Popov Form Computation for Matrices of Ore Polynomials

Mohamed Khochtali
University of Waterloo
Canada
mkhochta@uwaterloo.ca

Johan Rosenkilde, né Nielsen
Technical University of Denmark
Denmark
jsrn@jsrn.dk

Arne Storjohann
University of Waterloo
Canada
astorjoh@uwaterloo.ca

February 4, 2017

Abstract

Let $\mathsf{F}[\partial;\sigma,\delta]$ be a ring of Ore polynomials over a field. We give a new deterministic algorithm for computing the Popov form P of a nonsingular matrix $A \in \mathsf{F}[\partial;\sigma,\delta]^{n\times n}$. Our main focus is to ensure controlled growth in the size of coefficients from F in the case $\mathsf{F} = \mathsf{k}(z)$, and even $\mathsf{k} = \mathbb{Q}$. Our algorithms are based on constructing from A a linear system over F and performing a structured fraction-free Gaussian elimination. The algorithm is output sensitive, with a cost that depends on the orthogonality defect of the input matrix: the sum of the row degrees in A minus the sum of the row degrees in P. The resulting bit-complexity for the differential and shift polynomial case over $\mathbb{Q}(z)$ improves upon the previous best.

1 Introduction

Ore polynomial rings, also known as skew polynomial rings, are non-commutative generalisations of univariate polynomial rings, introduced by Øystein Ore [18]. They have a variety of applications, such as modelling recurrence relations and differential equations [18]. Row spaces of matrices over Ore polynomial rings arise in studying coupled systems of such equations. Computing normal forms of such matrices allow comparing systems and finding small or otherwise special elements in the spaces. In coding theory, Ore polynomials have been used to construct codes for protecting against a powerful notion of errors, rank-errors [8], and matrices over Ore polynomials arise in the decoding of such codes [21].

In this paper, we consider the computation of the Popov normal form of a non-singular matrix over an Ore polynomial ring (see the formal definition

of Popov form in Section 2). Our focus is when the base field F is infinite so coefficient growth is a concern, in particular $\mathsf{F} = \mathsf{k}(z)$ where k is some field, possibly also with coefficient growth in mind, for example $\mathsf{k} = \mathbb{Q}$. This is an important and natural setting for most of the aforementioned applications. In coding theory, the base field is usually a finite field, and faster methods than what we present are possible. However, recently interest has arisen in this area for when F is an algebraic number field [1].

An Ore polynomial ring is given by a base field F, an automorphism σ of F, and a "derivation" of σ : this is a map $\delta : \mathsf{F} \mapsto \mathsf{F}$ satisfying

$$\begin{aligned} \delta(a+b) &= \delta(a) + \delta(b) \\ \delta(ab) &= \sigma(a)\delta(b) + \delta(a)b \ . \end{aligned}$$

The Ore ring $\mathsf{F}[\partial; \sigma, \delta]$ is then given as the set of finite formal sums $a_0 + a_1\partial + \ldots + a_d\partial^d$, with $a_i \in \mathsf{F}$ and $a_d \neq 0$. Addition of two Ore polynomials is the usual element-wise addition, while multiplication is given from the following non-commutative rule for multiplying an $a \in \mathsf{F}$ with ∂ on the right:

$$\partial a = \sigma(a)\partial + \delta(a) \ .$$

We mention two particularly important examples of Ore polynomial rings where $\mathsf{F} = \mathsf{k}(z)$:

- *Differential polynomials* where $\sigma(z) = z$ and $\delta(f(z)) = f'(z)$ is the usual derivative with respect to z.
- The *shift case*, or time-dependence, where $\sigma(f(z)) = f(z+1)$ is the shift automorphism and $\delta = 0$.

An Ore ring $\mathsf{F}[\partial; \sigma, \delta]$ is both a left and right Euclidean rings, with division algorithms which essentially works as for usual univariate polynomials. Ore rings also admit unique left skew field of fractions. These facts mean that matrices over Ore rings behave mostly as we are used to: the notion of *rank* and *(non)-singularity* makes sense, in particular, and performing row or column operations on a matrix will not change its rank. Further, two matrices $M, M' \in \mathsf{F}[\partial; \sigma, \delta]^{n \times m}$ generate the same left row space if and only if there exists $U \in \mathrm{GL}_n(\mathsf{F}[\partial; \sigma, \delta])$ such that $M = UM'$, where $\mathrm{GL}_n(\mathsf{F}[\partial; \sigma, \delta])$ denotes the set of invertible $n \times n$ matrices over $\mathsf{F}[\partial; \sigma, \delta]$. See for example [6] for the basic skew algebra, or [3, 9] for discussions particular to the Ore polynomial case.

Computing reduced matrices and normal forms of matrices over an Ore polynomial ring $\mathsf{F}[\partial; \sigma, \delta]$ when F is infinite has previously been considered. Beckermann, Cheng and Labahn [3] and Cheng and Labahn [5] compute row reduced bases using an "order basis" approach as known from $\mathsf{F}[x]$ matrix arithmetic, and taking care of coefficient growth. Davies, Cheng and Labahn [19] show how computing the Popov form can be reduced to nullspace computation, a problem for which effective fraction-free techniques exist. Giesbrecht and Kim [9] compute the Hermite normal form of an Ore polynomial matrix by linearizing it to a larger matrix over F. The resulting problem can then be tackled completely

by the usual approaches over F matrices. We follow the same approach here for computing the Popov form. When F is a finite field, the problem is much simpler and faster methods exist [21, 20].

Cost estimates for our algorithm are given in Section 7. We summarize some of these cost estimates here and, to the best of our ability, compare with previous work. Consider computing the Popov form in the differential case when $\mathsf{F} = \mathsf{k}(z)$. We are given a non-singular input matrix $A \in \mathsf{k}(z)[\partial;\sigma,\delta]^{n\times n}$, that is, the entries of A are polynomials in ∂, the coefficients of which are rational functions from $\mathsf{k}(z)$ (c.f. Example 3.) We can assume, without loss of generality, by clearing denominators, that A is over $\mathsf{k}[z][\partial;\sigma,\delta]$. A running time estimate in terms of operations from k thus involves three paramaters: the dimension n; a bound d for $\deg_\partial A$; a bound e for $\deg_z A$. Our algorithm constructs from A a structured matrix of dimension $O(n^2d) \times O(n^2d)$ over $\mathsf{k}[z]$. We then perform a structured fraction free Gaussian elimination to recover the Popov form. The cost of the algorithm is $O(n^{\omega+2}d^3\,\mathsf{M}(n^2de))$ operations from k. Here, ω is an exponent for matrix multiplication, and M is a multiplication time: two polynomials from $\mathsf{k}[z]$ of degree strictly less than t can be multiplied in $\mathsf{M}(t)$ operations from k. Assuming $\theta = 3$ and a pseudo-liner multiplication time, and ignoring logarithmic factors, the cost of our algorithm is then on the order of n^7d^4e operations from k. For comparison, the fraction-free algorithm supporting [3, Corollary 7.7] seems to require on the order of $n^9d^4e^2$ operations from k to produce a row reduced form of A, while the algorithm in [5, Theorem 6.2] requires on the order of $n^8d^4e + n^7d^3e^2$ operations from k.

Now consider the case $\mathsf{k} = \mathbb{Q}$. Like before, we assume our input matrix is over $\mathbb{Z}[z][\partial;\sigma,\delta]$. Ignoring logarithmic factors, and again assuming pseudo-linear integer arithmetic, our algorithm requires on the order of $n^9d^5e\log\beta$ bit operations. Here, β is a paramater that depends on the magnitude of integer coefficients in A (see Theorem 20). The modular algorithm supporting [19, Theorem 6.3] seems to require about $n^{10}d^5e\log\beta + n^9d^4e^2\log\beta$ bit operations.

On the one hand, we point out that the algorithms of [3, 5] solve a considerably more general problem then we do in this paper: they can be applied to input matrices of arbitrary shape and rank and thus compute the rank of the input matrix as well as a left nullspace. Although we hope to consider the rank deficient case in the future, our analysis currently assumes the input matrix is nonsingular. On the other hand, the algorithms in [3, 5] only produce a row reduced form of A and not the canonical Popov forms. In many applications a row reduced form may be sufficient, but in some cases the canonical Popov form can be asymptotically smaller than a row reduced form [14, Appendix B].

Beyond the improved asymptotic worst case cost estimates we have reported above, our algorithm has two additional noteworthy features. First, in the shift case, the worst case running times we have reported above are improved by a factor of n: the linearized system has a special shape in this case which the algorithm is able to exploit. Second, for inputs that are are not too far from being row reduced the running time is asymptotically faster. The *orthogonality defect* of A is the difference between the sum of the row degrees in A and the sum of the row degrees in its Popov form P, denoted by $\mathsf{OD}(A)$. Our algorithms

are output sensitive in the parameter $\mathsf{OD}(A)$, which can be as small as 0 and as large as nd. If $n \leq \mathsf{OD}(A) \leq nd$ then the running times reported above are improved by a factor of $\mathsf{OD}(A)/(nd)$. For $\mathsf{OD}(A) < n$ further improvements are obtained. In the special case $\mathsf{OD}(A) = 0$, which means the input matrix is already reduced, the algorithm detects this and avoids the lions share of the computation, instead applying a fast normalization to transform the input to Popov form.

The rest of this paper is organised as follows. In Section 2 we define some notation and recall some important facts about matrices of Ore polynomials that are established in [3, 9]. In Section 3 we recall the linearization method of Labhalla, Lombardi and Marlin [16] for Hermite form computation over $\mathsf{k}[x]$. Section 4 extends the method to compute the Popov form of a nonsingular matrix of Ore polynomials. Section 5 gives the design and analysis of our algorithm for performing a structured block elimination of the linearized system. Section 6 shows how the elimination can be done in a fraction-free fashion and gives bounds on the sizes of intermediate expressions for some concrete cases of F, namely $\mathsf{F} = \mathsf{k}(z)$, $\mathsf{F} = \mathbb{Q}$ and $\mathsf{F} = \mathbb{Q}(z)$. Cost analysis for computing the Popov form over these rings is provided in Section 7. Section 8 concludes.

2 Preliminaries

First some notation. Let $\mathsf{F}[\partial;\sigma,\delta]$ be an Ore polynomial ring. The *degree* of a vector $\vec{v} \in \mathsf{F}[\partial;\sigma,\delta]^{1\times n}$ or matrix $A \in \mathsf{F}[\partial;\sigma,\delta]^{n\times m}$, denoted $\deg \vec{v}$ respectively $\deg A$, is the maximal degree of entries of $\vec{v}$ or A (we define $\deg 0 = -\infty$). If $\vec{v}$ is non-zero then by the *pivot* of $\vec{v}$, denoted $\mathrm{piv}(\vec{v})$, we mean the right-most entry of $\vec{v}$ which has $\deg \vec{v}$. The elements of $\vec{v}$ are denoted v_i for $i = 1, \ldots, n$. By $\mathrm{rdeg} A$ we mean the list $[d_1, d_2, \ldots, d_n]$ where d_i is the degree of $\mathrm{Row}_i A$, $1 \leq i \leq n$.

Definition 1. *Given a nonsingular $A \in \mathsf{F}[\partial;\sigma,\delta]^{n\times n}$ with $\mathrm{rdeg} A = [d_1, d_2, \ldots, d_n]$, the* leading matrix *of A, denoted $\mathsf{LM}(A) \in \mathsf{F}^{n\times n}$, is the matrix whose (i,j) entry is the coefficient of x^{d_i} of $A_{i,j}$. A is said to be* row reduced *if $\mathrm{rank}(\mathsf{LM}(A)) = n$.*

A canonical row reduced basis is provided by the Popov form. Although the Popov form can be defined for a matrix of arbitrary shape and rank, in this paper we focus on the case a non-singular matrix.

Definition 2. *A non-singular matrix $P \in \mathsf{F}[\partial;\sigma,\delta]^{n\times n}$ is in Popov form if $\mathrm{LM}(P)$ is unit lower triangular and the degrees of off-diagonal entries of P are strictly less than the degree of the diagonal entry in the same column.*

Note that the definition of Popov form implies the pivot index of row i is i, $1 \leq i \leq n$.

Example 3. *Let $A \in \mathbb{Z}_7[z][\partial;\sigma,\delta]^{3\times 3}$ be defined as follows*

$$A = \begin{bmatrix} 1+(5+z)\partial+\partial^2 & & 5+4\partial+4\partial^2 \\ 4+3\partial & 2z & 4\partial+4\partial^2 \\ 3 & & 4+4\partial \end{bmatrix} \to \overbrace{\begin{bmatrix} 2 & -\infty & \mathbf{2} \\ 1 & 0 & \mathbf{2} \\ 0 & -\infty & \mathbf{1} \end{bmatrix}}^{\textit{degree structure}}.$$

The Popov form $P \in \mathbb{Z}_7(z)[\partial;\sigma,\delta]^{3\times 3}$ *of* A *is*

$$\begin{bmatrix} 1+(z+2)\partial+\partial^2 & & 5 \\ \frac{2}{z} & 1 & \\ 6 & & 1+\partial \end{bmatrix} \to \overbrace{\begin{bmatrix} \mathbf{2} & -\infty & 0 \\ 0 & \mathbf{0} & -\infty \\ 0 & -\infty & \mathbf{1} \end{bmatrix}}^{\textit{degree structure}}.$$

A matrix in Popov form is row reduced but the converse is not true. The following is classical for $\mathsf{F}[x]$ matrices, see [15, Section 6.3.2]. For the extension to $\mathsf{F}[\partial;\sigma,\delta]$ matrices, see [3, Lemma A.1 (a)]. The last item is often called the Predictable Degree Property.

Theorem 4. *Let* $A \in \mathsf{F}[\partial;\sigma,\delta]^{n\times n}$ *be non-singular. Then the following are equivalent:*

1. *A is row reduced.*
2. *Among all matrices that are left equivalent to A, the list of row degrees of A, when sorted in non-decreasing order, will be lexicographically minimal.*
3. *For any $\vec{v} \in \mathsf{F}[\partial;\sigma,\delta]^{1\times n}$, we have*

$$\deg(\vec{v}A) = \max_{i=1,\ldots,n} (\deg \operatorname{Row}(A,i) + \deg v_i) \ .$$

Lemma 5. *If $A, U, P \in \mathsf{F}[\partial;\sigma,\delta]^{n\times n}$, all non-singular and U invertible and P in Popov form, then* $\deg U \le (n-1)\deg A$.

Proof. By Item 3 of Theorem 4, then $\deg U^{-1} \le \deg A$ since the degree of P is non-negative. By [9, Corollary 3.3] then $U = (U^{-1})^{-1}$ has degree at most $(n-1)\deg A$. □

The following notion is a measure for how far A is from being row reduced:

Definition 6. *Let $A \in \mathsf{F}[\partial;\sigma,\delta]$ and non-singular. The* orthogonality defect *of A, denoted $\mathsf{OD}(A)$, is given as $\sum \operatorname{rdeg} A - \sum \operatorname{rdeg} P$, where P is the Popov form of A.*

Lemma 7. *If A is row reduced then $\mathsf{OD}(A) = 0$.*

Proof. Follows immediately from Theorem 4, Item 2. □

3 Warm-up: Hermite form of $\mathsf{F}[x]$-matrices via linearization

Let $A \in \mathsf{F}[x]^{n\times n}$ be nonsingular with $\deg A = d$. We recall the method of Labhalla, Lombardi and Marlin [16] that reduces the problem of computing the Hermite form H of A over $\mathsf{F}[x]$ to that of transforming a matrix over F of dimension $(n^2d+n-dn)\times(n^2d+n)$ to reduced row echelon form.

Based on the parameters n and d, define the polynomial linearization $\phi_H : \mathsf{F}[x]^{*\times n} \mapsto \mathsf{F}^{*\times(n^2d+n)}$ by

$$\begin{aligned}\phi_H(v) &= \phi_H\left[\begin{array}{ccc} v_1 & \cdots & v_n \end{array}\right] \\ &= \left[\begin{array}{ccc|c|ccc} [v_1]_{nd} & \cdots & [v_1]_0 & \cdots & [v_n]_{nd} & \cdots & [v_1]_0 \end{array}\right] \in \mathsf{F}^{*\times(n^2d+n)},\end{aligned}$$

where $[v_i]_k$ denotes the coefficient of x^k of $v_i \in \mathsf{F}[x]^{*\times 1}$. The function ϕ_H maps each polynomial (modulo x^{nd+1}) to its coefficient vector of length $nd+1$, padded with zeroes if the polynomial has degree less than nd. For example, if $n = 2$ and $d = 2$ then

$$\phi_H(\left[\begin{array}{c|c} x^2+3x+6 & 2x+3 \end{array}\right]) = \left[\begin{array}{ccccc|ccccc} 0 & 0 & 1 & 3 & 6 & 0 & 0 & 0 & 2 & 3 \end{array}\right].$$

Now note that the F-linear vector space generated by the rows of

$$\phi_H\begin{bmatrix} x^{(n-1)d}A \\ \vdots \\ xA \\ A \end{bmatrix} \in \mathsf{F}^{(n^2d+n-dn)\times(n^2d+n)} \tag{1}$$

is equal to

$$\left\{\phi_H\left(\sum_{i=1}^{n} u_i \operatorname{row}(A,i)\right) \mid u_i \in \mathsf{F}[x],\ \deg u_i \le (n-1)d,\ 1 \le i \le n\right\}. \tag{2}$$

Since the unique unimodular matrix $U \in \mathsf{F}[x]^{n\times n}$ such that UA is in Hermite form has $\deg U \le (n-1)d$, it follows that the ϕ_H-linearized rows of the Hermite form of A are contained in (2). Labhalla, Lombardi and Marlin [16] show that the ϕ_H-linearized rows of the Hermite form of A will appear as rows in the reduced row echelon form of the linearization (1).

Example 8. *Consider the input*

$$A = \begin{bmatrix} 3x^2+6x+6 & 5x^2+3x+3 & 6x^2+x \\ 5x^2+5 & 6x^2+x & 5x^2+2x+6 \\ 3x+5 & 4x+5 & 5x^2+2x+1 \end{bmatrix} \in \mathsf{F}[x]^{3\times 3},$$

where $\mathsf{F} = \mathbb{Z}/(7)$. *Then*

$$\phi_H\begin{bmatrix} x^4A \\ \vdots \\ xA \\ A \end{bmatrix} = \left[\begin{array}{ccc|ccc|ccc} 3&6&6&&&&5&3&3&&&&6&1&0&&&\\ 5&0&5&&&&6&1&0&&&&5&2&6&&&\\ 0&3&5&&&&0&4&5&&&&5&2&1&&&\\ &3&6&6&&&&5&3&3&&&&6&1&0&&\\ &5&0&5&&&&6&1&0&&&&5&2&6&&\\ &0&3&5&&&&0&4&5&&&&5&2&1&&\\ &&3&6&6&&&5&3&3&&&&6&1&0&\\ &&5&0&5&&&6&1&0&&&&5&2&6&\\ &&0&3&5&&&0&4&5&&&&5&2&1\\ &&&3&6&6&&&5&3&3&&&&6&1&0\\ &&&5&0&5&&&6&1&0&&&&5&2&6\\ &&&0&3&5&&&0&4&5&&&&5&2&1 \end{array}\right] \in \mathsf{F}^{15\times 21}. \tag{3}$$

Note that up to a row permutation, the matrix (3) has the shape

$$\left[\begin{array}{ccccccccc|ccccccccc|ccccc}
3&6&6&&&&&&&5&3&3&&&&&&&6&1&&&\\
&3&6&6&&&&&&&5&3&3&&&&&&&6&1&&\\
&&3&6&6&&&&&&&5&3&3&&&&&&&6&1&\\
&&&3&6&6&&&&&&&5&3&3&&&&&&&6&1\\
\hline
5&0&5&&&&&&&6&1&&&&&&&5&2&6&&\\
&5&0&5&&&&&&&6&1&&&&&&&5&2&6&\\
&&5&0&5&&&&&&&6&1&&&&&&&5&2&6\\
&&&5&0&5&&&&&&&6&1&&&&&&&5&2&6\\
\hline
&3&5&&&&&&&&4&5&&&&&&&5&2&1&&\\
&&3&5&&&&&&&&4&5&&&&&&&5&2&1&\\
&&&3&5&&&&&&&&4&5&&&&&&&5&2&1\\
&&&&3&5&&&&&&&&4&5&&&&&&&5&2&1
\end{array}\right] \in \mathsf{F}^{15\times 21},$$

which better illustrates the Sylvestor-like structure. The reduced row echelon form of (3) is

$$\left[\begin{array}{ccccccccc|cccccccc|ccccc}
1&&&&&&&&&4&&&&&&&0&2&1&&&4&0\\
&1&&&&&&&&0&&&&&&&3&6&1&&&0&4\\
&&1&&&&&&&0&&&&&&&4&0&6&&&5&3\\
&&&1&&&&&&0&&&&&&&3&0&0&&&0&4\\
&&&&1&&&&&0&&&&&&&4&0&0&&&4&3\\
&&&&&1&&&&0&&&&&&&3&0&0&&&0&4\\
&&&&&&1&&&0&&&&&&&4&0&0&&&0&0\\
\hline
&&&&&&&1&&&1&&&&&&1&2&2&&&3&6\\
&&&&&&&&1&&&1&&&&&6&0&2&&&4&1\\
&&&&&&&&&&&&1&&&&1&0&0&&&0&6\\
&&&&&&&&&&&&&1&&&6&0&0&&&6&1\\
&&&&&&&&&&&&&&1&1&1&0&0&&&0&6\\
\hline
&&&&&&&&&&&&&&&&&&&1&&6&0\\
&&&&&&&&&&&&&&&&&&&&1&6&0\\
&&&&&&&&&&&&&&&&&&&&&1&6&0
\end{array}\right].$$

The ϕ_H-linearized rows of the Hermite form correspond to the last nonzero row in each horizontal slice, namely rows 7, 12 and 15. Indeed,

$$\phi_H^{-1}\left[\begin{array}{ccccccccc|cccccccc|ccccc}
&&&&&&&1&&0&&&&&&&4&0&0&&&0&0\\
\hline
&&&&&&&&&&&&&&1&1&0&0&&&0&6\\
\hline
&&&&&&&&&&&&&&&&&&&1&6&0
\end{array}\right] = \begin{bmatrix} 1 & 4 & 0 \\ & x+1 & 6 \\ & & x^2+6x \end{bmatrix}$$

is the Hermite form of A.

4 Popov form of $\mathsf{F}[\partial;\sigma,\delta]$ via linearization

In this section we apply the linearization technique to compute the Popov form of $\mathsf{F}[\partial;\sigma,\delta]$ polynomials. Here we deal only with the overall correctness of the approach and its structural properties, while the latter sections give precise algorithms for the steps as well as cost bounds.

Define the linearization $\phi_P : \mathsf{F}[\partial;\sigma,\delta]^{*\times n} \mapsto \mathsf{F}^{*\times(n^2d+n)}$ by

$$\begin{aligned}\phi_P(\vec{v}) &= \phi_P\left[\begin{array}{ccc} v_1 & \cdots & v_n \end{array}\right] \\ &= \left[\begin{array}{ccc|c|ccc} [v_n]_{nd} & \cdots & [v_1]_{nd} & \cdots & [v_n]_0 & \cdots & [v_1]_0 \end{array}\right] \in \mathsf{F}^{*\times(n^2d+n)}.\end{aligned}$$

Note that output of ϕ_P is simply a permutation of the $\mathsf{F}[\partial;\sigma,\delta]$-equivalent of ϕ_H in §3.

Let now A_{lin} be given as the ϕ_P-image of the vectors

$$x^j \operatorname{Row}_i(A) \qquad \text{for } i = 1,\dots,n \text{ and } j = 0,\dots,nd - \deg \operatorname{Row}_i(A)\ ,$$

ordered by descending degrees and breaking ties by the i index. In other words, consider for $t = 0,\dots,nd$ the matrix

$$\hat{B}_t := \operatorname{diag}(x^{t-\deg \operatorname{Row}_1(A)}\operatorname{Row}_1(A),\dots,x^{t-\deg \operatorname{Row}_n(A)}\operatorname{Row}_n(A)).$$

Each row of $\hat{B}_t$ has degree exactly t, but some elements now have negative-degree terms. Let $B_t \in \mathsf{F}^{*\times(n^2d+n)}$ be given as the ϕ_P-image of the rows of $\hat{B}_t$ which have no negative-degree terms. Then A_{lin} consists of putting the B_t on top of each other. More precisely, we can write A_{lin} uniquely in block upper triangular form as

$$A_{\mathsf{lin}} = \begin{bmatrix} \frac{B_{nd}}{\frac{B_{nd-1}}{\frac{\vdots}{B_0}}} \end{bmatrix} = \begin{bmatrix} C_{nd} & * & \cdots & * \\ & C_{nd-1} & \cdots & * \\ & & \ddots & \vdots \\ & & & C_0 \end{bmatrix}\ , \tag{4}$$

where each $C_* \in \mathsf{F}^{*\times n}$ has no zero rows. Note that the rowspace of A_{lin} is in one-to-one correspondence, through ϕ_P, with the set

$$\left\{ \sum_{i=1}^{n} u_i \operatorname{row}(A,i) \;\middle|\; u_i \in \mathsf{F}[\partial;\sigma,\delta],\ \deg u_i \le nd - \deg \operatorname{Row}_i A \right\}. \tag{5}$$

Lemma 9. *If A is non-singular, then A_{lin} has full row rank and row dimension $n^2d + n - \sum \operatorname{rdeg} A$. For $d \le t \le nd$, then B_t (and C_t) has exactly n rows.*

Proof. A_{lin} has full row rank, since any F-linear relation between rows of A_{lin} maps to an $\mathsf{F}[\partial;\sigma,\delta]$-linear relation between rows of A through ϕ_P and by the definition of A_{lin}. No such relation exist since A is non-singular. The i'th row of A is represented in exactly $nd - \deg \operatorname{Row}_i(A) + 1$ of the B_t, so the row dimension of A_{lin} becomes as claimed. This also shows that B_t has n rows when $t \ge d$ since every row of A is represented. □

We say the pivot of a vector $\vec{v} \in \mathsf{F}^{1\times(n^2d+n)}$ is the index of the left-most non-zero element of $\vec{v}$ (not to be confused with pivot of a $\mathsf{F}[\partial;\sigma,\delta]$ vector, see Section 2). Define $\eta : \{1,\dots,n\} \times \mathbb{Z}_{\ge 0} \to \{1,\dots,n^2d+n\}$ as the map between (pivot, degree) of $\mathsf{F}[\partial;\sigma,\delta]^{1\times n}$ vectors and the pivot in $\mathsf{F}^{1\times(n^2d+n)}$ vectors induced by ϕ_P, that is,

$$\eta(i,d') = n(nd - d') + n + 1 - i\ .$$

For a vector $\vec{v} \in \mathsf{F}^{1\times(n^2d+n)}$ we say that the *P-pivot* and *P-degree* of $\vec{v}$ are the first and second components of the 2-tuple $\eta^{-1}(i)$, where i is $\vec{v}$'s pivot.

Now let R_{lin} be the reduced row echelon form of A_{lin}. Then R_{lin} can also be written uniquely in block upper triangular form as

$$R_{\mathsf{lin}} = \begin{bmatrix} T_{nd} & * & \cdots & * \\ & T_{nd-1} & \cdots & * \\ & & \ddots & \vdots \\ & & & T_0 \end{bmatrix}, \tag{6}$$

where each $T_* \in \mathsf{F}^{*\times n}$ has no zero rows. Note that those rows in R_{lin} with P-degree t are contained in the submatrix of R_{lin} occupied by T_t. Because R_{lin} is in echelon form, for any given degree t and pivot i, there is at most one row in R_{lin} with P-degree t and P-pivot i, and any row in the row space of R_{lin} with P-degree t and P-pivot i will be a linear combination of this row, and possibly rows below it.

Theorem 10. *Let A be non-singular, and let R_{lin} be the reduced row echelon form of A_{lin}. Then the Popov form P of A is the matrix whose i'th row is the ϕ_P^{-1}-image of the row of R_{lin} with minimal P-degree having P-pivot i.*

Proof. The unique unimodular matrix $U \in \mathsf{F}[\partial;\sigma,\delta]^{n\times n}$ with $UA = P$ has $\deg U \leq (n-1)d$ by Lemma 5, and therefore the ϕ_P-linearized rows of P are contained in the row space of R_{lin}. We will prove that they in fact appear directly as rows of R_{lin}. By the minimality of the row degrees of the Popov form, Item 2 of Theorem 4, the rows chosen as in the theorem must therefore be exactly those rows of R_{lin}.

So for $1 \leq i \leq n$, consider the i'th row $\vec{p}$ of P, which has pivot i. Since $\phi_P(\vec{p})$ is in the row space of R_{lin}, there must be exactly one row $\vec{r}_k$ of R_{lin} with the same pivot, with row index k. If $\vec{w}$ is the unique vector over F satisfying $\vec{w}R_{\mathsf{lin}} = \phi_P(\vec{p})$ then clearly $w_k = 1$ and $w_j = 0$ for $j < k$. We claim $w_j = 0$ also for $j > k$ in which case $\vec{r}_k = \phi_P(\vec{p})$ as we wanted to prove. Suppose, to arrive at a contradiction, that $w_j \neq 0$ for some $j > k$, and let $\vec{r}_j$ be the j'th row of R_{lin}. Since all other rows of R_{lin} are zero at the pivot position of $\vec{r}_j$, that means $\deg p_{i,j'} \geq d'$, where j', d' are the P-pivot respectively P-degree of $\vec{r}_j$. On the other hand, since the $\phi_P^{-1}(\vec{r}_k)$ is in the row space of A and has pivot j', the minimality of the degrees of the Popov form implies $d' \geq \deg p_{j',j'}$. But then $\deg p_{i,j'} \geq \deg p_{j',j'}$, which contradicts that P is in Popov form. We conclude that $w_j = 0$ for $j > k$, and hence $\phi_P(\vec{p}) = \vec{r}_k$. □

Example 11. *For clarity, we exemplify the approach with a usual polynomial ring, i.e. $\sigma = \mathsf{id}$ and $\delta = 0$. Consider the input $A \in \mathsf{F}[x]^{3\times 3}$ from Example 8,*

$\mathsf{F} = \mathbb{Z}/(7)$. *Then*

$$A_{\mathsf{lin}} = \phi_P \begin{bmatrix} x^4 A \\ \vdots \\ xA \\ A \end{bmatrix} = \left[\begin{array}{cccccccccccccccccccc}
6&5&3&1&3&6&0&3&6&&&&&&&&&&&\\
5&6&5&2&1&0&6&0&5&&&&&&&&&&&\\
5&0&0&2&4&3&1&5&5&&&&&&&&&&&\\
&&&6&5&3&1&3&6&0&3&6&&&&&&&&\\
&&&5&6&5&2&1&0&6&0&5&&&&&&&&\\
&&&5&0&0&2&4&3&1&5&5&&&&&&&&\\
&&&&&&6&5&3&1&3&6&0&3&6&&&&&\\
&&&&&&5&6&5&2&1&0&6&0&5&&&&&\\
&&&&&&5&0&0&2&4&3&1&5&5&&&&&\\
&&&&&&&&&6&5&3&1&3&6&0&3&6\\
&&&&&&&&&5&6&5&2&1&0&6&0&5\\
&&&&&&&&&5&0&0&2&4&3&1&5&5
\end{array}\right] \in \mathsf{F}^{15\times 21}. \tag{7}$$

The row reduced echelon form of A_{lin} *is*

$$R_{\mathsf{lin}} = \left[\begin{array}{cccccccccccccccccccc}
1&&&&&6&&&6&&&&&&2&&0&0\\
&1&2&&&4&&&6&&&&&&3&&6&5\\
\hline
&&&1&&0&&&6&&&&&&2&&4&1\\
&&&&1&2&&&4&&&&&&2&&5&3\\
\hline
&&&&&&1&&0&&&&&&6&&0&0\\
&&&&&&&1&2&&&&&&0&&0&0\\
\hline
&&&&&&&&&1&&&&&6&&0&0\\
&&&&&&&&&&1&&&&0&&6&5\\
&&&&&&&&&&&1&&&0&&4&1\\
\hline
&&&&&&&&&&&&1&&6&&0&0\\
&&&&&&&&&&&&&1&6&&1&2\\
&&&&&&&&&&&&&&4&&3&6\\
\hline
&&&&&&&&&&&&&&&1&6&5\\
&&&&&&&&&&&&&&&&4&1\\
\hline
&&&&&&&&&&&&&&&&&1&2
\end{array}\right].$$

We pick out the the rows of R_{lin} *with minimal degree having P-pivot 1,2, respectively 3. The Popov form of A is thus*

$$\phi_P^{-1} \left[\begin{array}{cccccccc|ccc}
&&&&&&&1&4&&1\\
&&&&&&&&&1&2\\
&&&&&&&&0&&0\\
\hline
&&&&&&1&&&&\\
&&&&&6&&&&&\\
&&&&1&&&&&&\\
\end{array}\right] = \begin{bmatrix} x+1 & 0 & 4 \\ 2 & 1 & 0 \\ 0 & 0 & x^2+6x \end{bmatrix} \in \mathsf{F}[x]^{3\times 3}.$$

We now consider some structural properties of A_{lin} and R_{lin}. Recall that we have written A_{lin} as a block upper triangular matrix

$$A_{\mathsf{lin}} = \left[\begin{array}{ccc|ccc}
C_{nd} & \cdots & * & * & \cdots & * \\
& \ddots & \vdots & \vdots & \ddots & \vdots \\
& & C_{k+1} & * & \cdots & * \\
\hline
& & & C_k & \cdots & * \\
& & & & \ddots & \vdots \\
& & & & & C_0
\end{array}\right]$$

with each C_* of column dimension n and with no zero rows. Let $\{a_1, a_2, \ldots, a_n\}$ be the multi-set of row degrees of A. The following lemma follows from the definition of A_{lin}.

Lemma 12. *For $k = 0, 1, \ldots, nd$, the trailing submatrix*

$$\begin{bmatrix} C_k & \cdots & * \\ & \ddots & \vdots \\ & & C_0 \end{bmatrix}$$

of A_{lin} has row dimension $(k+1)n - \sum_{i=1}^{n} \min(a_i, k+1)$.

We have also written R_{lin} in a block upper triangular form as

$$R_{\mathsf{lin}} = \left[\begin{array}{ccc|ccc} T_{nd} & \cdots & * & * & \cdots & * \\ & \ddots & \vdots & \vdots & \ddots & \vdots \\ & & T_{k+1} & * & \cdots & * \\ \hline & & & T_k & \cdots & * \\ & & & & \ddots & \vdots \\ & & & & & T_0 \end{array}\right]$$

where each $T_* \in \mathsf{F}^{*\times n}$ has no zero rows. Let $\{p_1, p_2, \ldots, p_n\}$ be the multi-set of row degrees in the Popov form of A. The following lemma follows as a corollary of Theorem 10.

Lemma 13. *For $k = 0, 1, \ldots, nd$, the trailing submatrix*

$$\begin{bmatrix} T_k & \cdots & * \\ & \ddots & \vdots \\ & & T_0 \end{bmatrix} \tag{8}$$

of R_{lin} has row dimension at most $(k+1)n - \sum_{i=1}^{n} \min(p_i, k+1)$.

For $k = 0, 1, \ldots, nd$, define OD_k to be the nullity (row dimension of the left nullspace) of the principal submatrix

$$\begin{bmatrix} C_{nd} & \cdots & * \\ & \ddots & \vdots \\ & & C_{k+1} \end{bmatrix} \tag{9}$$

of A_{lin}. Recall that $\mathsf{OD}(A) := \sum \mathrm{rdeg} A - \sum \mathrm{rdeg} P$.

Theorem 14. *For $k = 0, 1, \ldots, nd$ we have $\mathsf{OD}_k \leq \mathsf{OD}(A)$.*

Proof. Let P be a permutation, and U be a unit lower nonsingular matrix over F that such that premultiplying (9) by UP transforms it to echelon form

$$\left[\begin{array}{c} R_{k+1} \\ \hline \end{array}\right]$$

with OD_k zero rows. Then applying $\mathrm{diag}(U, I)$ to A_{lin} yields

$$\left[\begin{array}{c|c} U & \\ \hline & I \end{array}\right] \left[\begin{array}{ccc|ccc} C_{nd} & \cdots & * & * & \cdots & * \\ & \ddots & \vdots & \vdots & \ddots & \vdots \\ & & C_{k+1} & * & \cdots & * \\ \hline & & & C_k & \cdots & * \\ & & & & \ddots & \vdots \\ & & & & & C_0 \end{array}\right] = \left[\begin{array}{ccc|ccc} & & & * & \cdots & * \\ & R_{k+1} & & E_k & \cdots & * \\ \hline & & & C_k & \cdots & * \\ & & & & \ddots & \vdots \\ & & & & & C_0 \end{array}\right], \quad (10)$$

where $E_k \in \mathsf{F}^{\mathsf{OD}_k \times n}$. Considering that A_{lin} has full row rank, the row dimension of the submatrix

$$\left[\begin{array}{ccc} E_k & \cdots & * \\ C_k & \cdots & * \\ & \ddots & \vdots \\ & & C_0 \end{array}\right] \quad (11)$$

of the matrix on the right of (10) will be equal to the row dimension of the trailing submatrix (8) of R_{lin}. Lemmas 12 and 13 now give

$$\begin{aligned} \mathsf{OD}_k &\leq \sum_{i=1}^{n} \min(a_i, k+1) - \sum_{i=1}^{n} \min(p_i, k+1) \\ &= \sum_{i=1}^{n} (\min(a_i, k+1) - \min(p_i, k+1)). \end{aligned}$$

Assume now that $a_1 \leq a_2 \leq \cdots \leq a_n$ and $p_1 \leq p_2 \leq \cdots p_n$. Then $a_i - p_i \geq 0$ for $i = 1, 2, \ldots, n$ by Item 2 of Theorem 4, and

$$\min(a_i, k+1) - \min(p_i, k+1) \begin{cases} = & a_i - p_i & \text{if } a_i \leq k+1 \\ = & 0 & \text{if } a_i > k+1 \text{ an } p_i \geq k+1 \\ < & a_i - p_i & \text{if } a_i > k+1 \text{ and } p_i < k+1 \end{cases}.$$

Thus $\min(a_i, k+1) - \min(p_i, k+1) \leq a_i - p_i$ in all cases, establishing the result. □

5 Block elimination of the linearized system

Let $A \in \mathsf{F}[\partial; \sigma, \delta]^{n \times n}$ be non-singular with $\deg A \leq d$. In this section we show how to perform a structured Gaussian elimination of the linearized system A_{lin} over F. We first consider in Section 5.1 the problem of transforming A to Popov form when A is already row reduced, the so called normalization problem [22]. Then we consider the general case in Section 5.2.

5.1 Normalization of an already row reduced matrix

We can detect if A is row reduced by testing its leading coefficient matrix for non-singularity. Suppose A is already row reduced. Let U be the unique matrix

such that $UA = P$ is in Popov form. By the predictable degree property [3, Lemma A.1 (a)] then $\deg \mathrm{Col}_i U \leq d - \deg \mathrm{Row}_i A$, $1 \leq i \leq n$. Consider the following submatrix of A_{lin} (4) comprised of the last $n(d+1) - \sum \mathrm{rdeg} A$ rows:

$$\bar{A}_{\mathsf{lin}} = \begin{bmatrix} B_d \\ \vdots \\ B_0 \end{bmatrix} = \left[\begin{array}{ccc|ccc} & & & C_d & \cdots & * \\ & & & & \ddots & \vdots \\ & & & & & C_0 \end{array}\right].$$

Note that the rowspace of $\bar{A}_{\mathsf{lin}}$ is in one-to-one correspondence, through ϕ_P, with the set

$$\left\{ \sum_{i=1}^{n} u_i \, \mathrm{Row}_i A \;\middle|\; u_i \in \mathsf{F}[\partial; \sigma, \delta],\ \deg u_i \leq d - \deg \mathrm{Row}_i A \right\}.$$

As a corollary of Lemma 9 we have that $\bar{A}_{\mathsf{lin}}$ has full row rank $n(d+1) - \sum \mathrm{rdeg} A$. The next theorem is a corollary of Theorem 10.

Theorem 15. *Let A be non-singular and row reduced, and let $\bar{R}_{\mathsf{lin}}$ be the reduced row echelon form of $\bar{A}_{\mathsf{lin}}$. Then the Popov form P of A is the matrix S whose i'th row is the ϕ_P^{-1}-image of the row of $\bar{R}_{\mathsf{lin}}$ with minimal degree having P-pivot i.*

Because A is row reduced, by Lemma 7 we have $\mathsf{OD}(A) = 0$, so by Theorem 14 each block C_* in $\bar{A}_{\mathsf{lin}}$ will have full row rank. Since the right block of $\bar{A}_{\mathsf{lin}}$ has column dimension $n(d+1)$, performing standard Gauss Jordan elimination would cost $O((nd)^3)$ operations from F to produce $\bar{R}_{\mathsf{lin}}$ in its entirety. We can save a factor of d by avoiding the complete computation of $\bar{R}_{\mathsf{lin}}$. Instead, first compute an echelon form of $\bar{A}_{\mathsf{lin}}$ by applying Gaussian elimination to each full rank slice B_*. Gaussian elimination of a single B_* has cost $O(n^3 d)$, yielding a total cost for all slices of $O(n^3 d^2)$ operations in F. Then use back substitution to reduce the n rows whose ϕ_P^{-1}-image has minimal degree and P-pivot i, $1 \leq i \leq n$. This costs an additional $O(n^3 d^2)$. Finally, scale these n rows so their pivots are equal to one. We obtain the following result.

Theorem 16. *Let $A \in \mathsf{F}[\partial; \sigma, \delta]^{n \times n}$ be non-singular and row reduced, with $\deg A \leq d$. The Popov form of A can be computed within the following cost:*

- *Computing $\partial^k \, \mathrm{Row}_i A$ for $1 \leq k \leq d - \deg \mathrm{Row}_i A$, $1 \leq i \leq n$.*
- *An additional $O(n^3 d^2)$ operations from F.*

5.2 General case

Now assume that A is not already row reduced, so that $\mathsf{OD} := \mathsf{OD}(A) > 0$. The key observation is that since $\deg P \leq d$, the ϕ_P-linearization of the rows of P will be contained in the row space of the trailing submatrix (8) of R_{lin} for $k = d$. This implies that the rows of R_{lin} occupied by $T_{nd}, T_{nd-1}, \cdots, T_{d+1}$ are not required.

Our algorithm for performing the elimination of A_{lin} has three phases. The first phase computes the matrix (11) for $k = d$, whose rowspace is equal to that of (8) for $k = d$. The second phase transforms this matrix to row echelon form. The third phase performs back substitution to reduce the n rows whose ϕ_P^{-1}-image has minimal degree and P-pivot i, $1 \leq i \leq n$.

Our main computation tool is the Gauss transform [23, Section 2.3]. Given as input a matrix

$$\begin{bmatrix} E_k \\ C_k \end{bmatrix} \in \mathsf{F}^{O(\mathsf{OD}+n)\times n}, \tag{12}$$

the so called Gauss transform algorithm [23, Algorithm 2.14] can be used to produce a permutation matrix P_k and unit lower triangular matrix U_k such that

$$\left[\begin{array}{c|c} F_k & \\ \hline N_k & I \end{array}\right]^{U_k} P_k \begin{bmatrix} E_k \\ C_k \end{bmatrix} = \left[\begin{array}{c} G_k \\ \hline \end{array}\right],$$

with $\begin{bmatrix} N_k \mid I \end{bmatrix} P_k$ and G_k are the left nullspace basis and a row echelon form, respectively, of the input matrix (12).

Phase 1: For convenience, let E_0 be the $0 \times n$ matrix. We will compute a Gauss transform as described above for $k = nd, nd-1, \ldots, d+1$. At the start of stage k we are exactly in the situation shown in (10). The key observation is that no entries in the rows occupied by R are required, and so the computation of these rows can be avoided. To go from stage k to $k+1$ we can thus apply only the nullspace to the next slice and obtain

$$\begin{bmatrix} N_k \mid I \end{bmatrix} P_k \left[\begin{array}{c|ccc} E_k & * & \cdots & * \\ C_k & * & \cdots & * \end{array}\right] = \begin{bmatrix} \mid E_{k-1} & \cdots & * \end{bmatrix}.$$

Continue this for $k = nd, nd-1, \ldots, d+1$.

Phase 2: For $k = d, d-1, \ldots, 0$, we apply the complete Gauss transform to the work matrix:

$$U_k P_k \left[\begin{array}{c|ccc} E_k & * & \cdots & * \\ C_k & * & \cdots & * \end{array}\right] = \left[\begin{array}{c|ccc} G_k & * & \cdots & * \\ & E_{k-1} & \cdots & * \end{array}\right].$$

Repeating this for $k = d, d-1, \ldots, 0$, we have computed the row echelon form

$$G = \begin{bmatrix} G_d & \cdots & * \\ & \ddots & \vdots \\ & & G_0 \end{bmatrix} \in \mathsf{F}^{*\times n(d+1)}.$$

Phase 3: Identify for $i = n, n-1, \ldots, 1$, the row in G whose ϕ_P^{-1}-image has minimal degree and P-pivot i, and use back substitution to zero out the entries in this row which are above a pivot, similar to how we proceeded in Section 5.1. Finally, scale these n rows to make their pivots equal to one.

Example 17. *Consider the input matrix*

$$A = \begin{bmatrix} 7x^2+3x+8 & 9x^2+7x+4 & x^2+2x+2 \\ 3x^2+4 & 7x^2+6x+8 & 5x^2+10x \\ 3x^2+2x+5 & 7x^2+5x+1 & 4x^2+8x+5 \end{bmatrix} \in \mathbb{Z}/(11)^{3\times 3}.$$

Then

$$A_{\mathsf{lin}} = \left[\begin{array}{cccccc|cccccccccccc} C_6 & * & * & * & * & * & & & & & & & & & & & \\ & C_5 & * & * & * & * & & & & & & & & & & & \\ & & C_4 & * & * & * & & & & & & & & & & & \\ & & & C_3 & * & * & & & & & & & & & & & \\ & & & & C_2 & * & & & & & & & & & & & \\ & & & & & * & & & & & & & & & & & \\ \hline & & & & & & 1&9&7&2&7&3&2&4&8 & & & \\ & & & & & & 5&7&3&10&6&0&0&8&4 & & & \\ & & & & & & 4&7&3&8&5&2&5&1&5 & & & \\ \hline & & & & & & & & & 1&9&7&2&7&3&2&4&8 \\ & & & & & & & & & 5&7&3&10&6&0&0&8&4 \\ & & & & & & & & & 4&7&3&8&5&2&5&1&5 \end{array}\right]$$

$$= \left[\begin{array}{cccccccccccccccccc} 1&9&7&2&7&3&2&4&8 & & & & & & & & & \\ 5&7&3&10&6&0&0&8&4 & & & & & & & & & \\ 4&7&3&8&5&2&5&1&5 & & & & & & & & & \\ \hline & & & 1&9&7&2&7&3&2&4&8 & & & & & & \\ & & & 5&7&3&10&6&0&0&8&4 & & & & & & \\ & & & 4&7&3&8&5&2&5&1&5 & & & & & & \\ \hline & & & & & & 1&9&7&2&7&3&2&4&8 & & & \\ & & & & & & 5&7&3&10&6&0&0&8&4 & & & \\ & & & & & & 4&7&3&8&5&2&5&1&5 & & & \\ \hline & & & & & & & & & 1&9&7&2&7&3&2&4&8 \\ & & & & & & & & & 5&7&3&10&6&0&0&8&4 \\ & & & & & & & & & 4&7&3&8&5&2&5&1&5 \end{array}\right]$$

For Phase 1, step $k = 6$ we compute and and apply the nullspace of C_6 to obtain

$$\left[\begin{array}{cccccc} E_5 & * & * & * & * & * \\ \hline C_5 & * & * & * & * & * \\ & C_4 & * & * & * & * \\ & & C_3 & * & * & * \\ & & & C_2 & * & * \end{array}\right] = \left[\begin{array}{cccccccccccccccccc} & & & 0&0&0&0&4&8 & & & & & & & & & \\ \hline & & & 1&9&7&2&7&3&2&4&8 & & & & & & \\ & & & 5&7&3&10&6&0&0&8&4 & & & & & & \\ & & & 4&7&3&8&5&2&5&1&5 & & & & & & \\ \hline & & & & & & 1&9&7&2&7&3&2&4&8 & & & \\ & & & & & & 5&7&3&10&6&0&0&8&4 & & & \\ & & & & & & 4&7&3&8&5&2&5&1&5 & & & \\ \hline & & & & & & & & & 1&9&7&2&7&3&2&4&8 \\ & & & & & & & & & 5&7&3&10&6&0&0&8&4 \\ & & & & & & & & & 4&7&3&8&5&2&5&1&5 \end{array}\right].$$

Phase 1, step $k = 5$ we compute and apply the nullspace of the 4×3 matrix occupied by E_5 and C_5 to obtain

$$\left[\begin{array}{cccccc} E_4 & * & * & * & * & \\ \hline C_4 & * & * & * & * & \\ & C_3 & * & * & * & \\ & & C_2 & * & * & \\ & & & * & * & \end{array}\right] = \left[\begin{array}{cccccccccccccccccc} & & & & 0&4&8 & & & & & & & & & & & \\ & & & & 0&0&0&0&4&8 & & & & & & & & \\ \hline & & & & 1&9&7&2&7&3&2&4&8 & & & & & \\ & & & & 5&7&3&10&6&0&0&8&4 & & & & & \\ & & & & 4&7&3&8&5&2&5&1&5 & & & & & \\ \hline & & & & & & & 1&9&7&2&7&3&2&4&8 & & \\ & & & & & & & 5&7&3&10&6&0&0&8&4 & & \\ & & & & & & & 4&7&3&8&5&2&5&1&5 & & \\ \hline & & & & & & & & & & 1&9&7&2&7&3&2&4&8 \\ & & & & & & & & & & 5&7&3&10&6&0&0&8&4 \\ & & & & & & & & & & 4&7&3&8&5&2&5&1&5 \end{array}\right].$$

We continue Phase 1 steps $k = 4, 3$ with nullspace application. Then we switch to Phase 2. For Phase 2 steps $k = 2, 1, 0$ we apply the entire Gauss transforms,

yielding the echelon form

$$\left[\begin{array}{c|c|c|c} & G_2 & * & * \\ \hline & & G_1 & * \\ \hline & & & G_0 \end{array}\right] = \left[\begin{array}{c|c|c|c} & 1\ 9\ 7 & 2\ 7\ 3 & 2\ 4\ 8 \\ & \ \ 1\ 10 & 3\ 8\ 2 & 0\ 0\ 0 \\ & \ \ \ \ 1 & 10\ 1\ 3 & 0\ 0\ 0 \\ \hline & & \ \ 1\ 2 & 0\ 0\ 0 \\ & & \ \ \ \ 1 & 10\ 1\ 3 \\ \hline & & & \ \ 1\ 2 \end{array}\right].$$

In Phase 3, we identify the Popov rows (rows 1*,* 5 *and* 6 *in this example) and then do back substitution:*

$$\left[\begin{array}{c|c|c|c} & 1 & 2 & 2\ \ 0 \\ & \ \ 1\ 10 & 3\ \ 8\ 2 & 0\ \ 0\ 0 \\ & \ \ \ \ 1 & 10\ 1\ 3 & 0\ \ 0\ 0 \\ \hline & & \ \ 1\ 2 & 0\ \ 0\ 0 \\ & & \ \ \ \ 1 & 10\ \ 1 \\ \hline & & & \ \ 1\ 2 \end{array}\right].$$

The Popov form of A *is thus*

$$\phi_P^{-1}\left[\begin{array}{cc|cc|cc} & & & 1 & 10 & 1 \\ & & & & 1 & 2 \\ \hline 1 & & 2 & & 2 & 0 \end{array}\right] = \left[\begin{array}{ccc} x+1 & 0 & 10 \\ 2 & 1 & 0 \\ 0 & 0 & x^2+2x+2 \end{array}\right].$$

The next theorem gives a cost analysis of the algorithm just described in terms of operations from F. The theorem give three cost estimates. First, we give an unconditional cost estimate based only on the input parameters n and d. Second, we give a refined cost estimate in terms of OD. Third, we consider the case of special Ore rings (such as the shift case) for which A_{lin} matrix may have the shape shown in Example 11 , that is, with a large block upper triangular submatrix of zeroes in the northeast corner: the cost estimates are improved by a factor of n in this case.

Theorem 18. *Let* $A \in \mathsf{F}[\partial;\sigma,\delta]^{n\times n}$ *be non-singular with* $\deg A \le d$*. The Popov form of* A *can be computed within the following costs.*

1. *General case:*
   - *Computing* $\partial^k \operatorname{Row}_i A$ *for* $1 \le k \le nd - \deg \operatorname{Row}_i A$*,* $1 \le i \le n$*.*
   - *Additional* $O(n^{\omega+2}d^3)$ *field operations from* F*.*
2. *A more refined cost is obtained by considering the parameter* OD*. Assume that* A *is not already row reduced, so that* $\mathsf{OD} := \mathsf{OD}(A) > 0$*. Then the number of additional operations is reduced to:*
   - $O(\mathsf{OD}^{\omega-2}n^4d^2)$ *if* $\mathsf{OD} < n$
   - $O(\mathsf{OD}\, n^{\omega+1}d^2)$ *if* $\mathsf{OD} \ge n$
3. *Finally, suppose that the Ore ring* $\mathsf{F}[\partial;\sigma,\delta]$ *has the property that for any nonzero element* $f \in \mathsf{F}[\partial;\sigma,\delta]$*, the trailing degree of* ∂f *is at least one more than the trailing degree of* f*. Then the* O*-estimates in parts 1 and 2 above for the additional operations are reduced by a factor of* n*.*

Proof. We first establish part 2 of the theorem. By Theorem 14, the row dimension of each nullspace N_* is bounded by OD. Instead of considering the three phases separately, we will partition the computational work done as follows. The nullspace N_k is applied for all k, $0 \leq k \leq nd$, but the unit lower triangular block F_k is applied only for $0 \leq k \leq d$. Also note that for $k \leq d$, the column dimension of the slice to which F_k is being applied to is bounded by $(d+1)n$. The application of the permutations P_* do not dominate the cost. We can thus partition the computational work as follows.

A Gauss transform: at most $nd+1$ times, compute a Gauss transform of a matrix bounded in dimension $O(\mathsf{OD}+n) \times n$.

B Nullspace application: at most $nd+1$ times, multiply an $O(\mathsf{OD}) \times n$ matrix by an $n \times O(n^2 d)$ matrix.

C Computing the echelon form: at most $d+1$ times, multiply an $O(n) \times O(n)$ matrix by a $O(n) \times O(nd)$ matrix.

D Back substitution: $O(n^3 d^2)$ operations.

Since the rank of the input matrix (12) is bounded by its column dimension n, the cost of computing (U_k, P_k) for a given k is bounded by $O((\mathsf{OD}+n)n^{\omega-1})$ by [23, Proposition 2.15]. This gives a total cost for (A) of $O((\mathsf{OD}+n)n^\omega d)$. Using an obvious block decomposition shows (C) can be done in time $O(n^\omega d^2)$.

It remains to bound the cost of (B). There are two cases, depending on whether $\mathsf{OD} \leq n$ or $\mathsf{OD} > n$. Using an obvious block decomposition shows a single nullspace application has cost $O(\mathsf{OD}^{\omega-2} n^3 d)$ if $\mathsf{OD} \leq n$ and $O(\mathsf{OD}\, n^\omega d)$ otherwise. The total cost for (B) is thus $O(\mathsf{OD}^{\omega-2} n^4 d^2)$ if $\mathsf{OD} \leq n$ and $O(\mathsf{OD}\, n^{\omega+1} d^2)$ if $\mathsf{OD} > n$. In both cases these upper bounds for the cost of (B) dominate the cost bounds for (A), (C) and (D).

Part 1 of the theorem follows by substituting the a priori upper bound $\mathsf{OD} \leq nd$ into the O-bound in part 2 for the case $\mathsf{OD} \geq n$.

For part 3, note that the (nonzero part) of the slice to which the nullspace is applied will now have dimension $O(nd)$ instead of $O(n^2 d)$. The cost estimates for the work in part (B) are thus reduced by a factor of n, but they still dominate the cost of parts (A), (C) and (D). □

6 Fraction free block elimination

Now consider the case when all entries in A_{lin} are coming from an integral domain, for example $\mathsf{F} = \mathsf{k}(z)$ for a field k but all entries are in $\mathsf{F}[z]$, or even $\mathbb{Z}[z]$ when $\mathsf{k} = \mathbb{Q}$. It is desirable in this setting to keep all intermediate quantities in the computation integral, while at the same time controlling their growth. The classic technology for this purpose in the linear algebra setting is fraction free Gaussian elimination [7, 2].

The Gauss transform algorithm [23, Algorithm 2.14] is actually designed to do fraction free Gaussian elimination, and because of its column recursive formulation, is well suited to the elimination of A_{lin}.

The incorporation of fraction free techniques into the algorithm supporting Theorem 18 is straightforward. For $k = nd, nd-1, \ldots, -1$, the fraction free Gauss transform algorithm also computes Δ_k, the minor of R_k in (10) comprised of its rank profiles columns. To start the process set $\Delta_{nd+1} = 1$ since R_{nd+1} is the 0×0 matrix, and recall that E_{nd} is the $0 \times n$ matrix. At step k we have the scaled matrix

$$\Delta_{k+1} \left[\begin{array}{c|ccc} E_k & * & \cdots & * \end{array} \right]$$

from the previous step, together with Δ_{k+1}. The rows of A_{lin} occupied by C_k are premultiplied by Δ_{k+1} to form the next slice

$$\Delta_{k+1} \left[\begin{array}{c|ccc} E_k & * & \cdots & * \\ C_k & * & \cdots & * \end{array} \right], \tag{13}$$

which will be fraction-free, that is, all entries are minors of A_{lin} of dimension bounded by one plus the rank of R_{k+1}. (We remark that only scaling the rows of A_{lin} that will be involved in the next elimination step is important for the complexity, and similar to [17].) At stage k, the fraction free Gauss transform takes as input (13), together with Δ_{k+1}, and returns as output the permutation P_k and the scaled matrix

$$\bar{U}_k = \left[\begin{array}{c|c} \bar{F}_k & \\ \hline \Delta_k N_k & \Delta_k I \end{array} \right], \tag{14}$$

together with Δ_k. The matrix $\bar{F}_k$ is equal to the unit lower triangular F_k from before but with each row scaled by a certain minor of A_{lin} which is known a priori to clear any denominators. The output (14) is also fraction-free, that is, all entries are minors of A_{lin} of dimension bounded by the rank of R_k. The nullspace applications in Phase 1 can now be done in a fraction free fashion as

$$\frac{1}{\Delta_{k+1}} \left(\left(\Delta_k \left[\begin{array}{c|c} N_k & I \end{array} \right] P_k \right) \left(\Delta_{k+1} \left[\begin{array}{c|ccc} E_k & * & \cdots & * \\ C_k & * & \cdots & * \end{array} \right] \right) \right)$$

yielding

$$\Delta_k \left[\begin{array}{c|ccc} & E_{k-1} & \cdots & * \end{array} \right].$$

In Phase 2 the entire Gauss transform is applied to obtain

$$\frac{1}{\Delta_{k+1}} \left(\bar{U}_k P_k \left(\Delta_{k+1} \left[\begin{array}{c|ccc} E_k & * & \cdots & * \\ C_k & * & \cdots & * \end{array} \right] \right) \right) = \Delta_k \left[\begin{array}{c|ccc} \bar{G}_k & * & \cdots & * \\ & E_{k-1} & \cdots & * \end{array} \right].$$

The back substitution in Phase 3 can be done iteratively in a fraction free fashion also [2].

Using the fraction free approach, all intermediate quantities arising during the elimination (i.e., the entries of (13) and (14)) will thus be minors of A_{lin}. We recall some well known a priori bounds for the size of these minors for some common cases. We will use size and $\overline{\mathsf{size}}$ for the bounds for A_{lin} and $\bar{A}_{\mathsf{lin}}$ respectively.

- $\mathsf{F} = \mathsf{k}[z]$ with $\deg_z A_{\mathsf{lin}}, \deg_z \bar{A}_{\mathsf{lin}} \leq e$. Multiplying the row dimension of A_{lin} and $\bar{A}_{\mathsf{lin}}$ by e gives explicit bounds for the degrees $\mathsf{size}_{\mathsf{k}[z]}$ and $\overline{\mathsf{size}}_{\mathsf{k}[z]}$ of minors of A_{lin} and $\bar{A}_{\mathsf{lin}}$ that satisfy

$$\mathsf{size}_{\mathsf{k}[z]} \in O(n^2 de) \quad \text{and} \quad \overline{\mathsf{size}}_{\mathsf{k}[z]} \in O(nde).$$

- $\mathsf{F} = \mathbb{Z}$ with the magnitude of entries of A_{lin} and $\bar{A}_{\mathsf{lin}}$ bounded by β. Hadamard's inequality [13, Corollary 7.82] gives an explicit bound $2^{\mathsf{size}_{\mathbb{Z}}}$ and $2^{\overline{\mathsf{size}}_{\mathbb{Z}}}$ for the magnitudes of minors of A_{lin} and $\bar{A}_{\mathsf{lin}}$ that satisfy

$$\mathsf{size}_{\mathbb{Z}} \in O(n^2 d \log(nd\beta)) \quad \text{and} \quad \overline{\mathsf{size}}_{\mathbb{Z}} \in O(nd \log(nd\beta)).$$

- $\mathsf{F} = \mathbb{Z}[z]$ with $\deg_z A \leq e$, and with the magnitude of integer coefficients of entries of A_{lin} and $\bar{A}_{\mathsf{lin}}$ bounded by β. Multiplying the determinant degree bound above with the logarithm base 2 of an explicit magnitude bound for the coefficients [10] gives

$$\mathsf{size}_{\mathbb{Z}[z]} \in O(n^4 d^2 e \log(nde\beta)) \quad \text{and} \quad \overline{\mathsf{size}}_{\mathbb{Z}[z]} \in O(n^2 d^2 e \log(nde\beta)).$$

Now let M be a multiplication time for $\mathsf{k}[z]$, that is, two polynomials from $\mathsf{k}[z]$ with degree strictly less than t can be multiplied in $\mathsf{M}(t)$ operations from k. Then over $\mathsf{k}[z]$ a cost estimate in terms of operations from k is obtained by multiplying the algebraic cost estimates of Theorem 18 by $\mathsf{M}(\mathsf{size}_{\mathsf{k}[z]})$. Note that the polynomial multiplication can be done modulo z^p for $p = 2\,\mathsf{size}_{\mathsf{k}[z]} + 1$ to control degrees during the fast matrix multiplications.

If M is a multiplication time for $\mathbb{Z}$, that is, two integers with bit-length bounded by t can be multiplied with $\mathsf{M}(t)$ bit operations, then a cost estimate in terms of bit operations for the cases $\mathbb{Z}$ and $\mathbb{Z}[z]$ are obtained by multiplying the algebraic cost estimates by $\mathsf{M}(\mathsf{size}_{\mathbb{Z}})$ and $\mathsf{M}(\mathsf{size}_{\mathbb{Z}[z]})$. Note that for the $\mathbb{Z}[z]$ case we can use Kronecker substitution [11] to reduce the integer polynomial multiplication to integer multiplication. Similar to the case $\mathsf{k}[z]$, the multiplication can be done modulo 2^p for an appropriate $p \in O(\mathsf{size}_{\mathbb{Z}[z]})$.

All these cost estimates can be improved (by logarithmic factors) by performing the matrix multiplications using a homomorphic imaging scheme. For example, if $\#\mathsf{k} > 2nd+1$, then two $n \times n$ matrices over $\mathsf{k}[x]$ with degree bounded by d can be multiplied using only $O(n^\omega d + n^2 \mathsf{M}(d))$ operations from k [4], instead of $O(n^\omega \mathsf{M}(d))$. For integer matrix multiplication we refer to [12].

7 Cost analysis for some common Ore rings

Let $A \in \mathsf{F}[\partial; \sigma, \delta]^{n \times n}$ be nonsingular with degree d. Theorem 18 gave cost estimates for computing the Popov form P of A in terms of operations from F. In this section we give refined cost estimates for some specializations of F, focusing on the differential and shift cases.

First consider the case $\mathsf{F} = \mathsf{k}(z)$. We will assume that A has entries over $\mathsf{k}[z]$. This can be achieved by clearing denominators, if necessary. As in [9], we will

assume that $\sigma(z) \in \mathsf{k}[z]$ and $\deg_z \delta(z) \leq 1$. Then $\partial z = \sigma(z)\partial + \delta(z) \in \mathsf{k}[z][\partial]$ and the degree in z in ∂ remains unchanged. The linearized systems will thus be over $\mathsf{k}[z]$, allowing application of our fraction-free algorithm.

In the following theorems recall that $\mathsf{OD} = \mathsf{OD}(A) := \sum \mathrm{rdeg} A - \sum \mathrm{rdeg} P$.

Theorem 19. *Let $A \in \mathsf{k}[z][\partial;\sigma,\delta]^{n\times n}$ be non-singular with $\deg A \leq d$ and $\deg_z A \leq e$. If A is row reduced, this can be detected and the Popov form of A computed in $O(n^3 d^2\, \mathsf{M}(nde))$ operations from k. If A is not row reduced, the Popov form can be computed in*

1. *$O(n^{\omega+2} d^3\, \mathsf{M}(n^2 de))$ field operations from k.*
2. *A more refined cost is obtained by considering the parameter OD. Then the number of operations from k is reduced to:*
   - *$O(\mathsf{OD}^{\omega-2} n^4 d^2\, \mathsf{M}(n^2 de))$ if $\mathsf{OD} < n$*
   - *$O(\mathsf{OD}\, n^{\omega+1} d^2\, \mathsf{M}(n^2 de))$ if $\mathsf{OD} \geq n$*
3. *Finally, suppose that $\mathsf{k}[z][\partial;\sigma,\delta]$ has the property that for any nonzero element $f \in \mathsf{k}[z][\partial;\sigma,\delta]$, the trailing degree of ∂f is at least one more than the trailing degree of f. Then the O-estimates in parts 1 and 2 above are reduced by a factor of n.*

Proof. To test if A is row reduced we can check if its leading coefficient matrix is nonsingular. This check will not dominate the cost. The theorem now follows from Theorems 15 and 18 and the estimates for $\overline{\mathsf{size}}_{\mathsf{k}[z]}$ and $\mathsf{size}_{\mathsf{k}[z]}$. □

Now consider the case $\mathsf{F} = \mathbb{Q}(z)$. As before, we will assume that A has entries over $\mathbb{Z}[z]$. Then A has entries polynomials in ∂ whose coefficients are polynomials in $\mathbb{Z}[z]$; let $||A||_\infty$ denote the largest in abolulte value of any (integer) coefficient of any of these $\mathbb{Z}[z]$ coefficients.

Theorem 20. *et $A \in \mathbb{Z}[z][\partial;\sigma,\delta]^{n\times n}$ be non-singular with $\deg A \leq d$ and $\deg_z A \leq e$. Suppose our Ore ring is either the differential polynomials (where $\sigma(z) = z$, $\delta(z) = 1$) or the shift polynomials (where $\sigma(z) = z+1$, $\delta(z) = 0$). If A is row reduced, this can be detected and the Popov form of A computed in $O(n^3 d^2)$ operations with integers bounded in length by $O^{\sim}(n^2 d^2 e(\log ||A||_\infty + e))$ bits. If A is not row reduced, the Popov form can be computed in*

1. *$O(n^{\omega+2} d^3)$ operations with integers bounded in length by $O^{\sim}(n^4 d^2 e(\log ||A||_\infty + e))$ bits.*
2. *A more refined cost is obtained by considering the parameter OD. Then the number of operations on integers is reduced to:*
   - *$O(\mathsf{OD}^{\omega-2} n^4 d^2)$ if $\mathsf{OD} < n$*
   - *$O(\mathsf{OD}\, n^{\omega+1} d^2)$ if $\mathsf{OD} \geq n$*
3. *In the shift case the O-estimates for numbers of operations in parts 1 and 2 above are reduced by a factor of n.*

Proof. As mentioned in Section 6, we can use Kronecker substition to reduce arithmetic operations from $\mathbb{Z}[z]$ to integer arithmetic. As before, we can test if A is row reduced by checking if its leading coefficient matrix is nonsingular using fraction-free gaussian elimination. From the proof of [9, Corollary 5.9] we have that $\log \beta := \log ||A_{\mathsf{lin}}||_\infty \in O(\log ||A|| + e \log(nd))$. The theorem now follows from Theorems 15 and 18 and the estimates for $\overline{\mathsf{size}}_{\mathbb{Z}[z]}$ and $\mathsf{size}_{\mathbb{Z}[z]}$. □

8 Conclusion

We give a new algorithm to compute the Popov form of a non-singular matrix $A \in \mathsf{F}[\partial; \sigma, \delta]^{n \times n}$ over an Ore polynomial ring. Our approach is to construct a matrix A_{lin} over F whose reduced row echelon form reveals the Popov form of A. Using structural properties of A_{lin}, and since we need only a small, a priori known part of the reduced echelon form, we are able to speed up the Gaussian elimination, especially in cases where A is already close to being row reduced. This approach combines immediately with existing fraction-free techniques which utilize fast matrix multiplication. We remark that a direct application of the iterative fraction-free elimination method of [17] would seem to yield the same cost estimates we report in Section 7 with θ replaced by 3.

For the shift and differential Ore rings we bound the bit complexity of our algorithm when the input matrix is over $\mathbb{Z}[z]$. Our algorithm is faster than [5] which computes only a row reduced form and not the Popov form (but, however, handles the important case of non-square and non-singular inputs which we do not.) Our approach still works for rectangular and rank deficient matrices (of any dimensions) but the complexity could be much higher. Making the algorithm as efficient for rectangular and rank deficient matrices is left for future work.

References

[1] D. Augot, P. Loidreau, and G. Robert. Rank metric and Gabidulin codes in characteristic zero. July 2013.

[2] E. H. Bareiss. Sylvester's identity and multistep integer-preserving Gaussian elimination. *Mathematics of Computation*, 22(103):565–578, 1968.

[3] B. Beckermann, H. Cheng, and G. Labahn. Fraction-free row reduction of matrices of Ore polynomials. *J. Symb. Comp.*, 41(5):513–543, 2006.

[4] A. Bostan and E. Schost. Polynomial evaluation and interpolation on special sets of points. *Journal of Complexity*, 21(4):420–446, 2005. Festschrift for the 70th Birthday of Arnold Schönhage.

[5] H. Cheng and G. Labahn. Modular computation for matrices of ore polynomials. *Computer Algebra 2006: Latest Advances in Symbolic Algorithms*, pages 43–66, 2007.

[6] P. K. Draxl. *Skew Fields*, volume 81. Cambridge Univ. Press, 1983.

[7] J. Edmonds. On systems of distinct linear representative. *J. Res. Nat. Bur. Standards*, 71B:241–245, 1967.

[8] E. M. Gabidulin. Theory of codes with maximum rank distance. *Problemy Peredachi Informatsii*, 21(1):3–16, 1985.

[9] M. Giesbrecht and M. S. Kim. Computing the Hermite form of a matrix of ore polynomials. *Journal of Algebra*, 376:341–362, 2012.

[10] A. J. Goldstein and R. L. Graham. A Hadamard-type bound on the coefficients of a determinant of polynomials. *SIAM Review*, 16:394–395, 1974.

[11] D. Harvey. Faster polynomial multipication via multipoint Kronecker substitution. *Journal of Symbolic Computation*, 44(10):1502–1510, 3009.

[12] D. Harvey and J. van der Hoeven. On the commplexity of integer matrix multiplication, 2014. https://hal.archives-ouvertes.fr/hal-01071191.

[13] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.

[14] C. P. Jeannerod, V. Neiger, E. Schost, and G. Villard. Fast computation of minimal interpolation bases in Popov form for arbitrary shifts. ACM Press, New York, 2016.

[15] T. Kailath. *Linear Systems*. Prentice Hall, Englewood Cliffs, N.J., 1980.

[16] S. E. Labhalla, H. Lombardi, and R. Marlin. Algorithmes de calcul de la réduction d'Hermite d'une matrice à coefficients polynomiaux. In *Comptes-Rendus de MEGA92, Nice, France*. Birkhauser, 1992.

[17] H. R. Lee and B. D. Saunders. Fraction free Gaussian elimination for sparse matrices. *Journal of Symbolic Computation*, 19(5):393–402, 1995.

[18] O. Ore. Theory of non-commutative polynomials. *Annals of Mathematics*, 34:480–508, 1933.

[19] G. L. P. Davies, H. Cheng. Computing popov form of general ore polynomial matrices. In *Proceedings of Milestones in Computer Algebra*, pages 149–156, 2008.

[20] S. Puchinger, S. Müelich, D. Mödinger, J. S. R. Nielsen, and M. Bossert. Decoding Interleaved Gabidulin Codes using Alekhnovich's Algorithm. In *Proc. of ACCT*, 2016. Preprint available as arXiv:1604.04397.

[21] S. Puchinger, J. S. R. Nielsen, W. Li, and V. Sidorenko. Row Reduction Applied to Decoding of Rank Metric and Subspace Codes. *Designs, Codes and Cryptography*, Feb. 2016. Submitted.

[22] S. Sarkar and A. Storjohann. Normalization of row reduced matrices. In A. Leykin, editor, *Proc. Int'l. Symp. on Symbolic and Algebraic Computation: ISSAC'11*, pages 297–303. ACM Press, New York, 2011.

[23] A. Storjohann. *Algorithms for Matrix Canonical Forms*. PhD thesis, Swiss Federal Institute of Technology, ETH–Zurich, 2000.