have rank r . For any $V \in \mathbb{K}[x]^{n \times k}$ such that FV has rank r , we have $\mathcal{K}(F) = \mathcal{K}(FV)$. As a corollary, if $J \subseteq \{1, \dots, n\}$ is such that $F_{*,J}$ has rank r , then $\mathcal{K}(F) = \mathcal{K}(F_{*,J})$.*

PROOF. The second statement follows from the first one, by building V from the columns of the identity matrix I_n with index in J .

Concerning the first statement, the rank assumption implies that the left kernels $\mathcal{K}(F)$ and $\mathcal{K}(FV)$ have the same rank $m - r$. Let $K_1 \in \mathbb{K}[x]^{(m-r) \times m}$ and $K_2 \in \mathbb{K}[x]^{(m-r) \times m}$ be bases of $\mathcal{K}(F)$ and $\mathcal{K}(FV)$, respectively. It is clear that $\mathcal{K}(F) \subseteq \mathcal{K}(FV)$, hence $K_1 = UK_2$ for some nonsingular $U \in \mathbb{K}[x]^{(m-r) \times (m-r)}$. The fact that kernel bases have unimodular column bases [38] ensures that U is unimodular, and thus $\mathcal{K}(F) = \mathcal{K}(FV)$. \square

LEMMA 3.2. *If $F \in \mathbb{K}[x]^{m \times n}$ and $G \in \mathbb{K}[x]^{\ell \times n}$ are two matrices which have the same right kernel, then F and G have the same column rank profile. As a corollary, if $I \subseteq \{1, \dots, m\}$ is such that $F_{I,*}$ has the same rank as F , then $F_{I,*}$ has the same rank profile as F .*

PROOF. The second statement follows from the first: applying Lemma 3.1 to F^T and $(F_{I,*})^T = (F^T)_{*,I}$ shows that these matrices have the same left kernel, i.e. F and $F_{I,*}$ have the same right kernel.

Let J_1 and J_2 be the rank profiles of F and G , respectively. Let $j \in \{1, \dots, n\}$ be such that $j \notin J_1$, meaning that there exists a vector $u = [u_1 \cdots u_j]^T \in \mathbb{K}[x]^{j \times 1}$ such that $u_j \neq 0$ and $F_{*,1..j}u = 0$. Since the right kernel of F is contained in that of G , it follows that $G_{*,1..j}u = 0$, and therefore $j \notin J_2$. We have proved $J_2 \subseteq J_1$, and the same arguments prove $J_1 \subseteq J_2$, by symmetry. Hence $J_2 = J_1$. \square

THEOREM 3.3. *Let $F \in \mathbb{K}[x]^{m \times n}$ have rank r , let $s \in \mathbb{Z}^m$, and let $K \in \mathbb{K}[x]^{(m-r) \times m}$ be an s -weak Popov basis of $\mathcal{K}(F)$. Let (π, δ) be the s -pivot profile of K , and let $\pi^c = \{1, \dots, m\} \setminus \pi$ be the indices of the columns of K which do not contain an s -pivot entry of K . Assume also that F factors as $F = SR$ where $R \in \mathbb{K}[x]^{r \times n}$ and $S \in \mathbb{K}[x]^{m \times r}$. Then,*

- (i) $F_{\pi^c,*} \in \mathbb{K}[x]^{r \times n}$ has rank r , which is the size of π^c ;
- (ii) $S_{\pi^c,*} \in \mathbb{K}[x]^{r \times r}$ is nonsingular and $|\delta| \leq \deg(\det(S_{\pi^c,*}))$, hence in particular $|\delta| \leq |\text{rdeg}(F_{\pi^c,*})| \leq r \deg(F)$;
- (iii) if $s \geq \text{rdeg}(F)$, then $|\text{rdeg}_s(K)| \leq |s|$.

Concerning the matrices R and S , note that they have rank r , since otherwise we would have $\text{rank}(F) = \text{rank}(SR) < r$. Taking a row basis of F for R proves the existence of such matrices.

Item (i) states that, from any shifted weak Popov basis of the left kernel of F , we can immediately deduce a set of $r = \text{rank}(F)$ rows of F which are $\mathbb{K}[x]$ -linearly independent. Item (iii) is the main degree property that was exploited in the design of the fastest known minimal kernel basis algorithm [39] (see [39, Thm. 3.4]), explaining also why this algorithm restricts to shifts such that $s \geq \text{rdeg}(F)$. Here we prove it as a consequence of the property in Item (ii), which gives more precise degree information in particular through better accounting for the rank of F .

We now prove Theorem 3.3.

PROOF. For Item (i) it suffices to prove that $\mathcal{K}(F_{\pi^c,*}) = \{0\}$. Let $v \in \mathbb{K}[x]^{1 \times r}$ be such that $vF_{\pi^c,*} = 0$. Construct $w \in \mathbb{K}[x]^{1 \times m}$ such that $w_{*,\pi} = 0$ and $w_{*,\pi^c} = v$. The vector w is in $\mathcal{K}(F)$, so $w = uK$ for some $u \in \mathbb{K}[x]^{1 \times (m-r)}$, and hence $0 = w_{*,\pi} = uK_{*,\pi}$. Thus $u = 0$ since $K_{*,\pi}$ is nonsingular, implying $v = w_{*,\pi^c} = uK_{*,\pi^c} = 0$.

To prove Item (ii) set

$$S_1 = S_{\pi^c,*} \in \mathbb{K}[x]^{r \times r} \text{ and } S_2 = S_{\pi,*} \in \mathbb{K}[x]^{(m-r) \times r}$$

as well as

$$K_1 = K_{*,\pi^c} \in \mathbb{K}[x]^{(m-r) \times r} \text{ and } K_2 = K_{*,\pi} \in \mathbb{K}[x]^{(m-r) \times (m-r)}.$$

Then S_1 is nonsingular since $F_{\pi^c,*} = S_1R$ has rank r . We are going to prove that K_2 is an s_{π} -weak Popov basis of $\mathcal{R}_{S_1}(S_2)$, where $s_{\pi} \in \mathbb{Z}^{m-r}$ is the subshift of s formed by its entries with index in π . From this, [27, Cor. 2.4] ensures that $\deg(\det(K_2)) = |\delta| \leq \deg(\det(S_1))$.

Since $F = SR$, with R full row rank, we have $\mathcal{K}(F) = \mathcal{K}(S)$ and so K is an s -weak Popov basis of $\mathcal{K}(S)$. From $KS = 0$ we obtain $K_1S_1 + K_2S_2 = 0$ and so the rows of K_2 are in $\mathcal{R}_{S_1}(S_2)$. It remains to show that any $p \in \mathcal{R}_{S_1}(S_2)$ is a $\mathbb{K}[x]$ -linear combination of the rows of K_2 . By definition of $\mathcal{R}_{S_1}(S_2)$, there exists $q \in \mathbb{K}[x]^{1 \times r}$ such that $pS_2 = qS_1$. Considering $v \in \mathbb{K}[x]^{1 \times m}$ such that $v_{*,\pi} = p$ and $v_{*,\pi^c} = -q$, we have $vS = v_{*,\pi^c}S_1 + v_{*,\pi}S_2 = -qS_1 + pS_2 = 0$, that is, $v \in \mathcal{K}(S)$. Thus, $v = uK$ for some $u \in \mathbb{K}[x]^{1 \times (m-r)}$, and we obtain $p = v_{*,\pi} = uK_{*,\pi} = uK_2$.

Let $F_1 = F_{\pi^c,*} \in \mathbb{K}[x]^{r \times n}$ and $t = \text{rdeg}(F_1) \in \mathbb{Z}^r$. In order to prove the last two bounds on $|\delta|$, observe that $|t| \leq r \deg(F)$ is clear since F_1 consists of r rows of F . It remains to show that $\deg(\det(S_1)) \leq |t|$. Let $U \in \mathbb{K}[x]^{n \times r}$ be such that F_1U is the $-t$ -column Popov form of F_1 . Since $\text{cdeg}_{-t}(F_1) \leq 0$, the minimality of the shifted column degrees of shifted reduced forms [35, Sec. 2.7] implies $\text{cdeg}_{-t}(F_1U) \leq 0$ as well. According to [37, Lem. 2.2], this translates as $\text{rdeg}(F_1U) \leq t$, and so $|\text{rdeg}(F_1U)| \leq |t|$. Since F_1U is $-t$ -column Popov, it is also row reduced, and therefore $|\text{rdeg}(F_1U)| = \deg(\det(F_1U))$ [18, Sec. 6.3.2]. It follows that $\deg(\det(F_1U)) \leq |t|$ and, using $F_1 = S_1R$, we obtain

$$\deg(\det(S_1)) + \deg(\det(RU)) = \deg(\det(S_1RU)) \leq |t|.$$

To prove Item (iii), recall that $\text{rdeg}_s(K) = (\delta_i + s_{\pi_i})_{1 \leq i \leq m-r}$, and therefore $|\text{rdeg}_s(K)| = |\delta| + |s_{\pi}|$. From Item (ii) we get that $|\text{rdeg}_s(K)| \leq |\text{rdeg}(F_{\pi^c,*})| + |s_{\pi}|$, and from the assumption $s \geq \text{rdeg}(F)$ we conclude that $|\text{rdeg}_s(K)| \leq |s_{\pi^c}| + |s_{\pi}| = |s|$. \square

Example 3.4. Following on from Example 2.6, consider the matrices F and K and the shift $s = \text{rdeg}(F) = (8, 5, 2, 8, 4)$. Since the s -pivot index of K is $\pi = (4, 5)$, the indices of the columns of K which do not contain an s -pivot entry are $\pi^c = (1, 2, 3)$.

Regarding Item (i), from the above s -pivot information we get that the rank of F is 3 and the rows $(1, 2, 3)$ of F are $\mathbb{K}[x]$ -linearly independent. The other kernel basis considered in Example 2.6 shows that the rows $(1, 2, 4)$ are also $\mathbb{K}[x]$ -linearly independent.

Regarding Item (ii), observe that the row degree of $F_{\pi^c,*}$ is $(8, 5, 2)$, so $|\text{rdeg}_s(F_{\pi^c,*})| = 15$. From Example 2.6, the s -pivot degree of K is $\delta = (8, 7)$, so $|\delta| = 15$. Furthermore here $r \deg(F) = 3 \cdot 8 = 24$. Thus, here we have $|\delta| = |\text{rdeg}(F_{\pi^c,*})| \leq r \deg(F)$.

Finally, regarding Item (iii), $|\text{rdeg}_s(K)| = |(8, 5)| = 13$, which is bounded from above by $|s| = 27$. \square

4 COMPUTING THE RANK PROFILE AND A KERNEL BASIS

In this section we give an improved version of the minimal kernel basis algorithm in [39]. In addition to the new algorithm we also include a proof of correctness and determine its complexity.

4.1 Algorithm

Our improvements of the algorithm, compared to the versions in in [39] [35, Sec. 11], is summarized as follows:

- (i) Besides a kernel basis, the algorithm also finds the *column rank profile* of F , without additional operations, based on the approach in [35, Sec. 11.1].
- (ii) The output kernel basis K is in *s-weak Popov form* instead of *s-reduced form*. This has the advantage of revealing the *s-pivot profile*, which can be used for example to further transform K into *s-Popov form* [25, Sec. 5]. Thanks to Item (i) of Theorem 3.3, this also reveals a set of $\text{rank}(F)$ rows of F that are $\mathbb{K}[x]$ -linearly independent, a property that we exploit in Algorithm 2.
- (iii) The algorithm *supports any input matrix* F , without assumption on its rank or dimensions. In comparison, the assumption $m \geq n$ is made in the complexity analysis in [35, 39], which implicitly requires that the input F have full column rank (indeed, if F is rank-deficient, the algorithm in these references cannot guarantee that the assumption $m \geq n$ is satisfied in recursive calls).
- (iv) The algorithm may use *any relation basis* (Lines 17 to 19), instead of restricting to approximant bases. As early experiments have showed [12, Sec. 4.2], this can lead to speed-ups at least by constant factors, for example by relying on well-chosen interpolation bases. Still, as seen in Theorem 4.1 and Section 4.3, for one specific point of the complexity analysis we restrict to relation bases modulo a diagonal matrix.

THEOREM 4.1. *Let $F \in \mathbb{K}[x]^{m \times n}$ have rank r , and let $s \in \mathbb{Z}_{\geq 0}^m$ such that $s \geq \text{rdeg}(F)$. The call $\text{KERNELBASIS-RANKPROFILE}(F, s)$ returns an *s-weak Popov basis* $K \in \mathbb{K}[x]^{(m-r) \times m}$ of $\mathcal{K}(F)$ and the *column rank profile* $(j_1, \dots, j_r) \in \mathbb{Z}_{>0}^r$ of F . Assuming that $m \in O(n)$ and that one chooses a matrix M at Line 18 which is diagonal with all entries of degree τ , this algorithm uses $O(m^{\omega-2}(m+n)(m+D))$ operations in \mathbb{K} , where $D = |s|$.*

Example 4.2. Let $F \in \mathbb{F}_2[x]^{5 \times 5}$ be the matrix from Example 2.6, and consider the shift $s = \text{rdeg}(F) = (8, 5, 2, 8, 4)$. At the top level of the recursion, Algorithm 1 first finds the kernel basis of the 5×2 submatrix $F_1 = F_{*,1,2}$, via a recursive call. With $5 \geq 2 \cdot 2$, this call runs Lines 7 to 13, with $\tau = \lceil \frac{2 \cdot 27}{3} \rceil = 18$. This eventually yields

$$K_1 = \begin{bmatrix} 0 & 1 & x^2 + 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \text{ and rank profile } J_1 = (1, 2).$$

In this case, using an approximant basis of F_1 at order τ yields the three above rows of the kernel, and two additional rows (this can be observed by running the code in Fig. 1 on Page 9). On this specific example it is easily observed that $\text{rank}(F_1) = 2$, so one may infer that K_1 and J_1 are directly deduced from A , without running the recursive call at Line 27.

Algorithm 1 KERNELBASIS-RANKPROFILE(F, s)

Input: a matrix $F \in \mathbb{K}[x]^{m \times n}$, a shift $s \in \mathbb{Z}_{\geq 0}^m$
Requirement: $s \geq \text{rdeg}(F)$ entrywise
Output: an *s-ordered weak Popov basis* $K \in \mathbb{K}[x]^{(m-r) \times m}$ of $\mathcal{K}(F)$ and the *column rank profile* $(j_1, \dots, j_r) \in \mathbb{Z}_{>0}^r$ of F

- 1: **if** $F = 0$ **then** ▷ kernel of zero is identity
- 2: **return** $I_m \in \mathbb{K}[x]^{m \times m}$, $() \in \mathbb{Z}_{>0}^0$
- 3: **if** $m = 1$ **then** ▷ kernel of nonzero $1 \times n$ matrix is empty
- 4: $j \in \{1, \dots, n\} \leftarrow$ index of first nonzero entry in F
- 5: **return** $[] \in \mathbb{K}[x]^{0 \times 1}$, $(j) \in \mathbb{Z}_{>0}^1$
- 6: **if** $m < 2n$ **then** ▷ “wide” matrix: divide and conquer on columns
- 7: $F_1 \in \mathbb{K}[x]^{m \times \lfloor n/2 \rfloor} \leftarrow F_{*, \{1, \dots, \lfloor n/2 \rfloor\}}$
- 8: $K_1 \in \mathbb{K}[x]^{\ell_1 \times m}$, $J_1 \in \mathbb{Z}_{>0}^{r_1} \leftarrow$
KERNELBASIS-RANKPROFILE(F_1, s)
- 9: $F_2 \in \mathbb{K}[x]^{\ell_2 \times \lfloor n/2 \rfloor} \leftarrow K_1 \cdot F_{*, \{\lfloor n/2 \rfloor + 1, \dots, n\}}$
- 10: $K_2 \in \mathbb{K}[x]^{\ell_2 \times \ell_1}$, $J_2 \in \mathbb{Z}_{>0}^{r_2} \leftarrow$
KERNELBASIS-RANKPROFILE($F_2, \text{rdeg}_s(K_1)$)
- 11: ▷ note: $r_1 = \text{rank}(F_1)$, $r_2 = \text{rank}(F_2)$, $\ell_1 = m - r_1$, and $\ell_2 = \ell_1 - r_2$
- 12: $J_2 \in \mathbb{Z}_{>0}^{r_2} \leftarrow$ shift J_2 by adding $\lfloor n/2 \rfloor$ to all entries
- 13: **return** $K_2 \cdot K_1$, $(J_1, J_2) \in \mathbb{Z}_{>0}^{r_1+r_2}$
- 14: ▷ from here we are in the case $F \neq 0$, $m \geq 2$, and $n \leq \frac{m}{2}$
- 15: ▷ minimize shift while preserving $\hat{s} \geq \text{rdeg}(F)$
- 16: $\mu \leftarrow \min(s - \text{rdeg}(F))$; $\hat{s} \leftarrow s - (\mu, \dots, \mu)$; $\tau \leftarrow \left\lceil \frac{2|\hat{s}|}{m-n} \right\rceil$
- 17: ▷ choose type of relations and compute relation basis
- 18: $M \in \mathbb{K}[x]^{n \times n} \leftarrow$ choose any matrix in Hermite normal form with $\min(\text{cdeg}(M)) \geq \tau$ ▷ for example $M = \text{diag}(x^\tau, \dots, x^\tau)$
- 19: $A \leftarrow \hat{s}$ -weak Popov basis of $\mathcal{R}_M(F)$
- 20: ▷ compute residual and indices I of rows already in kernel
- 21: $I \leftarrow \{i \in \{1, \dots, m\} \mid \text{rdeg}_{\hat{s}}(A_{i,*}) < \tau\}$ ▷ rows in kernel
- 22: $I^c \leftarrow \{1, \dots, m\} \setminus I$ ▷ rows expected not to be in kernel
- 23: $G \leftarrow A_{I^c,*} F M^{-1}$ ▷ if $M = \text{diag}(x^\tau, \dots, x^\tau)$, this is $x^{-\tau} A_{I^c,*} F$
- 24: if G has zero rows, update (I, I^c, G) accordingly
- 25: ▷ compute kernel of residual recursively and merge results
- 26: $t \leftarrow \text{rdeg}_{\hat{s}}(A_{I^c,*}) - (\gamma, \dots, \gamma)$ where $\gamma = \min(\text{cdeg}(M))$
- 27: $K_2, J \leftarrow \text{KERNELBASIS-RANKPROFILE}(G, t)$
- 28: $K \leftarrow$ matrix formed by both the rows of $K_2 \cdot A_{I^c,*}$ and those of $A_{I,*}$, sorted by increasing \hat{s} -pivot index
- 29: **return** K, J

Multiplying K_1 by the last three columns of F gives

$$F_2 = \begin{bmatrix} x^5 + x^4 + x + 1 & x + 1 & 0 \\ x^8 + 1 & x^4 + 1 & 0 \\ x^4 + 1 & 1 & 0 \end{bmatrix} \in \mathbb{K}[x]^{3 \times 3}.$$

Since this matrix has $m < 2n$, we then recurse along the first column of F_2 , with shift $\text{rdeg}_s(K_1) = (5, 8, 4)$. This gives a $(5, 8, 4)$ -weak Popov basis of the kernel of that column, as

$$K'_1 = \begin{bmatrix} x^3 & 1 & x^3 + 1 \\ 1 & 0 & x + 1 \end{bmatrix} \in \mathbb{K}[x]^{2 \times 3}.$$

This also provides the rank profile (1) of that column.

Multiplying K'_1 by the last two columns of F_2 gives a zero matrix, with the identity as the kernel basis. Hence $K_2 = K'_1$ is the sought $(5, 8, 4)$ -weak Popov basis of $\mathcal{K}(F_2)$, and the rank profile of F_2 is

$J_2 = (1)$. Then the latter is shifted to $J_2 = (1 + \lfloor 5/2 \rfloor) = (3)$, to keep track of the position of the column block F_2 in the input F .

Concatenating J_1 and J_2 yields the rank profile $(1, 2, 3)$ of F , and the product K_2K_1 is the kernel basis K given in Example 2.6. \square

4.2 Proof of correctness

In this subsection, we prove the correctness of Algorithm 1.

Cases $F = \mathbf{0}$ or $m = 1$. The correctness of Lines 1 to 5 is clear.

Case $2n > m$. Here the algorithm runs Lines 7 to 13 and returns. We assume correctness for the recursive calls at Lines 8 and 10. From $s \geq \text{rdeg}(F)$, we get $\text{rdeg}_s(K_1) \geq \text{rdeg}(K_1F) \geq \text{rdeg}(F_2)$ and the requirement of the call at Line 10 is satisfied. Lemma 2.5 implies that the matrix K_2K_1 is in s -weak Popov form. Furthermore,

$$K_2K_1F = K_2K_1[F_1 \quad F_{*,\{\lfloor n/2 \rfloor+1,\dots,n\}}] = K_2[\mathbf{0} \quad F_2] = \mathbf{0}.$$

To prove that K_2K_1 generates the kernel $\mathcal{K}(F)$, we let $p \in \mathcal{K}(F)$ and prove that $p = uK_2K_1$ for some $u \in \mathbb{K}[x]^{1 \times \ell_2}$. Since p is in $\mathcal{K}(F_1)$, and K_1 is a basis of the latter kernel, we have that $p = vK_1$ for some $v \in \mathbb{K}[x]^{1 \times \ell_1}$. By construction of F_2 , $pF = [\mathbf{0} \quad vF_2]$. Then $pF = \mathbf{0}$ implies $v \in \mathcal{K}(F_2)$ and, since K_2 is a basis of the latter kernel, we have $v = uK_2$ for some $u \in \mathbb{K}[x]^{1 \times \ell_2}$. This yields $p = uK_2K_1$. Thus K_2K_1 is an s -weak Popov basis of $\mathcal{K}(F)$.

In order to prove that (J_1, J_2) is the rank profile of F , the main observation is that since K_1 has full row rank, it can be completed into a nonsingular matrix $U = \begin{bmatrix} * \\ K_1 \end{bmatrix} \in \mathbb{K}[x]^{m \times m}$. Then,

$$UF = \begin{bmatrix} V & * \\ \mathbf{0} & F_2 \end{bmatrix} \text{ for some } V \in \mathbb{K}[x]^{r_1 \times \lfloor n/2 \rfloor}.$$

V has the same rank profile as $\begin{bmatrix} V \\ \mathbf{0} \end{bmatrix}$ and, since U is nonsingular, $UF_1 = \begin{bmatrix} V \\ \mathbf{0} \end{bmatrix}$ has the same rank profile as F_1 , which is J_1 . In particular, V has full row rank, and then the triangular form of UF implies that its rank profile is the concatenation of J_1 and of J_2 , the latter being the rank profile of F_2 shifted by adding $\lfloor n/2 \rfloor$ to all entries. Since UF and F have the same rank profile, F has rank profile (J_1, J_2) .

Case $n \leq m/2$. The algorithm runs Lines 15 to 29 and returns.

The basis A of $\mathcal{R}_M(F)$ at Line 19 is such that $AF = QM$ for some $Q \in \mathbb{K}[x]^{m \times n}$. Since both A and M are nonsingular, with M being upper triangular, implies that $Q = AFM^{-1}$ has the same rank profile as F . By the construction at Lines 20 to 24, we see that there is an $m \times m$ permutation matrix P such that

$$PQ = PAFM^{-1} = \begin{bmatrix} A_{I^c,*}FM^{-1} \\ A_{I,*}FM^{-1} \end{bmatrix} = \begin{bmatrix} G \\ \mathbf{0} \end{bmatrix}.$$

Thus G has the same rank profile as PQ , and hence the same rank profile as F . Thus, to conclude the proof for the rank profile, it suffices to verify that the recursive call at Line 27 computes the rank profile of G , which is true provided that t satisfies the requirements $t \geq \mathbf{0}$ and $t \geq \text{rdeg}(G)$. We prove this in the next paragraph.

Observe that the shift built at Line 16 satisfies $\hat{s} \geq \text{rdeg}(F)$. This implies $\text{rdeg}_{\hat{s}}(A_{I^c,*}) \geq \text{rdeg}(A_{I^c,*}F)$ and, defining $d = \text{cdeg}(M)$,

$$\begin{aligned} t &\geq \text{rdeg}(A_{I^c,*}F) - (\gamma, \dots, \gamma) = \text{rdeg}(GM) - (\gamma, \dots, \gamma) \\ &= \text{rdeg}_{(-\gamma, \dots, -\gamma)}(GM) \geq \text{rdeg}_{-d}(GM), \end{aligned}$$

where the last inequality comes from $(\gamma, \dots, \gamma) \leq d$. Now the fact that M is in Hermite form ensures that it is in $-d$ -reduced form with $\text{rdeg}_{-d}(M) = \mathbf{0}$, so that the predictable degree property yields

$\text{rdeg}_{-d}(GM) = \text{rdeg}(G)$. Thus $t \geq \text{rdeg}(G)$, and $t \geq \mathbf{0}$ follows since G has no zero row by construction.

This also ensures that K_2 is a t -weak Popov basis of $\mathcal{K}(G) = \mathcal{K}(A_{I^c,*}F)$. Lemmas 2.4 and 2.5 then imply that $K_2A_{I^c,*}$ is \hat{s} -weak Popov, with \hat{s} -pivot index a subset of that of $A_{I^c,*}$. Since the latter is disjoint from the \hat{s} -pivot index of $A_{I,*}$ and since $K_2A_{I^c,*}$ and $A_{I,*}$ are both \hat{s} -weak Popov, it follows that K is \hat{s} -weak Popov. Since s and \hat{s} only differ by a constant, K is s -weak Popov.

By construction, $A_{I,*}F = \mathbf{0}$ and $\mathbf{0} = K_2G = K_2A_{I^c,*}F$, and so $KF = \mathbf{0}$. It remains to prove that any $p \in \mathcal{K}(F)$ is a $\mathbb{K}[x]$ -linear combination of the rows of K . Since p is in $\mathcal{R}_M(F)$, we get $p = qA = q_I A_{I,*} + q_I^c A_{I^c,*}$ for some $q \in \mathbb{K}[x]^{1 \times m}$ and its subvectors q_I and q_I^c with indices in I and I^c , respectively. Then,

$$\mathbf{0} = pF = q_I A_{I,*}F + q_I^c A_{I^c,*}F = q_I^c A_{I^c,*}F,$$

which gives $q_I^c \in \mathcal{K}(A_{I^c,*}F) = \mathcal{K}(GM) = \mathcal{K}(G)$. Therefore $q_I^c = rK_2$ for some vector r , and we get $p = q_I A_{I,*} + rK_2A_{I^c,*}$.

Remark: As one can see above, M is required to be in Hermite normal form only for ensuring that the rank profile is not modified when right-multiplying by M . Hence, if one is only interested in a kernel basis, any column reduced matrix M will do.

4.3 Proof of complexity

The efficiency is based on three main ingredients. First, a fast algorithm for computing an s -weak Popov basis of $\mathcal{R}_M(F)$. Second, the fast multiplication of matrices which have unbalanced, but controlled, shifted row degrees. Third, the next lemma, which is a generalization and variant of [39, Thm. 3.6]: it states that the relation basis at Line 19 yields a substantial amount of kernel rows, effectively reducing the number of rows that remain to be found.

LEMMA 4.3. *Let $F \in \mathbb{K}[x]^{m \times n}$ have rank r , and let $s \in \mathbb{Z}_{\geq 0}^m$ such that $s \geq \text{rdeg}(F)$. Let $K \in \mathbb{K}[x]^{(m-r) \times m}$ be an s -reduced basis of $\mathcal{K}(F)$. For any $k > 0$, at most $\lfloor k \rfloor$ rows of K have s -degree more than or equal to $\tau = \lceil |s|/k \rceil$. Then, let $M \in \mathbb{K}[x]^{n \times n}$ be column reduced with $\min(\text{cdeg}(M)) \geq \tau$. For any s -reduced basis $A \in \mathbb{K}[x]^{m \times m}$ of $\mathcal{R}_M(F)$, at most $r + \lfloor k \rfloor$ rows of A are not in $\mathcal{K}(F)$.*

PROOF. If ρ is the number of rows of K whose s -degree is $\geq \tau$, then $|\text{rdeg}_s(K)| \geq \rho\tau$. Thus, from the bound $|\text{rdeg}_s(K)| \leq |s| \leq k\tau$ (see Item (iii) of Theorem 3.3), we get $\rho \leq k$. It follows that there are $\sigma = m - r - \rho \geq m - r - k$ rows of K whose s -degree is $< \tau$.

We claim that, as a consequence, there are at least σ rows of A which are in $\mathcal{K}(F)$. First, since all rows of K are also in $\mathcal{R}_M(F)$, by minimality of the s -degree of A , there are at least σ rows of A which have s -degree $< \tau$. The claim then follows from the fact that any $p \in \mathcal{R}_M(F)$ such that $\text{rdeg}_s(p) < \tau$ is in $\mathcal{K}(F)$. Indeed, $pF = qM$ for some $q \in \mathbb{K}[x]^{1 \times n}$. On the one hand, $\text{rdeg}(pF) \leq \text{rdeg}_s(p) < \tau$, since $s \geq \text{rdeg}(F)$. On the other hand, M is column reduced with $d = \text{cdeg}(M)$, and thus it is also $-d$ -reduced with $\text{rdeg}_{-d}(M) = \mathbf{0}$. Hence, assuming $pF \neq \mathbf{0}$ is nonzero (and therefore $q \neq \mathbf{0}$), and using $\tau \leq \min(d)$ as well as the predictable degree property, we obtain

$$\begin{aligned} \tau &> \text{rdeg}(pF) = \text{rdeg}(qM) = \tau + \text{rdeg}_{(-\tau, \dots, -\tau)}(qM) \\ &\geq \tau + \text{rdeg}_{-d}(qM) = \tau + \text{rdeg}(q) \geq \tau. \end{aligned}$$

This is a contradiction, hence $pF = \mathbf{0}$, i.e. $p \in \mathcal{K}(F)$.

In conclusion, A has at most $m - \sigma \leq r + k$ rows not in $\mathcal{K}(F)$. \square

Note that if we take $k \leq \frac{m-r}{2}$, then this number is $r+k \leq \frac{m+r}{2}$. For example, if $n < m$ and no information on r is known, one can take $k = \frac{m-n}{2}$, and then $r+k = \frac{m+n}{2}$. This is the choice made in Algorithm 1, which leads to $\tau = \lceil 2|s|/(m-n) \rceil$. Furthermore, since $2n \leq m$ in that algorithm, we obtain $r+k = \frac{m+n}{2} \leq \frac{3m}{4}$.

COROLLARY 4.4. *At Line 23 of Algorithm 1, the matrix \mathbf{G} has at most $\frac{3m}{4}$ rows.*

This lemma ensures that, when the algorithm enters Lines 15 to 29, then the number of rows becomes at most $3m/4$ in the recursive call (and the number of columns is unchanged). On the other hand, when the algorithm enters Lines 7 to 13, then the number of columns becomes at most $\lceil n/2 \rceil$ in each of the two recursive calls (and the number of rows remains bounded from above by m).

Note that we have proved in Section 4.2 that in each recursive call, the input shift is an upper bound on the row degrees of the input matrix. Now, we observe further that each of these shifts has a sum of entries at most $|s|$. This is clear at Line 8 which uses the input shift \mathbf{s} , and at Line 10 since the shift satisfies $|\text{rdeg}_s(\mathbf{K}_1)| \leq |s|$ according to Item (iii) of Theorem 3.3. Now, at Lines 26 and 27, the shift \mathbf{t} has entries at most those of the subtuple $\mathbf{s}_{\mathcal{I}_c}$, since

$$\text{rdeg}_s(\mathbf{A}) - (\gamma, \dots, \gamma) = \mathbf{s} + \boldsymbol{\delta} - (\gamma + \mu, \dots, \gamma + \mu)$$

where $\boldsymbol{\delta}$ is the \mathbf{s} -pivot degree of \mathbf{A} , with $\boldsymbol{\delta} \leq (\gamma, \dots, \gamma)$ under our assumption $\text{cdeg}(\mathbf{M}) = (\gamma, \dots, \gamma)$.

Recall the notation $D = |s|$, and note that $m \lceil D/m \rceil = \Theta(m+D)$.

Based on [39, Thm. 3.7], one can verify that the matrix products at Lines 9, 13, 23 and 28, use $O(m^{\omega-2}(m+n)(m+D))$ operations in \mathbb{K} . Note that the right-multiplication by \mathbf{M}^{-1} at Line 23 is only a matter of univariate polynomial exact division: the matrix $\mathbf{A}_{\mathcal{I}_c, *}$ \mathbf{F} is known to be a left multiple of \mathbf{M} by construction, and the latter matrix is diagonal by assumption.

It remains to observe that the computation of \mathbf{A} at Line 19 costs $O(m^{\omega-1}n\tau)$ operations in \mathbb{K} , since \mathbf{M} is diagonal [24, Thm. 1.4]. Since we are in the case $n \leq m/2$, we have $m-n \geq m/2$ and thus

$$\tau = \left\lceil \frac{2|s|}{m-n} \right\rceil \in O\left(1 + \frac{D}{m-n}\right) \subseteq O\left(1 + \frac{D}{m}\right).$$

Hence the above cost for computing \mathbf{A} is in $O(m^{\omega-2}n(m+D))$.

We conclude that all computations apart from recursive calls use a total of $O(m^{\omega-2}(m+n)(m+D))$ operations in \mathbb{K} , leading to the cost bound announced in Theorem 4.1.

5 FINDING THE COLUMN RANK PROFILE AND LINEARLY INDEPENDENT ROWS

Let \mathbf{F} be a polynomial matrix of rank r . This section presents a rank-sensitive algorithm to find both the rank profile of \mathbf{F} and a set of r rows of \mathbf{F} which are $\mathbb{K}[x]$ -linearly independent. In particular, this information locates an $r \times r$ nonsingular submatrix of \mathbf{F} .

5.1 Algorithm

The idea is to maintain a subset U of the top rows of \mathbf{F} , which are known to have full rank, and to incorporate new rows from the bottom part of \mathbf{F} . Precisely, U locates k rows with index in $(1, \dots, \theta-1)$, and the next step finds a set of rows of maximal rank in the matrix \mathbf{G} formed by joining these k rows $\mathbf{F}_{U, *}$ with the k rows with indices $(\theta, \dots, \theta+k-1)$ of \mathbf{F} (or only up to m if $\theta+k-1 \geq m$).

Algorithm 2 COLUMNRANKPROFILE(\mathbf{F}, θ, U)

Input: a matrix $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$, an integer $\theta \in \{1, \dots, m+1\}$, a list $U \subseteq \{1, \dots, \theta-1\}$ of size $k \geq 0$

Requirement: $k = 0$ or $\text{rank}(\mathbf{F}_{U, *}) = \text{rank}(\mathbf{F}_{1.., \theta-1, *}) = k$

Output: lists $I \subseteq \{1, \dots, m\}$ and $J \subseteq \{1, \dots, n\}$, both of size $r = \text{rank}(\mathbf{F})$, such that $\mathbf{F}_{I, J} \in \mathbb{K}[x]^{r \times r}$ is nonsingular and J is the rank profile of \mathbf{F}

```

1: if  $k = n$  then return  $U, (1, \dots, n)$ 
2: if  $k = 0$  then
3:    $i \leftarrow$  index of the first nonzero row of  $\mathbf{F}$ 
4:   return COLUMNRANKPROFILE( $\mathbf{F}, i+1, (i)$ )
5:  $\triangleright k > 0$  independent rows are known among the rows  $1, \dots, \theta-1$ 
6:  $\ell \leftarrow \min(k, m-\theta+1)$   $\triangleright$  now incorporate rows  $\theta, \dots, \theta+\ell-1$ 
7:  $V \leftarrow U \cup (\theta, \theta+1, \dots, \theta+\ell-1)$ 
8:  $\mathbf{G} \leftarrow \mathbf{F}_{V, *}, \in \mathbb{K}[x]^{(k+\ell) \times n}$ ;  $\mathbf{s} \leftarrow \text{rdeg}(\mathbf{G})$ 
9:  $\mathbf{K}, J \leftarrow \text{KERNELBASIS-RANKPROFILE}(\mathbf{G}, \mathbf{s})$ 
10:  $(\boldsymbol{\pi}, \boldsymbol{\delta}) \leftarrow \mathbf{s}$ -pivot profile of  $\mathbf{K}$ ;  $\boldsymbol{\pi}^c \leftarrow \{1, \dots, k+\ell\} \setminus \boldsymbol{\pi}$ 
11:  $U' \leftarrow \emptyset$ 
12: for  $i \in \boldsymbol{\pi}^c$  do
13:   if  $i \leq k$  then add the  $i$ th element of  $U$  to  $U'$ 
14:   else add  $\theta+i-k-1$  to  $U'$ 
15:  $\theta' \leftarrow \theta + \ell$ 
16: if  $\theta' > m$  then return  $U', J$ 
17: return COLUMNRANKPROFILE( $\mathbf{F}, \theta', U'$ )

```

Finding a set of rows of maximal rank of \mathbf{G} is done efficiently via Algorithm 1 and the property in Item (i) of Theorem 3.3, which locates independent rows from the \mathbf{s} -pivot index of the kernel basis. Since the call to Algorithm 1 also provides the column rank profile of \mathbf{G} , we eventually obtain I and J identifying a nonsingular submatrix of \mathbf{F} of size $r \times r$, with J the rank profile of $\mathbf{F}_{I, *}$. By Lemma 3.2, the latter is also the rank profile of \mathbf{F} .

Starting with $k = 1$ and U locating the first nonzero row of \mathbf{F} , this leads to a rank-sensitive algorithm, which at any stage considers a submatrix of \mathbf{F} with n columns and at most $2k \leq 2r$ rows.

THEOREM 5.1. *Let $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$ have rank r . Assume that $\text{rdeg}(\mathbf{F})$ is nondecreasing, and that (θ, U) satisfies the input requirements. Then COLUMNRANKPROFILE(\mathbf{F}, θ, U) uses $O(r^{\omega-2}n(m+D))$ operations in \mathbb{K} , where D is the sum of the nonnegative entries of $\text{rdeg}(\mathbf{F})$. It returns lists $I \subseteq \{1, \dots, m\}$ and $J \subseteq \{1, \dots, n\}$, both of size r , such that $\mathbf{F}_{I, J} \in \mathbb{K}[x]^{r \times r}$ is nonsingular and J is the rank profile of \mathbf{F} .*

Before proving the theorem we note that, if nothing particular is known about \mathbf{F} a priori, one can call this algorithm with $\theta = 1$ and $k = 0$ (meaning $U = \emptyset$). One can also permute the rows of \mathbf{F} to ensure that its row degrees are nondecreasing.

COROLLARY 5.2. *Given $\mathbf{F} \in \mathbb{K}[x]^{m \times n}$, one can locate an $r \times r$ nonsingular submatrix $\mathbf{F}_{I, J}$ of \mathbf{F} using $O(r^{\omega-2}n(m+D))$ operations in \mathbb{K} , where r is the rank of \mathbf{F} , J is the column rank profile of $\mathbf{F}_{I, *}$, and D is the sum of the degrees of the nonzero rows of \mathbf{F} .*

5.2 Proof of correctness

If $k = n$, the requirement $\text{rank}(\mathbf{F}_{U, *}) = n$ implies that the $n \times n$ matrix $\mathbf{F}_{U, 1..n}$ is nonsingular, proving the correctness of Line 1.

If $k = 0$, then the correctness Lines 2 to 4 follows from the fact the input requirements are satisfied for $\theta = i+1$ and $U = (i)$.

The integer ℓ at Line 6 is such that $0 \leq \ell \leq k$ and $\theta + \ell - 1 \leq m$. Then, at Line 7, the list V contains $k + \ell$ distinct indices, with the first k in U and the others in $\{\theta, \dots, m\}$. As a result, the matrix $\mathbf{G} = \mathbf{F}_{V,*}$ at Line 8 has rank between k and $k + \ell$. By Item (i) of Theorem 3.3, the list π^c computed at Lines 9 and 10 identifies $\text{rank}(\mathbf{G})$ rows of \mathbf{G} which are $\mathbb{K}[x]$ -linearly independent.

These rows provide a set of rows of \mathbf{F} which have maximal rank among its first $\theta + \ell - 1$ rows. The role of Lines 11 to 14 is simply to link the row indices in \mathbf{G} , as they appear in π^c , to the corresponding row indices in \mathbf{F} . This therefore provides U' such that $U' \subseteq \{1, \dots, \theta + \ell - 1\}$, whose size k' is between k and $k + \ell$, and such that $\mathbf{F}_{U',*}$ has full row rank with $\text{rank}(\mathbf{F}_{U',*}) = \text{rank}(\mathbf{F}_{1.., \theta + \ell - 1, *})$.

Then Line 15 updates θ to θ' , to reflect that we have now covered all rows with indices in $\{1, \dots, \theta + \ell - 1\}$, and that we can proceed with the remaining rows starting at index $\theta' = \theta + \ell$.

In the case where $\theta' = \theta + \ell > m$ (Line 16), all rows of \mathbf{F} have been processed and the algorithm can return U' and J . The correctness concerning U' has been discussed above, and the fact that J is the rank profile of $\mathbf{F}_{U',*}$ follows from Theorem 4.1.

Otherwise, if $\theta' \leq m$, we proceed with the remaining bottom part of \mathbf{F} recursively (Line 17). The properties of U' described above show that the input requirement are satisfied in this recursive call.

Thus, for correctness, it only remains to observe the algorithm terminates since all recursive calls involve a strictly larger set U (whose size is bounded by the rank r of \mathbf{F}), or a strictly larger index θ (which is bounded by $m + 1$).

Remark: the independent rows of \mathbf{F} found via the s-pivot profile of \mathbf{K} do not necessarily contain the k independent rows that were already identified by U ; in other words, U is not necessarily contained in U' . Having $U \subseteq U'$ would have guaranteed that I is the row rank profile of \mathbf{F} , yet the straightforward modification of this algorithm that would ensure $U \subseteq U'$ involves a different choice of shift s which would make Line 9 too costly.

5.3 Proof of complexity

The only costly operation performed in Algorithm 2, apart from recursive calls, is the computation of the kernel basis and rank profile of the matrix \mathbf{G} , via Algorithm 1 at Line 9. The main task for the complexity analysis is therefore to analyze how the dimensions and row degrees of \mathbf{G} evolve during the run of the algorithm.

Fix some input (\mathbf{F}, θ, U) , and assume the cardinality k of U is nonzero; otherwise, we are brought to this situation by Line 4 after at most mn zero tests performed by Line 3. For this input, let ρ be the number of recursive steps before arriving at a base case (either Line 1 or Line 16). Let $(\theta_0, U_0) = (\theta, U)$ be the original input and $(\theta_1, U_1), \dots, (\theta_\rho, U_\rho)$ be the input of the successive recursive calls when running the algorithm on (\mathbf{F}, θ, U) . Let also k_i be the cardinality of U_i for $1 \leq i \leq \rho$. Observe that $\theta_0 < \theta_1 < \dots < \theta_\rho$ and $k_0 \leq k_1 \leq \dots \leq k_\rho$, but recall from the remark in Section 5.2 that U_i is not necessarily a subset of U_{i+1} .

Let $\ell_i = \min(k_i, m - \theta_i + 1)$ as in Line 6, and \mathbf{G}_i be the matrix built at Line 8, which has $k_i + \ell_i$ rows. By Theorem 4.1, Line 9 costs

$$\tilde{O}\left((k_i + \ell_i)^{\omega-2}(k_i + \ell_i + n)(k_i + \ell_i + D_i)\right),$$

where $D_i = |\text{rdeg}(\mathbf{G}_i)|$. Now, since $k_i = \text{rank}(\mathbf{F}_{U_i,*}) \leq r \leq n$ and $k_i + \ell_i \leq 2k_i$, the above cost bound is within $O^*(r^{\omega-2}n(k_i + D_i))$. The

rest of this section shows that $\sum_{0 \leq i \leq \rho} k_i + D_i$ is in $O(m + D \log(r))$, which proves the complexity bound $O^*(r^{\omega-2}n(m + D))$.

Due to Line 15, $\theta_{i+1} = \theta_i + \ell_i$ for $0 \leq i < \rho$. Observe that $\ell_i = k_i$ for $i < \rho$; otherwise $\ell_i = m - \theta_i + 1$ and $\theta_{i+1} = \theta_i + \ell_i = m + 1$, hence the algorithm would stop at Line 16 before the ρ -th recursive call. It follows that $\theta_i = \theta_0 + \sum_{0 \leq j < i} k_j$ for $0 \leq i \leq \rho$. Line 16 ensures that $\theta_\rho \leq m$ in the last recursive call, hence $k_0 + \dots + k_{\rho-1} = \theta_\rho - \theta_0 \leq m$. We finally deduce $k_0 + \dots + k_\rho \leq m + k_\rho \leq m + r \leq 2m$.

It remains to prove $D_0 + \dots + D_\rho \in O(D \log(r))$. By construction, \mathbf{G}_i consists of $k_i + \ell_i$ rows among the first $\{\theta_i, \dots, \theta_i + \ell_i - 1\}$ rows of \mathbf{F} . Consequently, since $\text{rdeg}(\mathbf{F})$ is nondecreasing, $D_i \leq |\text{rdeg}(\mathbf{F}_{S_i,*})|$ where $S_i = \{\theta_i - k_i, \theta_i - k_i + 1, \dots, \theta_i + \ell_i - 1\}$. We claim that a given row $j \in \{1, \dots, m\}$ of \mathbf{F} may appear in at most $\lfloor \log_2(r) \rfloor + 2$ sets among S_0, \dots, S_ρ , that is, $\#\{0 \leq i \leq \rho \mid j \in S_i\} \leq \lfloor \log_2(r) \rfloor + 2$. Below we prove this claim, which concludes the proof since then

$$\begin{aligned} \sum_{0 \leq i \leq \rho} D_i &\leq \sum_{0 \leq i \leq \rho} |\text{rdeg}(\mathbf{F}_{S_i,*})| = \sum_{0 \leq i \leq \rho} \sum_{j \in S_i} \text{rdeg}(\mathbf{F}_{j,*}) \\ &= \sum_{1 \leq j \leq m} \sum_{\substack{0 \leq i \leq \rho \\ j \in S_i}} \text{rdeg}(\mathbf{F}_{j,*}) \leq D(\lfloor \log_2(r) \rfloor + 2). \end{aligned}$$

Since $(\min(S_i))_i$ and $(\max(S_i))_i$ are nondecreasing, having $j \in S_{i_1} \cap S_{i_2}$ for some $i_1 < i_2$ implies $j \in S_i$ for all $i_1 \leq i \leq i_2$. So we consider $j \in S_{i_1}, \dots, S_{i_1+c-1}$ for some $c > 0$ and $0 \leq i_1 \leq \rho - c + 1$.

The fact that $j \in S_i$ implies $j \leq \theta_i + \ell_i - 1 = \theta_{i+1} - 1$. On the other hand, $j \in S_{i+\gamma}$ for $0 \leq \gamma \leq c - 1$ implies $j \geq \theta_{i+\gamma} - k_{i+\gamma}$. We deduce $k_{i+\gamma} \geq \theta_{i+\gamma} - \theta_{i+1} + 1 = k_{i+\gamma-1} + \dots + k_{i+2} + k_{i+1} + 1$. Starting from $k_{i+1} \geq 1$, using this inequality iteratively for $\gamma = 2, \dots, c - 1$ shows that $k_{i+c-1} \geq 2^{c-2}$. Since $k_{i+c-1} \leq r$, we get $c \leq \lfloor \log_2(r) \rfloor + 2$.

6 TOPICS FOR FURTHER RESEARCH

Our algorithm can find the rank profile with a rank-sensitive cost. However the same cannot be said for such computations as kernel basis, column basis, and approximant/order bases; or for computing normal forms such as Hermite or Popov. We would like to make progress on filling this gap thanks to the new results in this paper. In addition, we are interested in applying our work in situations (including those mentioned in the introduction) where rank-sensitive algorithms would allow one to tackle significantly larger problems.

Another feature of our algorithms is that the complexity depends on the average row degree of the input matrix. However, they do not handle unbalanced column degrees, where the matrix might have average row degree close to the global degree but average column degree quite smaller. We would like to improve the cost towards the minimum of the average of both row and column degrees, or even on a notion of generic determinant bound [10, Sec. 6] generalized to rectangular matrices. We believe that *partial linearization* [30, Sec. 3] [10, Sec. 6] may lead to such an improvement.

Finally, while it has been popular in recent times to give $O^*(\cdot)$ complexities which hide log terms, there remains a strong interest in the more precise $O(\cdot)$ measure. In fact this is often the first audience question asked when an algorithm is presented with $O^*(\cdot)$ complexity. We would like to determine the logarithmic terms for the algorithms presented in this paper. Although technical, this seems feasible since the logarithmic factors in the complexity of the core tools are now well understood, specifically approximant bases [9, 16, 36] and multiplication with unbalanced degrees [15, 39].

Figure 1: SageMath code for Examples 2.6, 3.4 and 4.2. This code is written using SageMath (version 9.3 or later required) and illustrates many of the points in the three listed examples: running this code will show the matrices and some additional information. This code can also be easily adapted to make related experiments.

```

# For a detailed documentation of functionalities for univariate polynomial matrices, including
# the minimal_kernel_basis and minimal_approximant_basis methods used below, see
# https://doc.sagemath.org/html/en/reference/matrices/sage/matrix/matrix_polynomial_dense.html
# (the code below requires SageMath >=9.3; some other functionalities require SageMath 9.4 or 9.5)

pR.<x> = GF(2)[]
F = Matrix(pR, 5, 5, \
    [[x^2,x^3+1,x^8+x^6+x^4+x^3+x^2+x,x^4+1,x^3+1], \
    [0,x^4+1,x^5+x^4+x^3+x^2,x+1,x^2+1], \
    [0,x^2+1,x+1,0,1], \
    [0,0,x^8+1,x^4+1,0], \
    [0,0,x^4+1,1,0]])

print(f"Input matrix F:\n{F}\n")

K0 = F.minimal_kernel_basis()
print(f"Minimal kernel basis with shift s=0\n{K0}")
piv = [pi+1 for pi in K0.leading_positions()]
print(f"Its pivot indices are {piv}\n")

s = F.row_degrees()
K = F.minimal_kernel_basis(shifts=s)
print(f"Minimal kernel basis K with shift s=rdeg(F)\n{K}")
piv_s = [pi+1 for pi in K.leading_positions(shifts=s)]
print(f"Its s-pivot indices are {piv_s}\n")

F1 = F[:, :2] ; tau = 18
A = F1.minimal_approximant_basis(tau, shifts=s)
print(f"Approximant basis of first 2 columns at order {tau}:\n{A}")
# set I at Line 14 of Algorithm 1:
I = [i for i in range(5) if A[i, :].row_degrees(shifts=s)[0] < tau]
# --> gives 3 indices I == [1,3,4] so we have the whole kernel basis
# since here we do know rank(F[:, :2]) = 2 hence kernel rank 5-2==3
print(f"--> indices of rows in kernel: {[i+1 for i in I]}\n")
K1 = A[I, :]; t = K1.row_degrees(shifts=s)

F2 = K1 * F[:, :2]
print(f"Residual matrix F2 for second call:\n{F2}")
K2 = F2[:, 0].minimal_kernel_basis(shifts=t)

print(f"Kernel basis K1' of first column of F2:\n{K2}\n")

print(f"Test K1' * F2 == 0 --> {K2*F2 == 0}, so in fact K2 = K1'")

print(f"Verify K2*K1 is the above s-weak Popov matrix K --> {K2*K1 == K}")

```

REFERENCES

- [1] B. Beckermann. 1992. A reliable method for computing M-Padé approximants on arbitrary staircases. *J. Comput. Appl. Math.* 40, 1 (1992), 19–42. [https://doi.org/10.1016/0377-0427\(92\)90039-Z](https://doi.org/10.1016/0377-0427(92)90039-Z)
- [2] B. Beckermann and G. Labahn. 1994. A Uniform Approach for the Fast Computation of Matrix-Type Padé Approximants. *SIAM J. Matrix Anal. Appl.* 15, 3 (July 1994), 804–823. <https://doi.org/10.1137/S0895479892230031>
- [3] B. Beckermann, G. Labahn, and G. Villard. 1999. Shifted Normal Forms of Polynomial Matrices. In *Proceedings ISSAC 1999*. ACM, 189–196. <https://doi.org/10.1145/309831.309929>
- [4] B. Beckermann, G. Labahn, and G. Villard. 2006. Normal forms for general polynomial matrices. *J. Symbolic Comput.* 41, 6 (2006), 708–737. <https://doi.org/10.1016/j.jsc.2006.02.001>

- [5] H. Y. Cheung, T. C. Kwok, and L. C. Lau. 2013. Fast Matrix Rank Algorithms and Applications. *J. ACM* 60, 5, Article 31 (2013). <https://doi.org/10.1145/2528404>
- [6] J.-G. Dumas, C. Pernet, and Z. Sultan. 2017. Fast computation of the rank profile matrix and the generalized Bruhat decomposition. *J. Symbolic Comput.* 83 (2017), 187–210. <https://doi.org/10.1016/j.jsc.2016.11.011>
- [7] G. D. Forney, Jr. 1975. Minimal Bases of Rational Vector Spaces, with Applications to Multivariable Linear Systems. *SIAM Journal on Control* 13, 3 (1975), 493–520. <https://doi.org/10.1137/0313029>
- [8] K. O. Geddes, S. R. Czapor, and G. Labahn. 1992. *Algorithms for computer algebra*. Kluwer, Boston.
- [9] P. Giorgi, C.-P. Jeannerod, and G. Villard. 2003. On the complexity of polynomial matrix computations. In *Proceedings ISSAC 2003*. ACM, 135–142. <https://doi.org/10.1145/860854.860889>
- [10] S. Gupta, S. Sarkar, A. Storjohann, and J. Valeriate. 2012. Triangular x -basis decompositions and derandomization of linear algebra algorithms over $K[x]$. *J. Symbolic Comput.* 47, 4 (2012), 422–453. <https://doi.org/10.1016/j.jsc.2011.09.006>
- [11] C. Hermite. 1851. Sur l'introduction des variables continues dans la théorie des nombres. *Journal für die reine und angewandte Mathematik* 41 (1851), 191–216. <https://doi.org/10.1515/crll.1851.41.191>
- [12] S. G. Hyun, V. Neiger, and É. Schost. 2019. Implementations of efficient univariate polynomialmatrix algorithms and application to bivariate resultants. In *Proceedings ISSAC 2019*. ACM, 235–242.
- [13] S. G. Hyun, V. Neiger, and É. Schost. 2021. Algorithms for Linearly Recurrent Sequences of Truncated Polynomials. In *Proceedings ISSAC 2021*. ACM, 201–208. <https://doi.org/10.1145/3452143.3465533>
- [14] C.-P. Jeannerod. 2006. LSP matrix decomposition revisited. Research report 2006-28. Inria – LIP – Ens de Lyon. <http://www.ens-lyon.fr/LIP/Pub/Rapports/RR/RR2006/RR2006-28.pdf>
- [15] C.-P. Jeannerod, V. Neiger, É. Schost, and G. Villard. 2017. Computing minimal interpolation bases. *J. Symbolic Comput.* 83 (2017), 272–314. <https://doi.org/10.1016/j.jsc.2016.11.015>
- [16] C.-P. Jeannerod, V. Neiger, and G. Villard. 2020. Fast computation of approximant bases in canonical form. *J. Symbolic Comput.* 98 (2020), 192–224. <https://doi.org/10.1016/j.jsc.2019.07.011>
- [17] C.-P. Jeannerod, C. Pernet, and A. Storjohann. 2013. Rank-profile revealing Gaussian elimination and the CUP matrix decomposition. *J. Symbolic Comput.* 56 (2013), 46–68. <https://doi.org/10.1016/j.jsc.2013.04.004>
- [18] T. Kailath. 1980. *Linear Systems*. Prentice-Hall.
- [19] G. Labahn, V. Neiger, and W. Zhou. 2017. Fast, deterministic computation of the Hermite normal form and determinant of a polynomial matrix. *J. Complexity* 42C (2017), 44–71.
- [20] D. Lucas, V. Neiger, C. Pernet, D. S. Roche, and J. Rosenkilde. 2021. Verification protocols with sub-linear communication for polynomial matrix operations. *J. Symbolic Comput.* 105 (2021), 165–198. <https://doi.org/10.1016/j.jsc.2020.06.006>
- [21] C. C. MacDuffee. 1933. *The Theory of Matrices*. Springer-Verlag Berlin Heidelberg. <https://doi.org/10.1007/978-3-642-99234-6>
- [22] T. Mulders and A. Storjohann. 2003. On lattice reduction for polynomial matrices. *J. Symbolic Comput.* 35, 4 (2003), 377–401. [https://doi.org/10.1016/S0747-7171\(02\)00139-6](https://doi.org/10.1016/S0747-7171(02)00139-6)
- [23] V. Neiger. 2016. *Bases of relations in one or several variables: fast algorithms and applications*. Ph. D. Dissertation. École Normale Supérieure de Lyon. <https://tel.archives-ouvertes.fr/tel-01431413>
- [24] V. Neiger. 2016. Fast computation of shifted Popov forms of polynomial matrices via systems of modular polynomial equations. In *Proceedings ISSAC 2016*. ACM, 365–372. <https://doi.org/10.1145/2930889.2930936>
- [25] V. Neiger and C. Pernet. 2021. Deterministic computation of the characteristic polynomial in the time of matrix multiplication. *J. Complexity* 67 (2021), 101572. <https://doi.org/10.1016/j.jco.2021.101572>
- [26] V. Neiger, J. Rosenkilde, and G. Solomatov. 2018. Computing Popov and Hermite Forms of Rectangular Polynomial Matrices. In *Proceedings ISSAC 2018*. ACM, 295–302. <https://doi.org/10.1145/3208976.3208988>
- [27] V. Neiger and T. X. Vu. 2017. Computing canonical bases of modules of univariate relations. In *Proceedings ISSAC 2017*. ACM, 357–364. <https://doi.org/10.1145/3087604.3087656>
- [28] M. Newman. 1972. *Integral Matrices*. Academic Press.
- [29] A. Storjohann. 2000. *Algorithms for Matrix Canonical Forms*. Ph. D. Dissertation. Swiss Federal Institute of Technology – ETH.
- [30] A. Storjohann. 2006. Notes on computing minimal approximant bases. In *Challenges in Symbolic Computation Software (Dagstuhl Seminar Proceedings)*. <http://drops.dagstuhl.de/opus/volltexte/2006/776>
- [31] A. Storjohann and T. Mulders. 1998. Fast Algorithms for Linear Algebra Modulo N . In *Proceedings Algorithms – ESA' 98*. Springer, 139–150. https://doi.org/10.1007/3-540-68530-8_12
- [32] A. Storjohann and G. Villard. 2005. Computing the Rank and a Small Nullspace Basis of a Polynomial Matrix. In *Proceedings ISSAC 2005*. ACM, 309–316. <https://doi.org/10.1145/1073884.1073927>
- [33] A. Storjohann and S. Yang. 2015. A Relaxed Algorithm for Online Matrix Inversion. In *Proceedings ISSAC 2015*. ACM, 339–346. <https://doi.org/10.1145/2755996.2756672>
- [34] M. Van Barel and A. Bultheel. 1992. A general module theoretic framework for vector M -Padé and matrix rational interpolation. *Numer. Algorithms* 3 (1992), 451–462. <https://doi.org/10.1007/BF02141952>
- [35] W. Zhou. 2012. *Fast Order Basis and Kernel Basis Computation and Related Problems*. Ph. D. Dissertation. University of Waterloo.
- [36] W. Zhou and G. Labahn. 2012. Efficient Algorithms for Order Basis Computation. *J. Symbolic Comput.* 47, 7 (2012), 793–819. <https://doi.org/10.1016/j.jsc.2011.12.009>
- [37] W. Zhou and G. Labahn. 2013. Computing Column Bases of Polynomial Matrices. In *Proceedings ISSAC 2013*. ACM, 379–386. <https://doi.org/10.1145/2465506.2465947>
- [38] W. Zhou and G. Labahn. 2014. Unimodular Completion of Polynomial Matrices. In *Proceedings ISSAC 2014*. ACM, 413–420. <https://doi.org/10.1145/2608628.2608640>
- [39] W. Zhou, G. Labahn, and A. Storjohann. 2012. Computing Minimal Nullspace Bases. In *Proceedings ISSAC 2012*. ACM, 366–373. <https://doi.org/10.1145/2442829.2442881>