

Algorithms for Linearly Recurrent Sequences of Truncated Polynomials

Seung Gyu Hyun
University of Waterloo
Waterloo, ON, Canada

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

Éric Schost
University of Waterloo
Waterloo, ON, Canada

ABSTRACT

Linear recurrent sequences are those whose elements are defined as linear combinations of preceding elements, and finding recurrence relations is a fundamental problem in computer algebra. In this paper, we focus on sequences whose elements are vectors over the ring $\mathbb{A} = \mathbb{K}[x]/\langle x^d \rangle$ of truncated polynomials. Finding the ideal of their recurrence relations has applications such as the computation of minimal polynomials and determinants of sparse matrices over \mathbb{A} . We present three methods for finding this ideal: a Berlekamp-Massey-like approach due to Kurakin, one which computes the kernel of some block-Hankel matrix over \mathbb{A} via a minimal approximant basis, and one based on bivariate Padé approximation. We propose complexity improvements for the first two methods, respectively by avoiding the computation of redundant relations and by exploiting the Hankel structure to compress the approximation problem. Then we confirm these improvements empirically through a C++ implementation, and we discuss the above-mentioned applications.

CCS CONCEPTS

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

KEYWORDS

Linear recurrences; Berlekamp-Massey-Sakata; Approximant basis; Kurakin’s algorithm; Sparse matrix.

ACM Reference Format:

Seung Gyu Hyun, Vincent Neiger, and Éric Schost. 2021. Algorithms for Linearly Recurrent Sequences of Truncated Polynomials. In *Proceedings of the 2021 International Symposium on Symbolic and Algebraic Computation (ISSAC ’21)*, July 18–23, 2021, Virtual Event, Russian Federation. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3452143.3465533>

1 INTRODUCTION

Linear recurrences appear in many domains of computer science and mathematics, and computing recurrence relations efficiently is a fundamental problem in computer algebra. More specifically, given a sequence of elements in \mathbb{K}^r for some field \mathbb{K} and integer $r > 0$, we seek a representation of its *annihilator*, which is a polynomial ideal corresponding to all recurrence relations which are satisfied

by the sequence; the polynomials in the annihilator are said to *cancel* the sequence. In dimension $r = 1$, the Berlekamp-Massey algorithm [4, 26] computes the unique monic univariate polynomial of minimal degree that cancels the sequence. Sakata extended this algorithm first to dimension 2 [33] and then to the general case $r > 1$ [34]; see also Norton and Fitzpatrick’s extension to $r > 1$ [13]. Recent work includes variants of Sakata’s algorithm such as one which handles relations that are satisfied by several sequences simultaneously [35], approaches relating the problem to the kernel of a multi-Hankel matrix and exploiting either fast linear algebra [5] or a process similar to Gram-Schmidt orthogonalization [27], and an algorithm relying directly on multivariate polynomial arithmetic [6]. As for the representation of the output, all these algorithms compute a Gröbner basis or a border basis of the annihilator.

In this paper, we focus on computing recurrence relations for sequences whose elements are in \mathbb{A}^r , where $\mathbb{A} = \mathbb{K}[x]/\langle x^d \rangle$. This problem can be solved using a specialization of Kurakin’s algorithm [20, 21], as detailed in Section 3, where we explicitly describe the output generating set of the annihilator as a lexicographic Gröbner basis of some bivariate ideal. We derive a cost bound of $O(\delta d(n^2 \delta d + n^\omega d))$ operations in \mathbb{K} , where δ is the order of recurrence (see Section 2.1), and ω is an exponent for matrix multiplication over \mathbb{K} [1, 9, 24]. Because the Gröbner bases computed by Kurakin’s algorithm are often non-minimal, in Section 4 we propose a modified algorithm which aims at limiting as much as possible the computation of these extraneous generators. This lowers the cost to $O(\delta d^*(n^2 \delta d + n^\omega d))$, where d^* is a number arising in the algorithm as an upper bound on the cardinality d_{opt} of minimal Gröbner bases of the annihilator. In Section 7, we observe empirically that d^* is often close or equal to d_{opt} .

Despite the improvement, the above cost bound still has a dependence at least quadratic in the dimension n . Our interest in the case $n \gg 1$ is motivated among others by the following fact: given a zero-dimensional ideal $\mathcal{I} \in \mathbb{K}[x, y]$, one can recover a Gröbner basis of it via $\mathcal{I} = \text{Ann}(\mathbf{s})$ for some well-chosen $\mathbf{s} \in \mathbb{A}^N$ only if $\mathbb{K}[x, y]/\mathcal{I}$ has the *Gorenstein* property [16, 25]. When that is not the case, one can recover a basis of \mathcal{I} via the annihilator of *several* sequences simultaneously, which means precisely $n > 1$. For large n , we compute the annihilator via a minimal approximant basis of a block-Hankel matrix over \mathbb{A} constructed from \mathbf{s} . Computing this approximant basis via the algorithm PM-BASIS of [14] leads to a complexity of $O(\delta^\omega n d)$ operations in \mathbb{K} (Section 5.1). We then propose a novel improvement of this minimal approximant basis computation, based on a randomized compression of the input matrix which leverages its block-Hankel structure, reducing the cost to $O(\delta^2 n d + \delta^\omega d)$ operations in \mathbb{K} (Section 5.2).

The four above algorithms have been implemented in C++ using the libraries NTL [37] and PML [18], using Lazard’s structural

theorem [23] for generating examples of sequences; see Section 7 for more details. Our experiments on a prime field \mathbb{K} highlight a good match between cost bounds and practical running times, confirming also the benefit obtained from the improvements of both Kurakin's algorithm and the plain approximant basis approach.

Furthermore, in Section 6 we propose an algorithm with cost quasi-linear in the order δ , whereas the above cost bounds are at least quadratic. For $d \in O(\delta)$, we compute the annihilator via the bivariate Padé approximation algorithm of [28]: this uses $O(d^{\omega+1}\delta)$ operations in \mathbb{K} , at the price of restricting to $n \in O(1)$.

Finally, in Section 8 we mention applications to the computation of minimal polynomials and determinants of sparse matrices over \mathbb{A} . To design Wiedemann-like algorithms [39] for such matrices $A \in \mathbb{A}^{\mu \times \mu}$, we need to compute annihilators from sequences of the form $(u^T A^i v)_{i \geq 0} \in \mathbb{A}^{\mathbb{N}}$ for some vectors u and v ; several such sequences may be needed, leading to the case $n > 1$.

Sakata's 2-dimensional algorithm shares similarities with the case $n = 1$ of Kurakin's algorithm, and has the same complexity $O(\delta^2 d^2)$ [33, Thm. 3]. Apart from this, to the best of our knowledge previous work has $n = 1$ and considers r -dimensional sequences over \mathbb{K} for an arbitrary $r \geq 2$ [5, 6, 27]. Complexity in this r -variate context is often expressed using the degree D of the considered zero-dimensional ideal; here, $\delta \leq D \leq \delta d$ and a minimal Gröbner basis or a border basis will have at most $\min(\delta, d) + 1$ elements. The SCALAR-FGLM algorithm has cost $O(d_{\text{opt}} \delta^\omega d)$ [5, Prop. 16]. Both the Artinian border basis and POLYNOMIAL-SCALAR-FGLM algorithms [6, 27] cost $O(D^2 \delta d)$, which is $O(\delta^3 d)$ in the most favourable case $D = \delta$, and $O(\delta^3 d^3)$ when $D \in \Theta(\delta d)$ (which will be the case in our experiments, see Section 7). In all cases, a better complexity bound can be achieved by one of our algorithms outlined above.

While this is not reflected in the cost estimates above, Kurakin's algorithm and our modified version are still affected by the shape of the staircase of the computed Gröbner basis, due to early termination of the iterations and late additions; we leave a more refined complexity analysis with respect to D as future work.

2 LINEARLY RECURRENT SEQUENCES

In this section, we review key facts about linearly recurrent sequences and algorithmic tools used throughout the paper.

2.1 Recurrent sequences over $\mathbb{K}[x]/\langle x^d \rangle$

We consider the set $\mathcal{S} = (\mathbb{A}^n)^{\mathbb{N}}$ of (vector) sequences over the ring $\mathbb{A} = \mathbb{K}[x]/\langle x^d \rangle$ for some $d \in \mathbb{Z}_{>0}$, that is, sequences $s = (S_0, S_1, \dots)$ with each S_k in \mathbb{A}^n . Such a sequence is said to be *linearly recurrent* if there exist $\gamma \in \mathbb{N}$ and $p_0, \dots, p_\gamma \in \mathbb{A}$ with p_γ invertible such that

$$p_0 S_k + \dots + p_{\gamma-1} S_{k+\gamma-1} + p_\gamma S_{k+\gamma} = 0 \text{ for all } k \geq 0; \quad (1)$$

the *order* of s is the smallest such γ , denoted by δ hereafter. A polynomial $p_0 + \dots + p_\gamma y^\gamma$ in $\mathbb{A}[y]$ is said to *cancel* s if p_0, \dots, p_γ satisfies Eq. (1) (without requiring that p_γ be invertible). The set of canceling polynomials forms an ideal $\text{Ann}(s)$ in $\mathbb{A}[y]$, called the *annihilator* of s . Thus s is linearly recurrent of order δ if and only if there is a monic polynomial of degree δ in $\text{Ann}(s)$: such polynomials are called *generating polynomials* of s . Unlike for sequences over fields, here there may be canceling polynomials of degree less than δ , which prevents uniqueness of generating polynomials; and

there are sequences which are not linearly recurrent but still admit a nonzero canceling polynomial (i.e. $\text{Ann}(s) \neq \{0\}$).

Example 2.1. Consider $\mathbb{A} = \mathbb{K}[x]/\langle x^2 \rangle$ and the sequence $s = (1, 1+x, 1, 1+x, 1, 1+x, \dots)$ in $\mathbb{A}^{\mathbb{N}}$. Note that $xs = (x, x, x, x, \dots)$. This sequence has order $\delta = 2$, a generating polynomial is $y^2 - 1$, and a canceling polynomial of degree less than 2 is $x(y-1)$. One can verify that $\text{Ann}(s) = \langle y^2 - 1, x(y-1) \rangle$; in particular $y^2 + x(y-1) - 1$ is also a generating polynomial. For any sequence s in $\mathbb{K}^{\mathbb{N}}$ which is not linearly recurrent, the sequence xs in $\mathbb{A}^{\mathbb{N}}$ is not linearly recurrent but is canceled by x , i.e. $x \in \text{Ann}(xs) \setminus \{0\}$.

Like for sequences over fields, here canceling polynomials can be characterized as denominators of the (vector) generating series of the sequence, defined as $G_s = \sum_{k \geq 0} S_k y^{-k-1}$ in $(\mathbb{A}[[y^{-1}]])^n$. In what follows, the elements of $\mathbb{A}[y]^n$ are called polynomials, and for $g = (g_1, \dots, g_n) \in \mathbb{A}[y]^n$ we define $\deg(g) = \max_{1 \leq j \leq n} \deg(g_j)$.

LEMMA 2.2. *Let $s \in \mathcal{S}$, let G_s be its generating series, and let $p \in \mathbb{A}[y]$. Then, $p \in \text{Ann}(s)$ if and only if the series $pG_s \in (\mathbb{A}[[y^{-1}]])^n$ is a polynomial, in which case $\deg(pG_s) < \deg(p)$.*

In this paper, we want to compute a generating set for $\text{Ann}(s)$, for a linearly recurrent $s \in \mathcal{S}$, but for algorithms we typically only have access to a finite number of terms of the sequence. Suppose we have access to the partial sequence $s_e = (S_0, \dots, S_{e-1})$ in $\mathcal{S}_e = (\mathbb{A}^n)^e$, for some $e \in \mathbb{Z}_{>0}$. Similar to Eq. (1), a polynomial $p_0 + \dots + p_\gamma y^\gamma$ of degree $\gamma < e$ cancels s_e if

$$p_0 S_k + \dots + p_\gamma S_{k+\gamma} = 0 \text{ for all } 0 \leq k < e - \gamma. \quad (2)$$

Like for sequences over fields, here polynomials of degree γ which cancel s_e also cancel the whole sequence s , provided the discrepancy between e and γ is sufficiently large (namely, $e \geq \gamma + \delta$).

LEMMA 2.3. *Let $s \in \mathcal{S}$ be linearly recurrent of order δ . For any $e \in \mathbb{Z}_{>0}$ and any $p \in \mathbb{A}[y]$ with $\deg(p) \leq e - \delta$, one has $p \in \text{Ann}(s)$ if and only if p cancels s_e .*

2.2 Bivariate interpretation and generating sets

Uni-dimensional sequences of vectors in \mathbb{A}^n as above can be interpreted as two-dimensional sequences of vectors in \mathbb{K}^n , that is, sequences $\sigma = (\zeta_{i,j})_{i,j \geq 0}$ in $\mathfrak{S} = (\mathbb{K}^n)^{\mathbb{N}^2}$. This is based on the natural injection $\varphi : \mathbb{A}[y] \rightarrow \mathbb{K}[\alpha, \beta]$ with $(\varphi(x), \varphi(y)) = (\alpha, \beta)$.

Here we recall from [13, 33] that a polynomial $q = \sum_{i,j} q_{ij} \alpha^i \beta^j$ in $\mathbb{K}[\alpha, \beta]$ is said to cancel a sequence $\sigma = (\zeta_{i,j})_{i,j \geq 0} \in \mathfrak{S}$ if

$$\sum_{i,j} q_{ij} \zeta_{i+k_1, j+k_2} = 0 \text{ for all } k_1, k_2 \geq 0.$$

Then, let $s = (S_0, S_1, \dots) \in \mathcal{S}$, and define $\sigma = (\zeta_{i,j})_{i,j \geq 0} \in \mathfrak{S}$ such that $\zeta_{i,j} \in \mathbb{K}^n$ is the coefficient of degree $d-1-i$ of the truncated polynomial vector $S_j \in \mathbb{A}^n$ if $i < d$, and $\zeta_{i,j} = 0$ otherwise. Then, a polynomial $p \in \mathbb{A}[y]$ cancels s if and only if the polynomial $\varphi(p)$ cancels σ . Furthermore, the set of polynomials in $\mathbb{K}[\alpha, \beta]$ which cancel σ is an ideal of $\mathbb{K}[\alpha, \beta]$ which contains α^d , and this ideal is zero-dimensional if and only if s is linearly recurrent.

In what follows, we define $\bar{\varphi}(\mathcal{I}) = \{\{\varphi(p) \mid p \in \mathcal{I}\} \cup \{\alpha^d\}\}$ for any ideal \mathcal{I} of $\mathbb{A}[y]$, providing a correspondence between the ideals of $\mathbb{A}[y]$ and those of $\mathbb{K}[\alpha, \beta]$ containing α^d . For insight into possible "nice" generating sets for $\text{Ann}(s)$, we consider the lexicographic order $\preccurlyeq_{\text{lex}}$ with $\alpha \preccurlyeq_{\text{lex}} \beta$, and use the fact that Gröbner bases of the ideals in $\mathbb{K}[\alpha, \beta]$ for this order are well understood [23]. Below,

unless mentioned otherwise, we use \leq_{lex} when some term order is needed, e.g. leading terms and Gröbner bases.

Consider a zero-dimensional ideal \mathcal{I} in $\mathbb{K}[\alpha, \beta]$ that contains a power of α and let \mathcal{G} be its reduced Gröbner basis. Let

$$(\beta^{e_0}, \alpha^{d_1} \beta^{e_1}, \dots, \alpha^{d_{t-1}} \beta^{e_{t-1}}, \alpha^{d_t})$$

be the leading terms of the elements of \mathcal{G} listed in decreasing order, i.e. the e_i 's are decreasing and the d_i 's are increasing. We set $d_0 = e_t = 0$, and for $1 \leq i \leq t$ we set $\delta_i = d_i - d_{i-1}$, so that $d_i = \delta_1 + \dots + \delta_i$. Similarly, for $0 \leq i < t$ we set $e_i = e_i - e_{i+1}$. Then write $\mathcal{G} = \{g_0, \dots, g_t\}$, with g_i having leading term $\alpha^{d_i} \beta^{e_i}$; in particular $g_t = \alpha^{d_t} = \alpha^{\delta_1 + \dots + \delta_t}$ and g_0 is monic in β .

Lazard's Theorem states the following [23]: for $0 \leq i \leq t$ one can write $g_i = \alpha^{d_i} \hat{g}_i$, with \hat{g}_i monic of degree e_i in β . In addition, for $0 \leq i < t$, $\hat{g}_i = g_i / \alpha^{d_i}$ is in the ideal generated by

$$\langle \hat{g}_{i+1}, \alpha^{\delta_{i+2}} \hat{g}_{i+2}, \dots, \alpha^{\delta_{i+2} + \dots + \delta_t} \rangle = \left\langle \frac{g_{i+1}}{\alpha^{d_{i+1}}}, \frac{g_{i+2}}{\alpha^{d_{i+1}}}, \dots, \frac{g_t}{\alpha^{d_{i+1}}} \right\rangle;$$

in particular, α^{δ_1} divides g_1, \dots, g_t . Lazard also proved that a set of polynomials which satisfies these conditions is necessarily a minimal Gröbner basis.

With the above notation, a minimal Gröbner basis of \mathcal{I} has cardinality $t + 1$, with $t \leq \min(e_0, d_t)$ since $0 = d_0 < d_1 < \dots < d_t$ and $0 = e_t < \dots < e_1 < e_0$. Since for the reduced Gröbner basis \mathcal{G} each polynomial g_i is represented by at most $e_0 d_t$ coefficients in \mathbb{K} , the total size of \mathcal{G} in terms of field elements is at most $e_0 d_t \min(e_0, d_t)$. Finer bounds for the cardinality and size of \mathcal{G} could be given using the vector space dimension $\dim_{\mathbb{K}}(\mathbb{K}[\alpha, \beta] / \mathcal{I})$.

2.3 Univariate and bivariate approximation

For a univariate polynomial matrix $F \in \mathbb{K}[x]^{\mu \times \nu}$ and a positive integer d , we consider a free $\mathbb{K}[x]$ -module of rank μ defined as

$$\mathcal{A}_d(F) = \{p \in \mathbb{K}[x]^{1 \times \mu} \mid pF = 0 \text{ mod } x^d\};$$

its elements are called *approximants for F at order d* [2, 38]. Bases of such submodules can be represented as $\mu \times \mu$ nonsingular matrices over $\mathbb{K}[x]$ and are usually computed in so-called *reduced forms* [40] or the corresponding canonical *Popov forms* [31]. Extensions of these forms have been defined to accommodate degree weights or degree constraints, and are called *shifted reduced* or *Popov forms* [2, 3, 38]. The algorithm PM-BASIS [14] computes an approximant basis in shifted reduced form in time $O(\mu^{\omega-1}(\mu + \nu)d)$; using essentially two calls to this algorithm, one recovers the unique approximant basis in shifted Popov form within the same cost bound [19].

More generally, in the bivariate case with $F \in \mathbb{K}[\alpha, \beta]^{\mu \times \nu}$ and $(d, e) \in \mathbb{Z}_{>0}$, the set

$$\mathcal{A}_{d,e}(F) = \{p \in \mathbb{K}[\alpha, \beta]^{1 \times \mu} \mid pF = 0 \text{ mod } (\alpha^d, \beta^e)\}$$

is a $\mathbb{K}[\alpha, \beta]$ -submodule of $\mathbb{K}[\alpha, \beta]^{1 \times \mu}$ whose elements are called *approximants for F at order (d, e)*. Such submodules are usually represented by a \leq -Gröbner basis for some term order \leq on $\mathbb{K}[\alpha, \beta]^{1 \times \mu}$; for definitions of term orders and Gröbner bases for submodules we refer to [10]. For $\nu \leq \mu$ algorithms based on an iterative approach or on efficient linear algebra yield cost bounds in $O(\mu(vde)^2 + (vde)^3)$ and $O(\mu(vde)^{\omega-1} + (vde)^\omega)$ operations in \mathbb{K} respectively [12, 30], whereas a recent divide and conquer approach costs $O((M^\omega + M^2\nu)de)$, where $M = \mu \min(d, e)$ [28, Prop. 5.5]; in these cases the output is a minimal Gröbner basis.

3 KURAKIN'S ALGORITHM

In [20], Kurakin gives an algorithm based on the Berlekamp-Massey algorithm that computes the annihilators of a partial sequence over a ring R (and modules over R) that can be decomposed as a disjoint union $R = \{0\} \cup R_0 \cup \dots \cup R_{d-1}$ where

$$R_i = \{r_i r^* \mid r^* \in R \text{ invertible}\} \text{ for some } r_i \in R.$$

In this paper we consider $R = \mathbb{A} = \mathbb{K}[x] / \langle x^d \rangle$; in this case the canonical choice is $r_i = x^i$, with

$$R_i = \{x^i p^* \mid p^* \in \mathbb{A} \text{ with nonzero constant term}\}.$$

Consider a partial sequence $s_e \in \mathcal{S}_e$ of a linearly recurrent $s \in \mathcal{S}$ of order δ . Kurakin's algorithm computes d polynomials $P_i \in \mathbb{A}[y]$, $i = 0, \dots, d-1$, such that P_i is a canceling polynomial of s_e that has leading coefficient x^i and is minimal in degree among all canceling polynomials with leading coefficient x^i . Furthermore, one has $\text{Ann}(s) = \langle P_0, \dots, P_{d-1} \rangle$ provided $e \geq 2\delta$ [21, Thm. 1].

We first define three operations on sequences. Given a partial sequence s_e and $c \in \mathbb{A}$, $c \cdot s_e$ denotes multiplying c to every element in s_e , while $y^j \cdot s_e$ denotes a shift of j elements — that is, removing the first j elements. Given another partial sequence \hat{s}_e , the sum $s_e + \hat{s}_e$ returns the first $\min(e, \hat{e})$ elements of the two sequences added together element-wise.

Kurakin's algorithm iterates on $s = 0, \dots, e-1$, keeping track of polynomials $P_{i,s}$ as well as partial sequences $s_{e,i,s} = P_{i,s} \cdot s_e = \sum_{j=0}^{e-s} P_{i,s}[j] \cdot y^j \cdot s_e$, where $P_{i,s}[j]$ is the j -th coefficient of $P_{i,s}$. An invariant is that the leading coefficient of $P_{i,s}$ is x^i for all s . For each $s = 0, \dots, e-1$, the algorithm essentially attempts to either create a zero by using the partial sequences from previous iterations with equal number of leading zeros (similar to Gaussian elimination), or shift the sequence if we cannot cancel this element.

At each iteration s , let $\mathcal{I}[k]$ be the \mathbb{A} -submodule of \mathbb{A}^n generated by the elements $s_{e,i,s'}[k]$ for all $i = 0, \dots, d-1$ and $s' < s$ such that $s_{e,i,s'}$ has k leading zeros. Furthermore, let $\mathcal{P}[k, j]$ and $\mathcal{S}[k, j]$ be the corresponding polynomial and partial sequence to the j -th element in the basis of $\mathcal{I}[k]$, $\mathcal{I}[k, j]$. At iteration s , if $s_{e,i,s}$ has k leading zeros and $s_{e,i,s}[k] \in \mathcal{I}[k]$, then we can find coefficients such that $s_{e,i,s}[k] - \sum_j c_j \mathcal{I}[k, j] = 0$ and $s_{e,i,s} - \sum_j c_j \mathcal{S}[k, j]$ results in a sequence with at least $k+1$ zeros since both sequences had k leading zeros and we canceled $s_{e,i,s}[k]$. The algorithm terminates when all $s_{e,i,s} = 0$ (see Algorithm 1).

We track the subiterations by the index t for analysis; this does not play a role in the algorithm. Kurakin shows that the total number of subiterations across all s is $O(e)$ per polynomial, bringing the total to $O(ed)$ ([20, Thm. 2]). However, the analysis of the runtime in [20] treats all ring operations (including computing solution to line 12 of Algorithm 1) as constant time operations, which is unrealistic over \mathbb{A}^n . Thus, we will give a cost analysis in terms of number of field operations over \mathbb{K} .

We note that, since \mathbb{A}^n is a free $\mathbb{K}[x]$ -module of rank n (with a basis given by the canonical vectors of length n) and $\mathbb{K}[x]$ is a principal ideal domain, any of its $\mathbb{K}[x]$ -submodule is free of rank at most n . As a consequence, the number of generators of $\mathcal{I}[k]$ is at most n . This will allow us to bound the cost for solving submodule membership as well as the equation $s_{e,s,i}^{(t)}[k] - \sum_j c_j \mathcal{I}[k, j] = 0$.

Algorithm 1 KURAKIN(s_e)

Input: partial sequence s_e **Output:** minimal canceling polynomials of s_e

```
1: for  $i = 0, \dots, d-1$  do
2:   set  $P_{i,0} = x^i$  and  $s_{e,i,0} = x^i s_e$ 
3:   set  $k$  to be index of first non-zero element of  $s_{e,i,0}$ 
4:   if  $s_{e,i,0}[k] \neq 0$  then
5:     add  $s_{e,i,0}[k], P_{i,0}, s_{e,i,0}$  to  $\mathcal{I}[k], \mathcal{P}[k], \mathcal{S}[k]$  resp.
6:   for  $s = 1, \dots, e-1$  do
7:     for  $i = 0, \dots, d-1$  do
8:       set  $t = 0; P_{i,s}^{(t)} = yP_{i,s-1}$ ; and shift  $s_{e,i,s}^{(t)} = y \cdot s_{e,i,s-1}$ 
9:       if  $s_{e,i,s}^{(t)} = 0$  then continue to next  $i$ 
10:      set  $k$  to be the first non-zero index of  $s_{e,i,s}^{(t)}$ 
11:      if  $s_{e,i,s}^{(t)}[k] \notin \mathcal{I}[k]$  then continue to next  $i$ 
12:      solve for  $c_j$ 's such that  $s_{e,s,i}^{(t)}[k] - \sum_j c_j \mathcal{I}[k, j] = 0$ 
13:      set  $s_{e,i,s}^{(t+1)} = s_{e,i,s}^{(t)} - \sum_j c_j \mathcal{S}[k, j]$ 
14:      set  $P_{i,s}^{(t+1)} = P_{i,s}^{(t)} - \sum_j c_j \mathcal{P}[k, j]$ 
15:      go to line 9 with  $t = t + 1$ 
16:     for  $i = 0, \dots, d-1$  do
17:       set  $s_{e,i,s} = s_{e,i,s}^{(t)}$  and  $P_{i,s} = P_{i,s}^{(t)}$ 
18:       set  $k$  to be the index of first non-zero element of  $s_{e,i,s}$ 
19:       if  $s_{e,i,s}[k] \notin \mathcal{I}[k]$  then
20:         add  $s_{e,i,s}[k], P_{i,s}, s_{e,i,s}$  to  $\mathcal{I}[k], \mathcal{P}[k], \mathcal{S}[k]$  resp.
21:         reduce the basis of  $\mathcal{I}[k]$  if needed
22:     for  $i = 0, \dots, d-1$  do
23:       return  $P_{i,s}$  that makes  $s_{e,i,s} = 0$  for the first time
```

We can check membership $s_{e,i,s}[k] \in \mathcal{I}[k]$ and solve $s_{e,s,i}[k] - \sum c_j \mathcal{I}[k, j] = 0$ by finding the right approximant basis of

$$F = [\mathcal{I}[k, 0] \quad \dots \quad \mathcal{I}[k, n-1] \quad s_{e,s,i}[k]]$$

in Popov form. Since F has n rows and at most $n+1$ columns, we can compute this in cost $O(n^\omega d)$ [19]. The reduction in line 21 can be computed by the same approximant basis: if F has $n+1$ columns, there is a column in the approximant basis such that at least one entry has a nonzero constant term. By removing the corresponding $\mathcal{I}[k, j]$, we get a basis of $\mathcal{I}[k]$ of size n .

At lines 13 and 14, $\mathcal{S}[k, j]$ and $\mathcal{P}[k, j]$ have length and degree at most e resp., making the cost of these lines $O(n(ed)) = O(n^2 ed)$. Finally, using the fact that the total number of subiterations is bounded by $O(ed)$, we arrive at the total cost $O(ed(n^2 ed + n^\omega d))$.

We conclude by showing that the output of Algorithm 1 is indeed a basis of $\text{Ann}(s)$ and that it forms a lexicographic Gröbner basis.

THEOREM 3.1. *For each $i \in \{0, \dots, d-1\}$, let P_i be a canceling polynomial of s with leading coefficient x^i that is minimal in degree among all polynomials with leading coefficient x^i . Then one has $\text{Ann}(s) = \langle P_0, \dots, P_{d-1} \rangle$. Furthermore, $\{\varphi(P_0), \dots, \varphi(P_{d-1}), \alpha^d\}$ forms a Gröbner basis of $\bar{\varphi}(\text{Ann}(s))$ with respect to the lexicographic term order with $\alpha \leq_{\text{lex}} \beta$.*

PROOF. Suppose that there exists some $Q \in \mathbb{A}[y]$ with leading coefficient x^t that is in $\text{Ann}(s)$ but $Q \notin \langle P_0, \dots, P_{d-1} \rangle$. Note that for any polynomial in $\mathbb{A}[y]$, we can always make the leading coefficient to be some x^t by pulling out the minimal power of x from

the leading coefficient and multiplying by its inverse. Now, since we assumed minimality of degrees for P_i 's, $\deg(Q) > \deg(P_t)$ and $Q' = Q - y^{\deg Q - \deg P_t} P_t \in \text{Ann}(s)$ has degree less than Q . By normalizing the leading coefficient of Q' to be some $x^{t'}$, we can repeat the same process and keep decreasing the degree. This process must terminate when we encounter some Q' with leading coefficient $x^{t'}$ such that $\deg Q' < \deg P_{t'}$, or $Q' = 0$. Both cases lead to contradictions; thus, such Q cannot exist and $\text{Ann}(s) = \langle P_0, \dots, P_{d-1} \rangle$.

Next, let $\mathcal{G} = \{g_0, \dots, g_k\}$, $g_i \in \mathbb{K}[\alpha, \beta]$ with leading coefficient x^{d_i} , be the minimal reduced (lexicographic) Gröbner basis of $\bar{\varphi}(\text{Ann}(s))$. We can turn \mathcal{G} into another non-minimal Gröbner basis by adding the polynomials $\alpha^c g_i$, for $c = 1, \dots, d_{i+1} - 1$; we define the resulting basis as $\mathcal{G}' = \{g'_0, \dots, g'_d\}$, with $g'_d = \alpha^d$ and each g'_i has leading term $\alpha^i \beta^{r_i}$. Furthermore, define u_i as the degree of P_i such that $\varphi(P_i)$ has leading term $\alpha^i \beta^{u_i}$.

For $i = 0, \dots, d$, we have that $u_i \geq r_i$, otherwise \mathcal{G}' would not reduce $\varphi(P_i)$ to zero, which \mathcal{G}' must since $\varphi(P_i) \in \bar{\varphi}(\text{Ann}(s))$. We also have that $u_i \leq r_i$ due to the assumed minimality of degree for P_i 's. Thus, the leading terms of $\{\varphi(P_0), \dots, \varphi(P_{d-1}), \alpha^d\}$ generate the leading terms of $\bar{\varphi}(\text{Ann}(s))$. \square

4 LAZY ALGORITHM BASED ON KURAKIN'S

Kurakin's algorithm requires that we keep track of all d possible generators, regardless of the actual number of generators needed. For example, consider $s = (1, 1, 2, 3, 5, \dots) \in \mathbb{A}^{\mathbb{N}}$ with $\text{Ann}(s) = \langle y^2 - y - 1 \rangle$: Kurakin's algorithm returns $\{x^i(y^2 - y - 1), 0 \leq i < d\}$. In this section, we outline a modified version of Kurakin's algorithm that attempts to avoid as many extraneous computations as possible.

In the previous example, we can see that the polynomials associated with x^i , $i \geq 1$, were not useful. The next definition aims to qualify precisely the usefulness of the monomial x^i .

Definition 4.1. Let $P_{i,s}$ and $s_{e,i,s}$ be the polynomial and sequence at the end of step s associated with monomial x^i . A monomial x^{i_2} is useful wrt to x^{i_1} , $i_1 < i_2$, at step s if at least one of two conditions is true at the end of s :

U1. $P_{i_2,s} \neq x^{i_2-i_1} P_{i_1,s}$

U2. let k_{i_1} and k_{i_2} be the index of the first non-zero element of $s_{e,i_1,s}$ and $s_{e,i_2,s}$ resp., then $k_{i_1} \neq k_{i_2}$

Suppose a monomial x^{i_2} is not useful wrt x^{i_1} at step s , then by negating condition U1, we have $P_{i_2,s} = x^{i_2-i_1} P_{i_1,s}$. Due to negation of U2, $s_{e,i_2,s}$ is the zero sequence if and only if $s_{e,i_1,s}$ is the zero sequence; so either we return $P_{i_2,s} = x^{i_2-i_1} P_{i_1,s}$ or we do not terminate at this step for both monomials. Finally, since $k_{i_1} = k_{i_2}$ and $s_{e,i_2,s} = x^{i_2-i_1} s_{e,i_1,s}$, we always have that $s_{e,i_2,s}[k_{i_2}] = x^{i_2-i_1} s_{e,i_1,s}[k_{i_1}] \in (\langle s_{e,i_1,s}[k_{i_1}] \rangle \cup \mathcal{I}[k_{i_1}])$, meaning we can safely ignore $s_{e,i_2,s}[k_{i_2}]$ when updating $\mathcal{I}[k_{i_2}]$ at the end of step s . Thus, the negation of usefulness conditions U1 and U2 implies that any computation associated with x^{i_2} is not needed at step s .

However, as defined, U1 and U2 do not impose any conditions about the subiterations (indexed by t). The next lemma gives a different characterization of the usefulness conditions in terms of t .

LEMMA 4.2. *If x^{i_2} is useful wrt to x^{i_1} at some step s , then at some subiteration t of step s , one of $u1, u2, u3$ is true at the start of t :*

u1. $P_{i_2,s}^{(t)} \neq x^{i_2-i_1} P_{i_1,s}^{(t)}$

- u2. if $P_{i_2,s}^{(t)} = x^{i_2-i_1} P_{i_1,s}^{(t)}$, then $k_{i_2}^{(t)} \neq k_{i_1}^{(t)}$
u3. if $P_{i_2,s}^{(t)} = x^{i_2-i_1} P_{i_1,s}^{(t)}$ and $k_{i_2}^{(t)} = k_{i_1}^{(t)}$, then $s_{e,i_1,s}^{(t)}[k_{i_1}^{(t)}] \notin \mathcal{I}[k_{i_1}^{(t)}]$ and $s_{e,i_2,s}^{(t)}[k_{i_1}^{(t)}] \in \mathcal{I}[k_{i_1}^{(t)}]$

PROOF. We prove that if u1, u2, and u3 are false for every subiteration t and s , then U1 and U2 are false for x^{i_2} wrt x^{i_1} . Suppose the conditions u1, u2, and u3 are all false for every subiteration t at s . The negation of u1 forces $P_{i_2,s}^{(t)} = x^{i_2-i_1} P_{i_1,s}^{(t)}$ at the start of t , which sets the hypothesis of u2 true, implying $k_{i_2}^{(t)} = k_{i_1}^{(t)}$. Finally, since the hypothesis of u3 holds, we must have $s_{e,i_1,s}^{(t)}[k_{i_1}^{(t)}] \in \mathcal{I}[k_{i_1}^{(t)}]$ or $s_{e,i_2,s}^{(t)}[k_{i_1}^{(t)}] \notin \mathcal{I}[k_{i_1}^{(t)}]$. The two are mutually exclusive since $s_{e,i_2,s}^{(t)} = x^{i_2-i_1} s_{e,i_1,s}^{(t)}$, if $s_{e,i_1,s}^{(t)}[k_{i_1}^{(t)}] \in \mathcal{I}[k_{i_1}^{(t)}]$, then $s_{e,i_2,s}^{(t)}[k_{i_1}^{(t)}] \in \mathcal{I}[k_{i_1}^{(t)}]$. When $s_{e,i_1,s}^{(t)}[k_{i_1}^{(t)}] \in \mathcal{I}[k_{i_1}^{(t)}]$, we can update

$$P_{i_1,s}^{(t+1)} = P_{i_1,s}^{(t)} - \sum c_j \mathcal{I}[k_{i_1}^{(t)}, j]$$

$$P_{i_2,s}^{(t+1)} = x^{i_2-i_1} P_{i_1,s}^{(t)} - x^{i_2-i_1} \sum c_j \mathcal{P}[k_{i_1}^{(t)}, j] = x^{i_2-i_1} P_{i_1,s}^{(t+1)},$$

which was already implied by the assumption that u1 is false for all t . On the other hand, when $s_{e,i_2,s}^{(t)}[k_{i_1}^{(t)}] \notin \mathcal{I}[k_{i_1}^{(t)}]$, we also have $s_{e,i_1,s}^{(t)}[k_{i_1}^{(t)}] \notin \mathcal{I}[k_{i_1}^{(t)}]$, so the subiterations terminate and we must have $P_{i_2,s} = x^{i_2-i_1} P_{i_1,s}$ with $k_{i_2} = k_{i_1}$. This implies U1 and U2 also do not hold for step s . \square

While the converse is not true, we say a monomial x^{i_2} is *potentially useful* wrt x^{i_1} when at some step s and subiteration t , at least one of the conditions u1, u2, and u3 holds. Rather than iterating through $i = 0, \dots, d-1$, we keep a list of potentially useful monomials \mathcal{U} and iterate through $i \in \mathcal{U}$, with $\mathcal{U} = [0]$ initially. At each subiteration, we check to see if there exists $i' > i$, $i' \notin \mathcal{U}$ such that $x^{i'}$ satisfies one of u2 or u3, and add the smallest such i' to \mathcal{U} . Note that we need not check u1 since if u1 holds, then either u2 or u3 must have been true at some previous subiteration, thus i' is already included in \mathcal{U} . Condition u2 can be checked in $O(n)$ by checking the valuations of all entries in $s_{e,i,s}[k]$ at lines 4 and 10. Condition u3 can be checked in $O(\log d)$ membership computations via a binary search to find the minimal i' such that $x^{i'-i} s_{e,i,s}[k] \in \mathcal{I}[k]$ when $s_{e,i,s}[k] \notin \mathcal{I}[k]$ on line 11. Thus, the complexity for the subiterations do not change in terms of $O(\cdot)$. Defining $d^* = |\mathcal{U}| \leq d$, this brings the total cost to $O(ed^*(n^2ed + n^\omega d))$. While we do not know how far d^* is from the number d_{opt} of polynomials in the minimal lexicographic Gröbner basis of $\tilde{\varphi}(\text{Ann}(\mathbf{s}))$, we have observed empirically that d^* is often equal or close to d_{opt} (see Section 7).

5 VIA UNIVARIATE APPROXIMANT BASES

5.1 Approximants of a wide Hankel matrix

Extending the classical theory of linearly recurrent sequences over the field \mathbb{K} , another approach is to consider the left kernel of the block-Hankel matrix

$$H_{s,e} = \begin{bmatrix} S_0 & S_1 & \cdots & S_{e-1} \\ S_1 & S_2 & \cdots & S_e \\ \vdots & \ddots & \ddots & \vdots \\ S_e & S_{e+1} & \cdots & S_{2e-1} \end{bmatrix} \in \mathbb{A}^{(e+1) \times (en)}.$$

Indeed, if e is large enough, vectors in this kernel represent polynomials which cancel \mathbf{s} , and which even generate all of $\text{Ann}(\mathbf{s})$.

LEMMA 5.1. Let $\mathbf{s} \in \mathcal{S}$ be linearly recurrent of order δ , and define

$$\mathcal{K}_{s,e} = \{p = p_0 + \cdots + p_e y^e \in \mathbb{A}[y] \mid [p_0 \cdots p_e] H_{s,e} = 0\}$$

for $e \in \mathbb{N}$. Assume $e \geq \delta$. Then $\mathcal{K}_{s,e} = \text{Ann}(\mathbf{s}) \cap \mathbb{A}[y]_{\leq e}$, and in particular $\mathcal{K}_{s,e}$ is a generating set of $\text{Ann}(\mathbf{s})$.

PROOF. Let $p = p_0 + \cdots + p_e y^e \in \mathbb{A}[y]$ and $\gamma = \deg(p) \leq e$. Then $p \in \mathcal{K}_{s,e}$ if and only if $[p_0 \cdots p_e] H_{s,e} = 0$, and by definition of canceling partial sequences this exactly means that p cancels $\mathbf{s}_{e+\gamma}$. Now, $\deg(p) = \gamma \leq e + \gamma - \delta$ holds under the assumption $e \geq \delta$, hence p cancels $\mathbf{s}_{e+\gamma}$ if and only if $p \in \text{Ann}(\mathbf{s})$ by Lemma 2.3. It follows that $\mathcal{K}_{s,e}$ generates $\text{Ann}(\mathbf{s})$, since there exists a generating set of $\text{Ann}(\mathbf{s})$ whose polynomials all have degree at most δ . \square

Computing the left kernel of $H_{s,e}$ can be done via univariate approximation. Indeed, calling $F \in \mathbb{K}[x]^{(e+1) \times (en)}$ the natural lifting of $H_{s,e}$, an approximant basis of F at order d gives a generating set of that left kernel. As recalled in Section 2.3, using PM-BASIS, a basis of $\mathcal{A}_d(F)$ in shifted reduced or Popov form can be computed in $O(e^{\omega-1}(e+en)d) = O(e^\omega nd)$ operations in \mathbb{K} .

5.2 Speed-up by compression using structure

Now we show that, when n is large, one can speed up the above approach by a randomized ‘‘compression’’ of the matrix $H_{s,e}$. Precisely, taking a random constant matrix $C \in \mathbb{K}^{(en) \times (e+1)}$ and performing the right-multiplication FC , one obtains a square $(e+1) \times (e+1)$ matrix such that $\mathcal{A}_d(F) = \mathcal{A}_d(FC)$ holds with good probability. The cost of the approximant basis computation is thus reduced to $O(e^\omega d)$ operations in \mathbb{K} , and the right-multiplication can be done efficiently by leveraging the block-Hankel structure of F .

THEOREM 5.2. Algorithm 2 takes as input an integer $d \in \mathbb{Z}_{>0}$, vectors $F_0, \dots, F_{\mu+e-2} \in \mathbb{K}[x]^{1 \times n}$ of degree less than d , and a shift $w \in \mathbb{Z}_{>0}^\mu$, and uses $O(\mu e n d + \mu^\omega d)$ operations in \mathbb{K} to compute a w -Popov matrix $P \in \mathbb{K}[x]^{\mu \times \mu}$ of degree at most d . It chooses at most $\mu e n$ elements independently and uniformly at random from a subset of \mathbb{K} of cardinality κ , and P is the w -Popov basis of $\mathcal{A}_d(F)$ with probability at least $1 - \frac{\mu}{\kappa}$, where F is the block-Hankel matrix

$$F = \begin{bmatrix} F_0 & F_1 & \cdots & F_{e-1} \\ F_1 & F_2 & \cdots & F_e \\ \vdots & \ddots & \ddots & \vdots \\ F_{\mu-1} & F_\mu & \cdots & F_{\mu+e-2} \end{bmatrix} \in \mathbb{K}[x]^{\mu \times (en)}. \quad (3)$$

When applied to the computation of $\text{Ann}(\mathbf{s})$ with $\mu = e+1$, the cost becomes $O(e^2 n d + e^\omega d)$. Below we focus on the case of interest $\mu \leq en$, since when $en \in O(\mu)$ this w -Popov approximant basis is computed deterministically by PM-BASIS at a cost of $O(\mu^\omega d)$ operations in \mathbb{K} . Our approach is based on the following two lemmas.

LEMMA 5.3. Let $F \in \mathbb{K}[x]^{\mu \times \nu}$ and $d \in \mathbb{Z}_{>0}$. Let $C \in \mathbb{K}[x]^{\nu \times r}$ and $K \in \mathbb{K}[x]^{\nu \times (\nu-r)}$, for some $r \in \{0, \dots, \nu\}$, such that $FK = 0$ and $[C(0) \ K(0)] \in \mathbb{K}^{\nu \times \nu}$ is invertible. Then, $r \geq \rho$ where ρ is the rank of F , and $\mathcal{A}_d(F) = \mathcal{A}_d(FC)$.

PROOF. Let $N = [C \ K] \in \mathbb{K}[x]^{\nu \times \nu}$. The assumption that $N(0)$ is invertible ensures that N is nonsingular (since $\det(N)(0) = \det(N(0)) \neq 0$), and therefore K has full rank $\nu - r$. The assumption that the columns of K are in the right kernel of F , which has rank $\nu - \rho$, implies that $\nu - r \leq \nu - \rho$ and therefore $r \geq \rho$.

The inclusion $\mathcal{A}_d(F) \subset \mathcal{A}_d(FC)$ is obvious. For the other inclusion, let $p \in \mathcal{A}_d(FC)$, i.e. there exists $q \in \mathbb{K}[x]^{1 \times r}$ such that $pFC = x^d q$. It follows that $pFN = x^d [q \ 0]$, and thus

$$pF = x^d [q \ 0] N^{-1} = \frac{x^d [q \ 0] \text{Adj}(N)}{\det(N)}$$

where $\text{Adj}(N) \in \mathbb{K}[x]^{\nu \times \nu}$ is the adjugate of N . Our assumption $\det(N)(0) \neq 0$ means that x^d and $\det(N)$ are coprime, hence $\det(N)$ divides $[q \ 0] \text{Adj}(N)$, and $pF = 0 \pmod{x^d}$ follows. \square

LEMMA 5.4. *Let $F \in \mathbb{K}[x]^{\mu \times \nu}$ with rank ρ and $\mu \leq \nu$, and let $r \in \{\rho, \dots, \mu\}$. Let \mathcal{R} be a finite subset of \mathbb{K} of cardinality $\kappa \in \mathbb{Z}_{>0}$, and let $C \in \mathbb{K}^{\nu \times r}$ with entries chosen independently and uniformly at random from \mathcal{R} . Then, the probability that there exists $K \in \mathbb{K}[x]^{\nu \times (\nu-r)}$ such that $[C \ K(0)]$ is invertible and $FK = 0$ is at least $1 - \frac{r}{\kappa}$; furthermore if \mathbb{K} is finite and $\mathcal{R} = \mathbb{K}$, this probability is at least $\prod_{i=1}^r (1 - \kappa^{-i})$.*

PROOF. Consider a right kernel basis $B \in \mathbb{K}[x]^{\nu \times (\nu-\rho)}$ for F . Then B has unimodular row bases [41, Lem. 3.1], implying that there exists $V \in \mathbb{K}[x]^{(\nu-\rho) \times \nu}$ such that $VB = I_{\nu-\rho}$. In particular $V(0)B(0) = I_{\nu-\rho}$ and therefore $B(0)$ has full rank $\nu - \rho$. Define $K \in \mathbb{K}[x]^{\nu \times (\nu-r)}$ as the matrix formed by the first $\nu - r$ columns of B (recall $\nu - r \leq \nu - \rho$ by assumption). Then $FK = 0$. Furthermore $K(0)$ has rank $\nu - r$, hence the DeMillo-Lipton-Schwartz-Zippel lemma implies that $[C \ K(0)] \in \mathbb{K}^{\nu \times \nu}$ is singular with probability at most r/κ [11, 36, 42]. If \mathbb{K} is finite and $\mathcal{R} = \mathbb{K}$ then $[C \ K(0)]$ is invertible with probability exactly $\prod_{i=1}^r (1 - \kappa^{-i})$. \square

These lemmas lead to Algorithm 2 and Theorem 5.2; indeed computing FC has quasi-linear cost $O(\mu \text{end})$ thanks to the block-Hankel structure of F , and then the call $\text{PM-BASIS}(d, FC, w)$ costs $O(\mu^\omega d)$ operations as recalled in Section 2.3.

Algorithm 2 HANKEL-PM-BASIS(d, F, w)

Input: integers $d, \mu, e, n \in \mathbb{Z}_{>0}$, vectors $F_0, \dots, F_{\mu+e-2} \in \mathbb{K}[x]^{1 \times n}$ of degree less than d , a shift $w \in \mathbb{Z}_{>0}^\mu$

Output: a w -Popov matrix $P \in \mathbb{K}[x]^{\mu \times \mu}$ of degree at most d

- 1: $F \in \mathbb{K}[x]^{\mu \times (en)} \leftarrow$ form the block-Hankel matrix as in Eq. (3)
 - 2: **if** $\mu \geq en$ **then return** $\text{PM-BASIS}(d, F, w)$
 - 3: Choose $r \in \{\rho, \dots, \mu\}$ where ρ is the rank of F (by default, choose $r = \mu$ if no information is known on ρ)
 - 4: Fill a matrix $C \in \mathbb{K}^{(en) \times r}$ with entries chosen uniformly and independently at random from a subset of \mathbb{K} of cardinality κ
 - 5: Compute $FC \in \mathbb{K}[x]^{\mu \times r}$ (exploiting the Hankel structure of F)
 - 6: **return** $\text{PM-BASIS}(d, FC, w)$
-

Note that $1 - r/\kappa \geq 3/4$ as soon as $\kappa \geq 4\mu$ (which implies $\kappa \geq 4r$); furthermore $\prod_{i=1}^r (1 - \kappa^{-i}) \geq 3/4$ already for $\kappa = 7$. The randomization is of the Monte Carlo type, since the algorithm may return P which is not a basis of $\mathcal{A}_d(F)$. Still, since the expected w -Popov basis P of $\mathcal{A}_d(F)$ is unique, one can easily increase the probability of success by repeating the randomized computation and following

a majority rule. Another approach is to rely on the non-interactive, Monte Carlo certification protocol of [15], which has lower cost than Algorithm 2 but requires a larger field \mathbb{K} ; this first asks to compute the coefficient of degree d of PF , which here can be done via bivariate polynomial multiplication in time $O(\mu \text{end})$ thanks to the structure of F . For a given output P , this certification can be repeated for better confidence in P (in which case the coefficient of degree d of PF needs only be computed once).

6 VIA BIVARIATE PADÉ APPROXIMATION

Now, we propose another approach which directly uses the interpretation of canceling polynomials as denominators of the generating series of the sequence (see Lemma 2.2). The next lemma describes more precisely the link between the annihilator and these denominators when we have access to a partial sequence, that is, denominators of the generating series truncated at some order. One can also view this lemma as a description of the kernel of the univariate Hankel matrix $H_{s,e}$ via bivariate Padé approximation.

LEMMA 6.1. *Let $s \in \mathcal{S}$ be linearly recurrent of order δ , and for $e \in \mathbb{N}$ define $G = \sum_{j < 2e} S_j y^{2e-1-j} \in \mathbb{A}[y]^n$ and*

$$\mathcal{P}_{s,e} = \{p \in \mathbb{A}[y]_{\leq e} \mid pG = q \pmod{y^{2e}} \text{ for some } q \in \mathbb{A}[y]_{< e}^n\}.$$

Assume $e \geq \delta$. Then $\mathcal{P}_{s,e} = \text{Ann}(s) \cap \mathbb{A}[y]_{\leq e}$, and in particular $\mathcal{P}_{s,e}$ is a generating set of $\text{Ann}(s)$; furthermore for any $p \in \mathcal{P}_{s,e}$ the corresponding $q \in \mathbb{A}[y]_{< e}^n$ satisfies $\deg(q) < \deg(p)$.

PROOF. Let $p = p_0 + \dots + p_\gamma y^\gamma \in \mathbb{A}[y]_{\leq e}$ where $\gamma = \deg(p)$. Then $p \in \mathcal{P}_{s,e}$ if and only if the coefficient of pG of degree $2e - 1 - k$ is zero for $0 \leq k < e$. Since $\gamma \leq e \leq 2e - 1 - k$, this coefficient is

$$\text{Coeff}(pG, 2e - 1 - k) = \sum_{i=0}^{\gamma} p_i S_{2e-1-(2e-1-k-i)} = \sum_{i=0}^{\gamma} p_i S_{k+i} = 0.$$

Thus we have proved $\mathcal{P}_{s,e} = \mathcal{K}_{s,e}$, and Lemma 5.1 shows the claims in this lemma except the last one. Let $p \in \mathcal{P}_{s,e}$ and define q as the polynomial in $\mathbb{A}[y]_{< e}^n$ such that $pG = q \pmod{y^{2e}}$. Since $p \in \text{Ann}(s)$, Lemma 2.2 shows that pG_s is a polynomial. On the other hand the definitions of G and G_s yield $pG = y^{2e} pG_s - p \sum_{j \geq 2e} S_j y^{2e-1-j}$. Hence $-p \sum_{j \geq 2e} S_j y^{2e-1-j}$ is a polynomial, and since it has degree less than γ , and thus in particular less than $2e$, it is equal to q . \square

From G , define $F \in \mathbb{K}[\alpha, \beta]^{1 \times n}$ of bi-degree less than $(d, 2e)$ via the morphism φ from Section 2.2. Equip $\mathbb{K}[\alpha, \beta]$ with the lexicographic order \leq_{lex} , and let \leq be the corresponding term over position order on $\mathbb{K}[\alpha, \beta]^{n+1}$. Then a minimal \leq -Gröbner basis of the submodule of simultaneous Padé approximants

$$\{(p, q) \in \mathbb{K}[\alpha, \beta] \times \mathbb{K}[\alpha, \beta]^{1 \times n} \mid pF = q \pmod{(x^d, y^{2e})}\}$$

is computed in $O((n^\omega \min(d, e)^\omega + n^3 \min(d, e)^2)de)$ operations, using the algorithm of [28] (see also Section 2.3) with input matrix of size $(n+1) \times n$ formed by stacking the identity I_n below F . Lemma 6.1 shows that from this \leq -Gröbner basis one can find a minimal \leq_{lex} -Gröbner basis of $\bar{\varphi}(\text{Ann}(s))$ by selecting p for each (p, q) in the basis such that $\deg_\beta(q) < \deg_\beta(p)$.

While the PM-BASIS approach had cost quasi-linear in d and n , the method here is most efficient in an opposite parameter range: for $n \in O(1)$ and $d \leq e$ the above cost bound becomes $O(d^{\omega+1}e)$.

7 EXPERIMENTAL RESULTS

In this section, we compare timings for the algorithms in Sections 3 to 5, implemented in C++ using the libraries NTL [37] and PML [18] which provide high-performance support for univariate polynomials and polynomial matrices. We leave the implementation of the bivariate algorithm of Section 6 as future work. To control the cardinality and shape of the Gröbner basis, we use Lazard’s structural theorem (see Section 2.2). The shape of the monomial staircase is randomized with maximal β -degree δ and α^d included in the basis. After generating a random Gröbner basis \mathcal{G} of target degree and size, we use it to generate n sequences (with $e = 2\delta$ terms), using random initial conditions. Finally, we compute the annihilator of the sequence, which may not necessarily recover \mathcal{G} itself (see Section 8.1). Runtimes are showed below.

n	d	δ	d_{opt}	$D/d\delta$	K	LK	d^*	PM-B	HPM
1	64	256	1	1	62.8	0.93	1	1.06	NA
1	64	256	49	0.62	38.0	1.65	53	2.10	NA
1	128	512	16	0.92	>100	12	17	20.5	NA
1	128	32	12	0.91	7.85	0.078	12	0.029	NA
1	256	32	14	0.94	27.3	0.12	14	0.08	NA
1	256	128	27	0.92	>100	1.28	27	1.60	NA
1	512	256	29	0.96	>100	8.65	29	27.8	NA
2	17	256	2	0.5	14.1	0.91	2	0.33	0.29
3	12	512	4	0.4	6.93	1.40	4	2.47	1.86
8	16	256	1	1	54.1	3.16	1	0.56	0.25
32	16	256	1	1	>100	39.8	1	2.79	0.35
64	16	128	1	1	>100	>100	1	1.02	0.13

Table: Runtimes, in seconds, of algorithms Kurakin, Lazy Kurakin, direct PM-BASIS, and HANKEL-PM-BASIS, observed on AMD Ryzen 5 3600X 6-Core CPU with 16 GB RAM. The base field is $\mathbb{K} = \mathbb{F}_{9001}$.

As we claim in Section 4, d^* is often close or equal to d_{opt} . More interestingly, Lazy Kurakin outperforms Kurakin more than d/d^* would suggest. For example, for $\delta = 256, d = 64, d_{\text{opt}} = 49$, then $d/d^* \approx 1.2$ but Kurakin is 23 times slower than Lazy Kurakin. This is because the cost bound $O(ed^*(n^2ed + n^\omega d))$ for Lazy Kurakin assumes that d^* polynomials are tracked from the beginning of the algorithms. However, due to its lazy nature, polynomials are often added later in the algorithm and the bound of ed^* subiterations may significantly overestimate the true number of subiterations.

When δ, d, n are fixed, Kurakin’s algorithm performs worse for $d_{\text{opt}} = 1$ than $d_{\text{opt}} > 1$, although this is a favourable case for Lazy Kurakin. In this case, Kurakin’s algorithm computes $P_i = x^i P_0$ so there cannot be any early termination. Additionally, the size of the staircase is maximal ($D = ed$), so this is also the worst case for algorithms whose complexity depends directly on D . Lazy Kurakin’s algorithm somewhat remedies this by using the extra structure of \mathbb{A} and adding monomials in a lazy fashion. (When it is known that $\text{Ann}(s) = \langle P \rangle$, it is possible to design an algorithm that is quasilinear in e via structured system solving, see Section 8.2).

For scalar sequences over \mathbb{A} , i.e. $n = 1$, Lazy Kurakin’s algorithm seems to be the best choice when δ is large compared to d , whereas PM-Basis seems to be the best choice in the converse. When $e = 2\delta = d$, Lazy Kurakin outperforms PM-Basis, given that d^* is small. This is predicted by the theoretical complexities, as the former has complexity $O(e^3 d^*)$, while the latter has complexity $O(e^{\omega+1})$.

For $n > 1$, PM-BASIS and HANKEL-PM-BASIS clearly outperform Kurakin and Lazy Kurakin. This is as predicted since the complexity

of the former depends linearly on n , while the latter has a factor n^ω . The theoretical improvement of HANKEL-PM-BASIS over PM-BASIS is observed empirically, especially for the two cases of $n = 32, 64$.

8 APPLICATIONS TO SPARSE MATRICES

In this section, we outline two applications to sparse matrices $A \in \mathbb{A}^{n \times n}$: first, the computation of minimal polynomials of A , which are polynomials of minimal degree that cancel the matrix sequence $s_A = (A^0, A^1, A^2, \dots)$; second, the computation of the determinant of A . In what follows, we assume A has sparsity $O(n)$, i.e. it has $O(n)$ nonzero entries, and that the representation of A allows us to compute matrix-vector products at cost $O(nd)$. Our approach is based on Wiedemann’s [39], designed for matrices over fields.

8.1 Minimal polynomials of sparse matrices

Given a matrix A , the well-known Cayley-Hamilton theorem states that A cancels its own characteristic polynomial. This implies that the sequence of successive powers of A is linearly recurrent, and a polynomial of minimal degree that cancels this sequence is said to be a minimal polynomial of A . A different view one can take is that such canceling polynomials must cancel the n^2 linearly recurrent sequences $((A^i)_{j_1, j_2})_{i \geq 0}$ simultaneously for $1 \leq j_1, j_2 \leq n$. Then, as usual, we want to compute a Gröbner basis of the ideal of these canceling polynomials, denoted by $\text{Ann}(A)$.

Over \mathbb{A} , trying to deduce $\text{Ann}(A)$ from $\text{Ann}((u^T A^i v)_{i \geq 0})$, for random vectors $u, v \in \mathbb{A}^{n \times 1}$, presents a problem when $\text{Ann}(A)$ does not have the Gorenstein property [16, 25]. When $\text{Ann}(A)$ has the Gorenstein property, it has been showed that $\text{Ann}(A)$ can be recovered, with high probability, by using a bidimensional sequence with random initial conditions, provided \mathbb{K} has large characteristic [5]. When it does not have the property, $\text{Ann}(A)$ is still recoverable with a similar approach, but using several sequences [29]. Over various commutative rings, the problem of computing minimal polynomials of a matrix have been studied in [8, 17, 32]. However, the algorithms given in these works do not exploit sparsity.

Given matrix A as above, we start by choosing random $u_1, v \in \mathbb{A}^n$ and generating $s_{A,1} = (u_1^T A^i v)_{0 \leq i < 2n}$. Next, we apply one of the algorithms in the previous sections to compute $\text{Ann}(s_{A,1})$. If $\text{Ann}(s_{A,1}) = \text{Ann}(A)$, which can be checked probabilistically by checking if $\text{Ann}(s_{A,1})$ also cancels some validation sequence $((u')^T A^i v)_{0 \leq i < 2n}$, we terminate the process. Otherwise, we double the number of sequences by doubling the number of random u_i ’s and generating $s_{A,1}, \dots, s_{A,2^s}$. The cost of the process is $O(\tau n^2 d + \mathcal{L}(n, d, \tau))$, where τ is the number of sequences used and $\mathcal{L}(n, d, \tau)$ is the cost of finding the annihilators of a partial sequence of length n in $(\mathbb{K}[x]/\langle x^d \rangle)^\tau$. Note that this process must terminate. The crudest bound is when $\tau > n^2$ since then we could simply compute $\text{Ann}(A)$ directly. Another slightly more refined bound for the number of generic linear forms needed is $\tau \leq D$, where D is the size of the staircase of $\text{Ann}(A)$ [29, Prop. 1].

8.2 Determinant of sparse matrices

The determinant of a matrix is easily obtained from its minimal polynomial when the latter is equal to the characteristic polynomial. Wiedemann [39] calls such matrices *nonderogatory* and shows that preconditioning any matrix $B \in \mathbb{K}^{n \times n}$ with a random diagonal

matrix D results in a nonderogatory matrix with high probability. We will show that the same preconditioning can be applied to matrices over \mathbb{A} . Here, a particular role will be played by sequences $\mathbf{s} \in (\mathbb{A}^n)^{\mathbb{N}}$ such that $\text{Ann}(\mathbf{s}) = \langle P \rangle$, for some monic $P \in \mathbb{A}[y]$. Indeed, the next theorem shows that it is sufficient for the constant part of A to be nonderogatory in \mathbb{K} for A to be nonderogatory in \mathbb{A} and for the sequence of its powers to satisfy this property.

THEOREM 8.1. *Let $A_0 \in \mathbb{K}^{n \times n}$ be the constant part of A (i.e. for $x = 0$). If A_0 is nonderogatory, then $\text{Ann}(A) = \langle P \rangle$ for some monic $P \in \mathbb{A}[y]$ of degree n .*

PROOF. Let $P \in \mathbb{A}[y]$ be the minimal monic polynomial of the sequence $\mathbf{s}_A = (A^0, A^1, A^2, \dots)$, then $\deg(P) \leq n$ since A is $n \times n$. Now, A_0 is nonderogatory, so any canceling polynomial must have degree $\geq n$; thus, $\deg(P) = n$. Furthermore, if there exists another polynomial Q of degree n and leading coefficient x^i such that $Q \neq x^i P$, then $Q - x^i P$ is a canceling polynomial of degree less than n , contradicting the previous statement. Thus, $P, xP, \dots, x^{d-1}P$ are minimal in degree and, by Theorem 3.1, $\text{Ann}(A) = \langle P, xP, \dots, x^{d-1}P \rangle = \langle P \rangle$. \square

The above theorem allows us to use the same preconditioner as in [39]: a random constant diagonal matrix D . The preconditioning ensures that the ideal of canceling polynomial is generated by a single monic polynomial; thus, $\bar{\varphi}(\text{Ann}(AD))$ is Gorenstein and requires only a single linear form to be recovered. Furthermore, when it is known that the ideal is generated by a single polynomial, we can recover this polynomial in $\mathcal{O}(nd)$ by taking advantage of the fact that the constant part of the leading $n \times n$ submatrix of $H_{\mathbf{s}, 2n}$ is an invertible Hankel matrix [7]. Once we have P , we can compute $\det(A) = P(0)(\prod_i D_{i,i})^{-1}$. Under our sparsity assumption, the cost of this method is $\mathcal{O}(n^2d)$ for computing $(u^T A^i v)_{i \leq 2n}$, $\mathcal{O}(nd)$ for computing P , and $\mathcal{O}(n+d)$ for recovering the determinant from P , leading to the total cost of $\mathcal{O}(n^2d)$ operations in \mathbb{K} . This is to be compared with computing the determinant of A “at full precision”, i.e. by seeing A as a matrix over $\mathbb{K}[x]$, and then truncating the result modulo x^d : this costs $\mathcal{O}(n^\omega d)$ operations in \mathbb{K} [22].

REFERENCES

- [1] J. Alman and V. Vassilevska Williams. 2021. A Refined Laser Method and Faster Matrix Multiplication. In *Proceedings SODA 2021*. 522–539. <https://doi.org/10.1137/1.9781611976465.32>
- [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 (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 *ISSAC'99*. ACM, 189–196. <https://doi.org/10.1145/309831.309929>
- [4] E. Berlekamp. 1968. Nonbinary BCH decoding (Abstr.). *IEEE Trans. Inf. Theory* 14, 2 (1968), 242–242. <https://doi.org/10.1109/TIT.1968.1054109>
- [5] J. Berthomieu, B. Boyer, and J.-C. Faugère. 2017. Linear algebra for computing Gröbner bases of linear recursive multidimensional sequences. *J. Symb. Comput.* 83 (2017), 36–67. <https://doi.org/10.1016/j.jsc.2016.11.005>
- [6] J. Berthomieu and J.-C. Faugère. 2018. A Polynomial-Division-Based Algorithm for Computing Linear Recurrence Relations. In *ISSAC'18*. 79–86. <https://doi.org/10.1145/3208976.3209017>
- [7] A. Bostan, C.-P. Jeannerod, and É. Schost. 2008. Solving structured linear systems with large displacement rank. *Theor. Comput. Sci.* 407, 1 (2008), 155–181. <https://doi.org/10.1016/j.tcs.2008.05.014>
- [8] W. C. Brown. 2005. Null Ideals of Matrices. *Communications in Algebra* 33, 12 (2005), 4491–4504. <https://doi.org/10.1080/00927870500274820>
- [9] D. Coppersmith and S. Winograd. 1990. Matrix multiplication via arithmetic progressions. *J. Symb. Comput.* 9, 3 (1990), 251–280. [https://doi.org/10.1016/S0747-7171\(08\)80013-2](https://doi.org/10.1016/S0747-7171(08)80013-2)
- [10] D. A. Cox, J. Little, and D. O’Shea. 2005. *Using Algebraic Geometry (second edition)*. Springer-Verlag New-York, New York, NY. <https://doi.org/10.1007/b138611>
- [11] R. A. DeMillo and R. J. Lipton. 1978. A Probabilistic Remark on Algebraic Program Testing. *Inform. Process. Lett.* 7, 4 (1978), 193–195.
- [12] P. Fitzpatrick. 1997. Solving a Multivariable Congruence by Change of Term Order. *J. Symb. Comput.* 24, 5 (1997), 575–589. <https://doi.org/10.1006/jsc.1997.0153>
- [13] P. Fitzpatrick and G. H. Norton. 1990. Finding a basis for the characteristic ideal of an n -dimensional linear recurring sequence. *IEEE Trans. Inf. Theory* 36, 6 (1990), 1480–1487. <https://doi.org/10.1109/18.59953>
- [14] P. Giorgi, C.-P. Jeannerod, and G. Villard. 2003. On the complexity of polynomial matrix computations. In *ISSAC'03*. ACM, 135–142. <https://doi.org/10.1145/860854.860889>
- [15] P. Giorgi and V. Neiger. 2018. Certification of Minimal Approximant Bases. In *ISSAC'18*. ACM, 167–174. <https://doi.org/10.1145/3208976.3208991>
- [16] W. Gröbner. 1935. Über irreduzible Ideale in kommutativen Ringen. *Math. Ann.* 110, 1 (1935), 197–222.
- [17] C. Heuberger and R. Rissner. 2017. Computing J-ideals of a matrix over a principal ideal domain. *Linear Algebra Appl.* 527 (2017), 12–31. <https://doi.org/10.1016/j.laa.2017.03.028>
- [18] S. G. Hyun, V. Neiger, and É. Schost. 2019. Implementations of Efficient Univariate Polynomial Matrix Algorithms and Application to Bivariate Resultants. In *ISSAC'19*. ACM, 235–242. <https://doi.org/10.1145/3326229.3326272>
- [19] C.-P. Jeannerod, V. Neiger, and G. Villard. 2020. Fast computation of approximant bases in canonical form. *J. Symb. Comput.* 98 (2020), 192–224. <https://doi.org/10.1016/j.jsc.2019.07.011>
- [20] V. L. Kurakin. 1998. The Berlekamp–Massey algorithm over finite rings, modules, and bimodules. *Discrete Mathematics and Applications* 8, 5 (1998), 441–474.
- [21] V. L. Kurakin. 2000. Construction of the Annihilator of a Linear Recurring Sequence over Finite Module with the help of the Berlekamp–Massey Algorithm. In *FPSAC 2000*. Springer, 476–483. https://doi.org/10.1007/978-3-662-04166-6_45
- [22] G. Labahn, V. Neiger, and W. Zhou. 2017. Fast, deterministic computation of the Hermite normal form and determinant of a polynomial matrix. 42 (2017), 44–71. <https://doi.org/10.1016/j.jco.2017.03.003>
- [23] D. Lazard. 1985. Ideal Bases and Primary Decomposition: Case of Two Variables. *J. Symb. Comput.* 1, 3 (1985), 261–270.
- [24] F. Le Gall. 2014. Powers of Tensors and Fast Matrix Multiplication. In *ISSAC'14* (Kobe, Japan). ACM, 296–303. <https://doi.org/10.1145/2608628.2608664>
- [25] F. S. Macaulay. 1934. Modern algebra and polynomial ideals. In *Math. Proc. Camb. Philos. Soc.*, Vol. 30. Cambridge University Press, 27–46.
- [26] J. Massey. 1969. Shift-register synthesis and BCH decoding. *IEEE Trans. Inf. Theory* 15 (1969), 122–127.
- [27] B. Mourrain. 2017. Fast Algorithm for Border Bases of Artinian Gorenstein Algebras. In *ISSAC'17* (Kaiserslautern, Germany). ACM, 333–340. <https://doi.org/10.1145/3087604.3087632>
- [28] S. Naldi and V. Neiger. 2020. A Divide-and-Conquer Algorithm for Computing Gröbner Bases of Syzygies in Finite Dimension. In *ISSAC'20*. ACM, 380–387. <https://doi.org/10.1145/3373207.3404059>
- [29] V. Neiger, H. Rahkooy, and É. Schost. 2017. Algorithms for zero-dimensional ideals using linear recurrent sequences. In *CASC 2017*. Springer, 313–328.
- [30] V. Neiger and É. Schost. 2020. Computing syzygies in finite dimension using fast linear algebra. *J. Complexity* 60 (2020), 101502. <https://doi.org/10.1016/j.jco.2020.101502>
- [31] V. M. Popov. 1972. Invariant Description of Linear, Time-Invariant Controllable Systems. *SIAM Journal on Control* 10, 2 (1972), 252–264.
- [32] R. Rissner. 2016. Null ideals of matrices over residue class rings of principal ideal domains. *Linear Algebra Appl.* 494 (2016), 44–69. <https://doi.org/10.1016/j.laa.2016.01.004>
- [33] S. Sakata. 1988. Finding a minimal set of linear recurring relations capable of generating a given finite two-dimensional array. *J. Symb. Comput.* 5, 3 (1988), 321–337. [https://doi.org/10.1016/S0747-7171\(88\)80033-6](https://doi.org/10.1016/S0747-7171(88)80033-6)
- [34] S. Sakata. 1990. Extension of the Berlekamp–Massey algorithm to N dimensions. *Information and Computation* 84, 2 (1990), 207–239.
- [35] S. Sakata. 2009. The BMS Algorithm. In *Gröbner Bases, Coding, and Cryptography*. Springer, 143–163. https://doi.org/10.1007/978-3-540-93806-4_9
- [36] J. T. Schwartz. 1980. Fast Probabilistic Algorithms for Verification of Polynomial Identities. *J. ACM* 27, 4 (1980), 701–717. <https://doi.org/10.1145/322217.322225>
- [37] V. Shoup. 2020. NTL: A Library for doing Number Theory, version 11.4.3. <http://www.shoup.net>.
- [38] 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>
- [39] D. Wiedemann. 1986. Solving sparse linear equations over finite fields. *IEEE Trans. Inf. Theory* 32, 1 (1986), 54–62. <https://doi.org/10.1109/TIT.1986.1057137>
- [40] W. A. Wolovich. 1974. *Linear Multivariable Systems*. Applied Mathematical Sciences, Vol. 11. Springer-Verlag New-York.
- [41] W. Zhou and G. Labahn. 2013. Computing Column Bases of Polynomial Matrices. In *ISSAC'13*. ACM, 379–386. <https://doi.org/10.1145/2465506.2465947>
- [42] R. Zippel. 1979. Probabilistic algorithms for sparse polynomials. In *EUROSAM'79 (LNCS)*, Vol. 72. Springer, 216–226.