Implementations of efficient univariate polynomial matrix algorithms and application to bivariate resultants

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

Complexity bounds for many problems on matrices with univariate polynomial entries have been improved in the last few years. Still, for most related algorithms, efficient implementations are not available, which leaves open the question of the practical impact of these algorithms, e.g. on applications such as decoding some error-correcting codes and solving polynomial systems or structured linear systems. In this paper, we discuss implementation aspects for most fundamental operations: multiplication, truncated inversion, approximants, interpolants, kernels, linear system solving, determinant, and basis reduction. We focus on prime fields with a word-size modulus, relying on Shoup's C++ library NTL. Combining these new tools to implement variants of Villard's algorithm for the resultant of generic bivariate polynomials (ISSAC 2018), we get better performance than the state of the art for large parameters.

Keywords

Polynomial matrices, algorithms, implementation, resultant.

ACM Reference Format:

Seung Gyu Hyun, Vincent Neiger, and Éric Schost. 2019. Implementations of efficient univariate polynomial matrix algorithms and application to bivariate resultants. In *International Symposium on Symbolic and Algebraic Computation (ISSAC '19)*, July 15–18, 2019, Beijing, China. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3326229.3326272>

1 Introduction

Hereafter, $\mathbb{K}$ is a field and $\mathbb{K}[x]$ is the algebra of univariate polynomials over $\mathbb{K}$. Recent years have witnessed a host of activity on fast algorithms for polynomial matrices and their applications:

- Minimal approximant bases [15, 53] were used to compute kernel bases [54], giving the first efficient deterministic algorithm for linear system solving over $\mathbb{K}[x]$.
- Basis reduction [15, 16] played a key role in accelerating the decoding of one-point Hermitian codes [35] and in designing deterministic determinant and Hermite form algorithms [29].
- Progress on minimal interpolant bases [23, 24] led to the best known complexity bound for list-decoding Reed-Solomon codes and folded Reed-Solomon codes [24, Sec. 2.4 to 2.7].
- Coppersmith's block Wiedemann algorithm and its extensions [7, 26, 48] were used in a variety of contexts, from integer factorization [44] to polynomial system solving [22, 49].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ISSAC '19, July 15–18, 2019, Beijing, China

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6084-5/19/07.

<https://doi.org/10.1145/3326229.3326272>

At the core of these improvements, one also finds techniques such as high-order lifting [41] and partial linearization [42], [16, Sec. 6].

For many of these operations, no implementation of the latest algorithms is available and no experimental evidence has been given regarding their practical behavior. Our goal is to partly remedy this issue, by providing and discussing implementations for a core of fundamental algorithms such as multiplication, approximant and interpolant bases, etc., upon which one may implement higher level algorithms. As an illustration, we describe the performance of slightly modified versions of Villard's recent breakthroughs on bivariate resultant and characteristic polynomial computation [49].

Our implementation is based on Shoup's Number Theory Library (NTL) [40], and is dedicated to polynomial matrix arithmetic over $\mathbb{K} = \mathbb{F}_p$ for a word-size prime p . Particular attention was paid to performance issues, so that our library compares favorably with previous work for those operations where comparisons were possible. Our code is available at <https://github.com/vneiger/pml>.

Overview. Polynomial matrix algorithms rely on efficient arithmetic in $\mathbb{K}[x]$ and for matrices over $\mathbb{K}$; in Section 2, we review some related algorithms and their NTL implementations. Then, we describe our implementation of a key building block: multiplication.

Section 3 presents the next major part of our work, concerning algorithms for *approximant bases* [2, 15, 25, 53] and *interpolant bases* [3, 23, 24, 47]. We focus on a version of interpolants which is less general than in these references but allows for a more efficient algorithm. In particular, we show that with this version, both interpolant and approximant bases can be used interchangeably in several contexts, with interpolants sometimes achieving better performance than approximants. In Section 4, we discuss algorithms for minimal kernel bases, linear system solving, determinant, and basis reduction. Finally, using these tools, we study the practical behavior of the bivariate resultant algorithm of [49] (Section 5).

Below, cost bounds are given in an algebraic complexity model, counting all operations in the base field at unit cost. While standard, this point of view fails to describe parts of the implementation (CRT-based algorithms, such as the 3-primes FFT, cannot be described in such a manner), but we believe that this is a minor issue.

Implementation choices. NTL is a C++ library for polynomial and matrix arithmetic over rings such as $\mathbb{Z}$, $\mathbb{Z}/n\mathbb{Z}$, etc., and is often seen as a reference point for fast implementations in such contexts. Other libraries for these operations include for example FLINT [18] as well as FFLAS-FFPACK and LinBox [45, 46]. Currently, our implementation relies solely on NTL; this choice was based on comparisons of performance for the functionalities we need.

In our implementation, the base field is a prime finite field $\mathbb{F}_p$; we rely on NTL's `lzz_p` class. At the time of writing, on standard x86_64 platforms, NTL v11.3.1 uses `unsigned long`'s as its primary data type for `lzz_p`, supporting moduli up to 60 bits long.

For such fields, one can directly compare running times and cost bounds, since in the literature most polynomial matrix algorithms are analyzed in the algebraic complexity model. Besides, computations over $\mathbb{F}_p$ are at the core of a general approach consisting in solving problems over $\mathbb{Z}$ or $\mathbb{Q}$ by means of reduction modulo sufficiently many primes, which are chosen so as to satisfy several, partly conflicting, objectives. We may want them to support Fourier transforms of high orders. Linear algebra modulo each prime should be fast, so we may wish them to be small enough to support vectorized matrix arithmetic with SIMD instructions. On the other hand, using larger primes allows one to use fewer of them, and reduces the likelihood of unlucky choices in randomized algorithms.

As a result, while all NTL `lzz_p` moduli are supported, our implementation puts an emphasis on three families: small FFT primes that support AVX-based matrix multiplication (such primes have at most 23 bits); arbitrary size FFT primes (at most 60 bits); arbitrary moduli (at most 60 bits). Very small fields such as $\mathbb{F}_2$ or $\mathbb{F}_3$ are supported, but we did not make any specific optimization for them.

Experiments. All runtimes below are in seconds and were measured on an Intel Core i7-4790 CPU with 32GB RAM, using the version 11.3.1 of NTL. Unless specified otherwise, timings are obtained modulo a random 60 bit prime. Runtimes were measured on a single thread; currently, most parts of our code do not explicitly exploit multi-threading. Tables below only show a few selected timings, with the best time(s) in bold; for more timings, see <https://github.com/vneiger/pml/tree/master/benchmarks>.

2 Basic polynomial and matrix arithmetic

We review basic algorithms for polynomials and matrices, and related complexity results that hold over an abstract field $\mathbb{K}$, and we describe how we implemented these operations. Hereafter, for $d \geq 0$, $\mathbb{K}[x]_d$ is the set of elements of $\mathbb{K}[x]$ of degree less than d .

2.1. Polynomial multiplication. Multiplication in $\mathbb{K}[x]$ and Fast Fourier Transform (FFT) are cornerstones of most algorithms in this paper. Let $M : \mathbb{N} \rightarrow \mathbb{N}$ be a function such that polynomials of degree at most d in $\mathbb{K}[x]$ can be multiplied in $M(d)$ operations in $\mathbb{K}$. If $\mathbb{K}$ supports FFT, we can take $M(d) \in O(d \log(d))$, and otherwise, $M(d) \in O(d \log(d) \log \log(d))$ [11, Chapter 8]; as in this reference, we assume that $d \mapsto M(d)/d$ is increasing. A useful variant of multiplication is the *middle product* [5, 17]: for integers c and d , and F in $\mathbb{K}[x]_c$ and G in $\mathbb{K}[x]_{c+d}$, `MIDDLEPRODUCT`(F, G, c, d) returns the slice of the product FG with coefficients of degrees $c, \dots, c + d - 1$; a common case is with $c = d$. The direct approach computes the whole product and extracts the slice. Yet, the *transposition principle* [27] yields a more efficient approach, saving a constant factor (roughly a factor 2 when $c = d$, if FFT multiplication is used).

Polynomial matrix algorithms frequently use fast evaluation and interpolation at multiple points. In general, subproduct tree techniques [11, Chapter 10] allow one to do evaluation and interpolation of polynomials in $\mathbb{K}[x]_d$ at d points in $O(M(d) \log(d))$ operations. For special sets of points, one can do better: if we know α in $\mathbb{K}$ of order at least d , then evaluation and interpolation at the geometric progression $(1, \alpha, \dots, \alpha^{d-1})$ can both be done in time $O(M(d))$ [6].

In NTL, multiplication in $\mathbb{F}_p[x]$ uses either naive, Karatsuba, or FFT techniques, depending on p and on the degree (NTL provides FFT primes with roots of unity of order 2^{25} , and supports arbitrary

user-chosen FFT primes). FFT multiplication uses the TFFT algorithm of [19] and Harvey’s improvements on arithmetic mod p [20]. For primes p that do not support Fourier transforms, multiplication is done by means of either 3-primes FFT techniques [11, Chapter 8] or Schönhage and Strassen’s algorithm. We implemented middle products for naive, Karatsuba and FFT multiplication, closely following [5, 17], as well as evaluation/interpolation algorithms for general sets of points and for geometric progressions.

2.2. Matrix multiplication. Let ω be such that $n \times n$ matrices over any ring can be multiplied by a bilinear algorithm doing $O(n^\omega)$ ring operations. The naive algorithm does exactly n^3 multiplications. First improvements due to Winograd and Waksman [50, 51] reduced the number of operations to $n^3/2 + O(n^2)$ if 2 is a unit. Strassen’s and Winograd’s recursive algorithms [43, 52] have $\omega = \log_2(7)$; the best known bound is $\omega \leq 2.373$ [8, 30]. Note that, using blocking, rectangular matrices of sizes $(m \times n)$ and $(n \times p)$ can be multiplied in $O(mnp \min(m, n, p)^{\omega-3})$ ring operations. NTL implements its own arithmetic for matrices over $\mathbb{F}_p$ and chooses one of several implementations depending on the bitsize of p , the matrix dimensions, the available processor instructions, etc.

2.3. Polynomial matrix multiplication. In what follows, we write $MM(n, d)$ for a function such that two $n \times n$ matrices of degree at most d can be multiplied in $MM(n, d)$ operations in $\mathbb{K}$; we make the assumption that $d \mapsto MM(n, d)/d$ is increasing for all n .

From the definitions above we obtain $MM(n, d) \in O(n^\omega M(d))$, which is in $O(n^\omega d)$. Using evaluation/interpolation at $1, \alpha, \dots, \alpha^{2d}$ or at roots of unity, one obtains the following bounds on $MM(n, d)$:

- $O(n^\omega d + n^2 M(d))$ if an element α in $\mathbb{K}$ of order more than $2d$ is known [6, Thm. 2.4].
- $O(n^\omega d + n^2 d \log(d))$ if $\mathbb{K}$ supports FFT in degree $2d$.

We also mention a polynomial analogue of an integer matrix multiplication algorithm from [10] which uses evaluation/interpolation, done plainly via multiplication by (inverse) Vandermonde matrices. Then, the corresponding part of the cost (e.g. $O(n^2 M(d))$ for geometric progressions) is replaced by the cost of multiplying matrices over $\mathbb{K}$ in sizes roughly $(d \times d)$ by $(d \times n^2)$; this is in $O(n^2 d^{\omega-1})$ if $d \leq n^2$. For moderate values of d , where $M(d)$ is not in the FFT regime, this allows us to leverage fast matrix multiplication over $\mathbb{K}$.

We implemented and compared various algorithms for matrix multiplication over $\mathbb{F}_p[x]$. For matrices of degree less than 5, we use dedicated routines based on Karatsuba’s and Montgomery’s formulas [32]; for matrices of small size (up to 10, depending on p), we use Waksman’s algorithm. For other inputs, most of our efforts were spent on variants of the evaluation/interpolation scheme.

For FFT primes, we use evaluation/interpolation at roots of unity. For general primes, we use either evaluation/interpolation at geometric progressions (if such points exist in $\mathbb{F}_p$), or our adaptation of the algorithm of [10], or 3-primes multiplication (as for polynomials, we lift the product from $\mathbb{F}_p[x]$ to $\mathbb{Z}[x]$, where it is done modulo up to 3 FFT primes). No single variant outperformed or underperformed all others for all sizes and degrees, so thresholds were experimentally determined to switch between these options, with different values for small (less than 23 bits) and for large primes.

Middle product versions of these algorithms were implemented, and are used in approximant basis algorithms (Section 3.1) and Newton iteration (Section 4.3). Multiplier classes are available: they

do precomputations on a matrix $\mathbf{A}$ to accelerate repeated multiplications by $\mathbf{A}$; they are used in Dixon’s algorithm (Section 4.3).

The table below shows timings for our multiplication and LinBox’ one, for random $m \times m$ matrices of degree d and two choices of prime p . The global comparison showed running times that are either similar or in favor of our implementation.

m	d	20 bit FFT prime			60 bit prime		
		ours	Linbox	ratio	ours	Linbox	ratio
8	131072	1.1198	1.5930	0.70	3.577	13.59	0.26
32	4096	0.4283	0.5092	0.84	2.000	5.330	0.38
128	1024	1.7292	2.1126	0.82	15.73	23.13	0.68
512	128	4.3533	4.3837	0.99	41.57	50.62	0.82

3 Approximant bases and interpolant bases

These bases are matrix generalizations of Padé approximation and play an important role in many higher-level algorithms. For $\mathbf{F}$ in $\mathbb{K}[x]^{m \times n}$ and M non-constant in $\mathbb{K}[x]$, they are bases of the $\mathbb{K}[x]$ -module $\mathcal{A}_M(\mathbf{F})$ of all $\mathbf{p}$ in $\mathbb{K}[x]^{1 \times m}$ such that $\mathbf{p}\mathbf{F} = 0 \bmod M$. Specifically, *approximant bases* are for $M = x^d$ and *interpolant bases* for $M = \prod_i (x - \alpha_i)$ for d distinct points $\alpha_1, \dots, \alpha_d$ in $\mathbb{K}$. (Here, we do not consider more general cases from the literature, for example with several moduli $M_1, \dots, M_n$, one for each column of $\mathbf{p}\mathbf{F}$.)

Since $\mathcal{A}_M(\mathbf{F})$ is free of rank m , such a basis is represented row-wise by a nonsingular $\mathbf{P}$ in $\mathbb{K}[x]^{m \times m}$. The algorithms below return $\mathbf{P}$ in *s-ordered weak Popov form* (also known as *s-quasi Popov form* [4]), for a given *shift* $\mathbf{s} = (s_1, \dots, s_m)$ in $\mathbb{Z}^m$. Shifts allow us to set degree constraints on the sought basis $\mathbf{P}$, and they inherently occur in a general approach for finding bases of solutions to equations (approximants, interpolants, kernels, etc.). Approximant basis algorithms often require $\mathbf{P}$ to be in *s-reduced form* [47]; although the *s-ordered weak Popov form* is stronger, obtaining it involves minor changes in these algorithms, without impact on performance according to our experiments. Recent literature shows that this stronger form reveals valuable information for further computations with $\mathbf{P}$ [23, 25], in particular for finding *s-Popov* bases [4].

From the shift $\mathbf{s}$, the *s-degree* of $\mathbf{p} = [p_i]_i \in \mathbb{K}[x]^{1 \times m}$ is defined as $\text{rdeg}_s(\mathbf{p}) = \max_{1 \leq i \leq m} (\deg(p_i) + s_i)$, which extends to matrices: $\text{rdeg}_s(\mathbf{P})$ is the list of *s-degrees* of the rows of $\mathbf{P}$. Then, the *s-pivot* of $\mathbf{p}$ is its rightmost entry p_i such that $\text{rdeg}_s(\mathbf{p}) = \deg(p_i) + s_i$, and a nonsingular matrix $\mathbf{P}$ is in *s-ordered weak Popov form* if the *s-pivots* of its rows are located on the diagonal.

To simplify cost bounds below, we make use of the function $\text{MM}'(m, d) = \sum_{i=0}^{\log_2(d)} 2^i \text{MM}(m, d/2^i) \in O(\text{MM}(m, d) \log(d))$.

3.1. Approximant bases. For $\mathbf{F}$ in $\mathbb{K}[x]^{m \times n}$ and d in $\mathbb{Z}_{>0}$, an *approximant basis* for $(\mathbf{F}, d)$ is a nonsingular $m \times m$ matrix whose rows form a basis of $\mathcal{A}_{x^d}(\mathbf{F})$. We implemented minor variants of the algorithms M-BASIS (iterative, via matrix multiplication) and PM-BASIS (divide and conquer, via polynomial matrix multiplication) from [15]. The lowest-level function (M-BASIS-1 with the signature in Algorithm 1), handles order $d = 1$ in time $O(\text{rank}(\mathbf{F})^{\omega-2} mn)$; here, working modulo X , the matrix $\mathbf{F}$ is over $\mathbb{K}$. Our implementation follows [25, Algo. 1], which returns an *s-Popov* basis, using only an additional row permutation compared to the algorithm in [15].

This form of the output of M-BASIS-1 suffices to ensure that M-BASIS and PM-BASIS return bases in *s-ordered weak Popov form*. Our implementation of M-BASIS follows the original design [15]

Algorithm 1: M-BASIS-1($\mathbf{F}, \mathbf{s}$)

Input: matrix $\mathbf{F}$ in $\mathbb{K}^{m \times n}$, shift $\mathbf{s}$ in $\mathbb{Z}^m$

Output: the *s-Popov* approximant basis for $(\mathbf{F}, 1)$

with d iterations, each computing the *residual* $\mathbf{R}$ and updating $\mathbf{P}$ via multiplication by a basis $\mathbf{Q}$ obtained by M-BASIS-1 on $\mathbf{R}$. We also follow [15] for PM-BASIS, using a threshold T such that M-BASIS is called if $d \leq T$. Building PM-BASIS directly upon M-BASIS-1, i.e. choosing $T = 1$, has the same cost bound but is slower in practice.

Input: matrix $\mathbf{F}$ in $\mathbb{K}[x]^{m \times n}$, order d in $\mathbb{Z}_{>0}$, shift $\mathbf{s}$ in $\mathbb{Z}^m$

Output: an *s-ordered weak Popov* approximant basis for $(\mathbf{F}, d)$

Algorithm 2: M-BASIS($\mathbf{F}, d, \mathbf{s}$)

1. $\mathbf{P} \leftarrow$ identity matrix in $\mathbb{K}[x]^{m \times m}$, and $\mathbf{t} \leftarrow$ copy of $\mathbf{s}$
2. For $k = 0, \dots, d - 1$:
 - a. $\mathbf{R} \in \mathbb{K}^{m \times n} \leftarrow$ coefficient of $\mathbf{P}\mathbf{F}$ of degree k
 - b. $\mathbf{Q} \in \mathbb{K}[x]^{m \times m} \leftarrow$ M-BASIS-1($\mathbf{R}, \mathbf{t}$)
 - c. $\mathbf{P} \leftarrow \mathbf{Q}\mathbf{P}$, and then $\mathbf{t} \leftarrow \text{rdeg}_s(\mathbf{P})$
3. Return $\mathbf{P}$

Algorithm 3: PM-BASIS($\mathbf{F}, d, \mathbf{s}$)

1. if $d \leq T$ return M-BASIS($\mathbf{F}, d, \mathbf{s}$)
2. $\mathbf{P}_1 \leftarrow$ PM-BASIS($\mathbf{F} \bmod x^{\lceil d/2 \rceil}, \lceil d/2 \rceil, \mathbf{s}$)
3. $\mathbf{R} \leftarrow \text{MIDDLEPRODUCT}(\mathbf{P}_1, \mathbf{F}, \lceil d/2 \rceil, \lfloor d/2 \rfloor)$
4. $\mathbf{t} \leftarrow \text{rdeg}_s(\mathbf{P}_1)$
5. $\mathbf{P}_2 \leftarrow$ PM-BASIS($\mathbf{R}, \lfloor d/2 \rfloor, \mathbf{t}$)
6. return $\mathbf{P}_2\mathbf{P}_1$

These algorithms use $O((m^\omega + m^{\omega-1}n)d^2)$ and $O((1 + \frac{n}{m})\text{MM}'(m, d))$, respectively [15]. Some implementation details are discussed in Section 3.2. The next table compares timings for LinBox’ and our implementations of PM-BASIS for a 20 bit FFT prime (LinBox’ implementation is not optimized for large primes and general primes).

m	n	d	ours	Linbox	ratio
8	4	131072	6.6754	15.5743	0.43
32	16	8192	4.4185	7.1150	0.62
128	64	2048	18.0030	28.7113	0.63
512	256	256	39.6255	42.4051	0.93

We also implemented [25, Algo. 3] which returns *s-Popov* bases and is about twice slower than PM-BASIS; making this overhead negligible for some usual cases is future work. For completeness, we handle general approximants (with one modulus per column of $\mathbf{F}$) by an iterative approach from [2, 47]; faster algorithms are more complex [23–25] and use partial linearization techniques.

These techniques from [42, 53] yield cost bounds in $O(m^{\omega-1}nd)$, which is a $\Theta(\frac{m}{n})$ speedup compared to PM-BASIS. Implementing them is work in progress. experimental code, which focuses for simplicity on $n = 1$ and “generic” inputs for which the degrees in $\mathbf{P}$ can be predicted, revealed significant speedups:

m	n	d	PM-BASIS	PM-BASIS with linearization
4	1	65536	1.6693	1.26891
16	1	16384	1.8535	0.89652
64	1	2048	2.2865	0.14362
256	1	1024	36.620	0.20660

Approximant bases are often applied to solve block-Hankel systems [28]. In two specific settings, we have compared this approach

to the one which uses structured matrix algorithms; we are not aware of previous comparisons of this kind. We obtain the following running times, using the NTL-based solver from [21].

m	d	Setting 1		Setting 2	
		PM-BASIS	solver	PM-BASIS	solver
5	8000	0.996	8.23	2.19	3.820
12	1000	0.687	6.18	2.33	2.28
30	500	2.84	42.5	19.5	11.5

Setting 1: we call PM-BASIS on $[F^T - I_m]^T$ at order $2d$ with shift $(0, \dots, 0)$, where F is an $m \times m$ matrix of degree $2d-1$, and we solve a system with $m \times m$ Hankel blocks of size $d \times d$ (the structured solver returns a random solution to the system). Our experiments show a clear advantage for approximant algorithms. The asymptotic costs being similar, the effects at play here are constant factor differences: approximant basis algorithms seem to be somewhat simpler and to better leverage the main building blocks (matrix arithmetic over $\mathbb{K}$ and univariate polynomial arithmetic).

Setting 2 is the vector rational reconstruction problem. We call PM-BASIS on $[F^T - I_m]^T$ at order $(m+1)d$ with shift $(0, \dots, 0)$, where F is a $1 \times m$ vector of degree $(m+1)d-1$, and we solve a block system with $1 \times m$ Hankel blocks of size $md \times d$. The cost bounds are $O(m^{\omega+1}d)$ and $O(m^\omega d)$, respectively. Approximants are faster up to dimension about 15, which is explained by the arguments in the previous paragraph. For larger dimensions, as predicted by the cost estimates, the block-Hankel solver is more efficient.

3.2. Interpolant bases. For matrices $E = (E_1, \dots, E_d)$ in $\mathbb{K}^{m \times n}$ and pairwise distinct points $\alpha = (\alpha_1, \dots, \alpha_d)$ in $\mathbb{K}$, consider

$$\mathcal{I}_\alpha(E) = \{p \in \mathbb{K}[x]^{1 \times m} \mid p(\alpha_i)E_i = 0 \text{ for } 1 \leq i \leq d\}.$$

An *interpolant basis* for (E, α) is a matrix whose rows form a basis of the $\mathbb{K}[x]$ -module $\mathcal{I}_\alpha(E)$. Note that $\mathcal{I}_\alpha(F(\alpha_1), \dots, F(\alpha_d))$ coincides with $\mathcal{A}_M(F)$, for F in $\mathbb{K}[x]^{m \times n}$ and $M = \prod_{i=1}^d (x - \alpha_i)$.

This definition is a specialization of those in [3, 24], which consider n sets of points, one for each of the n columns of $E_1, \dots, E_d$: here, these sets are all equal. This more restrictive problem allows us to give faster algorithms than those in these references, by direct adaptations of the approximant basis algorithms presented above. Besides, Sections 4.1 and 4.2 will show that interpolant bases can often play the same role as approximant bases in applications.

These adaptations are described in Algorithms 4 and 5, where $\alpha_{i..j}$ stands for the sublist $(\alpha_i, \alpha_{i+1}, \dots, \alpha_j)$. In the next proposition, we assume that $MM(n, d)$ is in $\Omega(n^2 M(d))$ (instead, one may add an extra term $O(n^2 M(d) \log(d))$ in the cost).

PROPOSITION 3.1. *Algorithm 5 is correct. For input evaluation points in geometric progression, it costs $O(MM'(m, d))$ if $n \leq m$ and $O(MM'(m, d) + m^{\omega-1}nd \log(d))$ otherwise. For general evaluation points, an extra cost $O(m^2 M(d) \log^2(d))$ is incurred.*

(Correctness follows from Items (i) and (iii) of [25, Lem. 2.4]; the cost analysis is standard for such divide and conquer algorithms.)

Our current code uses the threshold $T = 32$ in the divide and conquer PM-BASIS and PM-INTBASIS: beyond this point, they are faster than the iterative M-BASIS and M-INTBASIS. Unlike in most other functions, where elements of $\mathbb{K}[x]^{m \times n}$ are represented as matrices of polynomials (Mat<Vec<zz_p>> in NTL), in M-BASIS and M-INTBASIS we see them as polynomials with matrix coefficients (Vec<Mat<zz_p>>). Indeed, since these algorithms involve

Input: matrices $E = (E_1, \dots, E_d)$ in $\mathbb{K}^{m \times n}$, evaluation points $\alpha = (\alpha_1, \dots, \alpha_d)$ in $\mathbb{K}$, shift s in $\mathbb{Z}^m$

Output: an s -ordered weak Popov interpolant basis for (E, α)

Algorithm 4: M-INTBASIS(E, α, s)

1. $P \leftarrow$ identity matrix in $\mathbb{K}[x]^{m \times m}$, and $t \leftarrow$ copy of s
2. For $k = 0, \dots, d-1$:
 - a. $R \in \mathbb{K}^{m \times n} \leftarrow P(\alpha_k)E_k$
 - b. $Q \in \mathbb{K}[x]^{m \times m} \leftarrow$ M-BASIS-1(R, t)
 - c. $P \leftarrow QP$, and then $t \leftarrow \text{rdeg}_s(P)$
3. Return P

Algorithm 5: PM-INTBASIS(E, α, s)

1. if $d \leq T$ return M-INTBASIS(E, α, s)
2. $P_1 \leftarrow$ PM-INTBASIS($E_{1.. \lceil d/2 \rceil}, \alpha_{1.. \lceil d/2 \rceil}, s$)
3. $R \leftarrow (P_1(\alpha_{\lceil d/2 \rceil+1})E_{\lceil d/2 \rceil+1}, \dots, P_1(\alpha_d)E_d)$
4. $t \leftarrow \text{rdeg}_s(P_1)$
5. $P_2 \leftarrow$ PM-INTBASIS($R, \alpha_{\lceil d/2 \rceil+1..d}, t$)
6. return $P_2 P_1$

only matrix arithmetic over $\mathbb{K}$ (recall that $\deg(Q) \leq 1$), this turns out to be more cache-friendly and faster.

We implemented two variants for approximant bases: either the residual R is computed from P and F at each iteration, or we initialize a list of residuals with a copy of F and we update the whole list at each iteration using Q . The second variant improves over the first when $n > m/2$, with significant savings when n is close to m . For interpolant bases, this does not lead to any gain.

Timings are showed in the next table, for Algorithms M-BASIS (M), M-INTBASIS (M-I), PM-BASIS (PM), PM-INTBASIS for general points (PM-I) and for geometric points (PM-Ig). For approximants, we take a random input in $\mathbb{K}[x]^{m \times n}$ of degree $d-1$; for interpolants, we take d random matrices in $\mathbb{K}^{m \times n}$. We focus on the common case $m \approx 2n$, which arises for example in kernel algorithms (Section 4.2) and in fraction reconstruction, itself used in basis reduction (Section 4.5) and in the resultant algorithm of [49] (Section 5).

m	n	d	M	M-I	d	PM	PM-I	PM-Ig
4	2	32	1.60e-4	1.42e-4	32768	1.06	6.81	1.47
16	8	32	1.98e-3	1.55e-3	4096	1.82	5.51	1.92
32	16	32	0.0104	7.59e-3	2048	3.90	8.18	3.56
64	32	32	0.0502	0.0354	1024	8.1	12.2	6.38
128	64	32	0.374	0.253	1024	45	56.7	33.3
256	128	32	2.92	1.83	1024	288	292	198

Concerning iterative algorithms, we observe that interpolants are slightly faster than approximants, which is explained by the cost of computing the residual R : it uses one Horner evaluation of P and one matrix product for interpolants, whereas for approximants it uses about $\min(k, \deg(P))$ matrix products at iteration k .

As for the divide and conquer algorithms, interpolant bases with general points are slower, in some cases significantly, than the other two algorithms: although the complexity analysis predicted a disadvantage, we believe that our implementation of multipoint evaluation at general points could be improved to reduce this gap. For the other two algorithms, the comparison is less clear. There could be many factors at play here, but the main differences lie in the base case (Step 1) which calls the iterative algorithm, and in the

computation of residuals (Step 3) which uses either middle products or geometric evaluation. It seems that FFT-based polynomial multiplication performs slightly better than geometric evaluation for small matrices and slightly worse for large matrices.

4 Higher-level algorithms

In this section we consider kernel bases, system solving, determinants, and basis reduction; we discuss algorithms which rely on multiplication, through approximant/interpolant bases and lifting techniques. For many of these algorithms, we are not aware of previous implementations or experimental comparisons.

4.1. A note on matrix fraction reconstruction. Given $\mathbf{H}$ in $\mathbb{K}(x)^{n \times n}$, a *left fraction description* of $\mathbf{H}$ is a pair of polynomial matrices $(\mathbf{Q}, \mathbf{R})$ in $\mathbb{K}[x]^{n \times n}$ such that $\mathbf{H} = \mathbf{Q}^{-1}\mathbf{R}$. It is *minimal* if $\mathbf{Q}$ and $\mathbf{R}$ have unimodular left matrix GCD and $\mathbf{Q}$ is in reduced form (*right fraction descriptions* are defined similarly). Besides, $\mathbf{H}$ is said to be *strictly proper* if the numerator of each of its entries has degree less than the corresponding denominator.

Such a description of $\mathbf{H}$ is often computed from the power series expansion of $\mathbf{H}$ at sufficient precision, using an approximant basis. Yet, for resultant computations in Section 5.2, we would like to use an interpolant basis to obtain this description from sufficiently many values of $\mathbf{H}$. We now state the validity of this approach; this is a matrix version of rational function reconstruction [11, Chap. 5.7].

PROPOSITION 4.1. *Let $\mathbf{H}$ be in $\mathbb{K}(x)^{n \times n}$ be strictly proper and suppose $\mathbf{H}$ admits left and right fraction descriptions of degrees at most D , for some $D \in \mathbb{Z}_{>0}$. For M in $\mathbb{K}[x]$ of degree at least $2D$ and such that all denominators in $\mathbf{H}$ are invertible modulo M , define the matrix $\mathbf{F} = [\mathbf{H} \bmod M \quad -\mathbf{I}_n]^T \in \mathbb{K}[x]^{2n \times n}$. Then, if $\mathbf{P} \in \mathbb{K}[x]^{2n \times 2n}$ is a 0-ordered weak Popov basis of $\mathcal{A}_M(\mathbf{F})$, the first n rows of $\mathbf{P}$ form a matrix $[\mathbf{Q} \quad \mathbf{R}]$ such that $(\mathbf{Q}, \mathbf{R})$ is a minimal left fraction description of $\mathbf{H}$, with $\mathbf{Q}$ in 0-ordered weak Popov form.*

The proof given in [15, Lem 3.7] for the specific $M = x^{2D+1}$ extends to any modulus M ; using an ordered weak Popov form (rather than a reduced form) allows us both to know *a priori* that the first n rows are those of degree at most D , and to use degree $2D$ instead of $2D + 1$ (since $\deg(\mathbf{R}) < \deg(\mathbf{Q})$ is ensured by this form).

In particular, if $M = \prod_{i=1}^{2D} (x - \alpha_i)$ for pairwise distinct points $(\alpha_1, \dots, \alpha_{2D})$, the interpolant basis algorithms in Section 3.2 give a minimal left fraction description of $\mathbf{H}$ from $\mathbf{H}(\alpha_1), \dots, \mathbf{H}(\alpha_{2D})$.

4.2. Kernel basis. We implemented two kernel basis algorithms: the first one, based on Lemma 4.2, finds the kernel basis from a single approximant basis at sufficiently large order; the second one, from [54], uses several approximant bases at small order and combines recursively obtained kernel bases via multiplication. With a minor modification and no performance impact, the latter returns an s -ordered weak Popov basis. In both cases, we designed variants which rely on interpolant bases instead of approximant bases.

LEMMA 4.2. *Let $\mathbf{F}$ be in $\mathbb{K}[x]^{m \times n}$ of degree $d \geq 0$, let $\mathbf{s}$ be in $\mathbb{N}^m$, and let δ in $\mathbb{Z}_{>0}$ be an upper bound on the s -degree of any s -reduced left kernel basis of $\mathbf{F}$; for example, $\delta = nd + \max(\mathbf{s}) - \min(\mathbf{s}) + 1$. Let M be in $\mathbb{K}[x]$ of degree at least $\delta + d$, and $\mathbf{P}$ in $\mathbb{K}[x]^{m \times m}$ be an s -reduced basis of $\mathcal{A}_M(\mathbf{F})$. Then, the submatrix of $\mathbf{P}$ formed by its rows of s -degree less than δ is an s -reduced left kernel basis for $\mathbf{F}$.*

For a proof, see <https://hal.archives-ouvertes.fr/hal-01995873v1/document>.

In particular, one may find $\mathbf{P}$ via PM-INTBASIS at $d + \delta$ points or via PM-BASIS at order $d + \delta$; for $n \leq m$, this costs $O(\text{MM}'(m, d + \delta))$. The approximant-based direct approach is folklore [54, Sec. 2.3], yet explicit statements in the literature focus on shifts linked to the degrees in $\mathbf{F}$, with better bounds δ (see [54, Lem. 3.3], [33, Lem. 4.3]).

The algorithm of [54] is more efficient, at least when the entries of $\mathbf{s}$ are close to the corresponding row degrees of $\mathbf{F}$; for a uniform shift, it costs $O(m^\omega \lceil nd/m \rceil)$ operations. We obtained significant practical improvements over the plain implementation of [54, Algo. 1] thanks to the following observation: if $n \leq m/2$, for a vast majority of input $\mathbf{F}$, the approximant basis at Step 2 of [54, Algo. 1], computed at order more than $2s$, contains the sought kernel basis. Furthermore, this can be easily tested by checking well-chosen degrees, and then the algorithm can exit early, avoiding the further recursive calls. We took advantage of this via the following modifications: we use order $2s + 1$ rather than $3s$ (see [54, Rmk. 3.5] for a discussion on this point), and when $n > m/2$ we directly reduce the number of columns via the divide and conquer scheme in [54, Thm. 3.15].

The use of approximants here follows the idea in Lemma 4.2: row vectors of small degree which are in $\mathcal{A}_M(\mathbf{F})$ for a large degree M must be in the kernel of $\mathbf{F}$. Thus, one can directly replace approximant bases with interpolant bases in [54, Algo. 1], up to modifying Step 8 accordingly (dividing by the appropriate polynomial M).

Timings for both approaches are showed in the next table, for a random $\mathbf{F}$ of degree d . Except for the last row, the shift is uniform and, as expected, [54, Algo. 1] is faster than the direct approach; the differences between interpolant and approximant variants follow those observed in Section 3. The last row corresponds to a shift yielding the kernel basis in Hermite form and shows, as expected, that the direct approach is faster for shifts that are far from uniform.

We note that [54, Algo. 1] may use partial linearization if it computes matrix products with unbalanced degrees or approximant bases with $n \ll m$. We have not yet implemented this part of the algorithm, which may lead to slowdowns for some rare inputs.

m	n	d	direct		[54, Algo. 1]	
			approx.	int.	approx.	int.
8	4	8192	7.22	6.60	2.16	2.49
8	7	8192	14.1	14.4	4.64	5.63
32	16	1024	86.3	63.1	3.75	3.51
32	31	1024	142	118	8.27	8.09
128	64	256	2720	1827	16.8	11.8
128	127	256	>1h	>1h	43.8	35.6
16	1	512	5.68	5.31	11.5	11.2

4.3. Linear system solving. For systems $\mathbf{A}\mathbf{v} = \mathbf{b}$, with $\mathbf{A}$ in $\mathbb{K}[x]^{m \times n}$, $\mathbf{b}$ in $\mathbb{K}[x]^{m \times 1}$ and $\mathbf{v}$ in $\mathbb{K}(x)^{n \times 1}$, we implemented two families of algorithms. The first one uses lifting techniques, assuming $\mathbf{A}$ is square, nonsingular, with $\mathbf{A}(0)$ invertible; the algorithm returns a pair $(\mathbf{u}, f)$ in $\mathbb{K}[x]^{n \times 1} \times \mathbb{K}[x]$ such that $\mathbf{A}\mathbf{u} = f\mathbf{b}$ and f has minimal degree. The second one uses a kernel basis and works for any input $\mathbf{A}$; under the assumptions above, it has a similar output.

Lifting techniques. Under the above assumptions, our lifting algorithm is standard: if $\mathbf{A}$ and $\mathbf{b}$ have degree at most d , we first compute the truncated inverse $\mathbf{S} = \mathbf{A}^{-1} \bmod x^{d+1}$ by matrix Newton iteration [38]. Then, we use Dixon's algorithm [9] to compute $\mathbf{v} \bmod x^{2nd} = \mathbf{A}^{-1}\mathbf{b} \bmod x^{2nd}$; it consists of roughly $2n$ steps, each involving a matrix-vector product using either $\mathbf{A}$ or $\mathbf{S}$. Then, vector

rational reconstruction is applied to recover $(\mathbf{u}, f)$ from $\mathbf{v}$. The cost of this algorithm is $O(\text{MM}(n, d))$ for the truncated inverse of $\mathbf{A}$ and $O(n^3 M(d))$ for Dixon's algorithm; overall this is in $O(n^3 d)$.

To reduce the exponent in n , Storjohann introduced the *high-order lifting* algorithm [41]. The core of this algorithm is the computation of $\Theta(\log(n))$ slices $S_0, S_1, \dots$ of the power series expansion of $\mathbf{A}^{-1}$, where the coefficients of S_i are the coefficients of degree $(2^i - 1)d - 2^i + 1, \dots, (2^i + 1)d - 2^i - 1$ in $\mathbf{A}^{-1}$. These matrices are computed recursively, each step involving 4 matrix products; the other steps of the algorithm, that use these S_i to compute $\mathbf{v} \bmod x^{2nd}$, are cheaper, so the runtime is $O(\text{MM}(n, d) \log(n)) \subset O(n^\omega d)$.

Using kernel bases. For this second approach, let $\mathbf{A}$ be any matrix in $\mathbb{K}[x]^{m \times n}$ and $\mathbf{b}$ be in $\mathbb{K}[x]^{m \times 1}$. The algorithm simply computes $\mathbf{K} \in \mathbb{K}[x]^{(n+1) \times k}$, a right kernel basis of the augmented matrix $[\mathbf{A} \mid \mathbf{b}] \in \mathbb{K}[x]^{m \times (n+1)}$. The matrix $\mathbf{K}$ generates, via $\mathbb{K}(x)$ -linear combinations of its columns, all solutions $\mathbf{v} \in \mathbb{K}(x)^{n \times 1}$ to $\mathbf{A}\mathbf{v} = \mathbf{b}$.

In particular, if $\mathbf{K}$ is empty (i.e. $k = 0$, which requires $m \geq n$), or if the last row of $\mathbf{K}$ is zero, then the system has no solution. Furthermore, if $\mathbf{A}$ is square and nonsingular, $\mathbf{K}$ has a single column $[\mathbf{u}^\top \mid f]^\top$, where $\mathbf{u} \in \mathbb{K}[x]^{n \times 1}$ and $f \in \mathbb{K}[x]$, with f of minimal degree (otherwise, $\mathbf{K}$ would not be a basis).

In this context, the fastest known kernel algorithm is [54, Algo. 1]. To exploit it best, we choose the input shift $\mathbf{s} = (d, d)$, where $d = \deg(\mathbf{b})$ and $\mathbf{d} \in \mathbb{N}^n$ is the tuple of column degrees of $\mathbf{A}$ (zero columns of $\mathbf{A}$ are discarded while computing $\mathbf{d}$).

Implementation. We implemented the approaches described above: lifting with Dixon's algorithm, high-order lifting, and via kernel. The table below shows timings for randomly chosen $m \times m$ matrix $\mathbf{A}$ and $m \times 1$ vector $\mathbf{b}$, both of degree d . In this case the lifting algorithms apply (with high probability). On such inputs, Dixon's algorithm usually does best. High-order lifting, although theoretically faster, is outperformed, mainly because it performs $\Theta(\log(n))$ matrix products (we will however see that this algorithm still plays an important role for basis reduction). The kernel based approach is moderately slower than Dixon's algorithm, but has the advantage of working without any assumption on $\mathbf{A}$.

m	d	Dixon	high-order lifting	kernel
16	1024	1.53	2.39	2.07
32	1024	4.94	13.8	9.45
64	1024	19.5	94.7	46.5
128	512	55.2	266	108

4.4. Determinant. We implemented four algorithms, taking as input a square $m \times m$ matrix $\mathbf{A}$.

The most basic one uses expansion by minors, which turns out to be the fastest option up to dimension about 6.

The second one assumes that we have an element α in $\mathbb{K}$ of order at least $2md + 1$, and uses evaluation/interpolation at the geometric progression $1, \alpha^2, \dots, \alpha^{2md}$; this costs $O(m^3 M(d) + m^{\omega+1} d)$ operations in $\mathbb{K}$. For dimensions between 7 and about 20, it is often the fastest variant, sometimes competing with the third.

The third one consists in solving a linear system with random right-hand side of degree d ; this yields a solution $(\mathbf{u}, f)$ with f the sought determinant up to a constant [36]. For dimensions exceeding 20, this is the fastest method (assuming that $\mathbf{A}(0)$ is invertible).

The last one, from [29], is based on triangularization and runs in $O(m^\omega d)$. Our implementation currently only supports the generic case where the Hermite form of $\mathbf{A}$ has diagonal $(1, \dots, 1, \det(\mathbf{A}))$, or in other words, all so-called row bases computed in that algorithm are the identity. This allows us to circumvent the temporary lack of a row basis implementation, while still being able to observe meaningful timings. Indeed, we believe that timings for a complete implementation called on such "generic" input will be similar to the timings presented here in many cases of interest, where one can easily detect whether the algorithm has made a wrong prediction about the row basis being identity (for example if the input matrix is reduced, which is the case if it is the denominator of a minimal fraction description). This recursive determinant algorithm calls expansion by minors as a base case for small dimensions; for larger dimensions, it is generally slightly slower than the third method.

The timings below are for a random $m \times m$ matrix $\mathbf{A}$ of degree d .

m	d	minors	evaluation	linsolve	triangular
4	65536	0.7751	2.014	8.460	0.7826
16	4096	∞	4.14	5.023	7.38
32	4096	∞	34.5	25.4	41.6
64	2048	∞	127	68.6	100
128	512	∞	244	96.6	99.0

4.5. Basis reduction. Our implementation of the algorithm of [15] takes as input a nonsingular matrix $\mathbf{A} \in \mathbb{K}[x]^{m \times m}$ of degree d such that $\mathbf{A}(0)$ is invertible, and returns a reduced form of $\mathbf{A}$. The algorithm first computes a slice S of $2d$ consecutive coefficients of degree about md in the power series expansion of $\mathbf{A}^{-1}$, then uses PM-Basis to reconstruct a fraction description $\mathbf{S}^{-1} = \mathbf{R}^{-1}\mathbf{Q}$, and then returns $\mathbf{R}$. A Las Vegas randomized version is mentioned in [15], to remove the assumption on $\mathbf{A}(0)$: we will implement it for large enough $\mathbb{K}$, but for smaller fields this requires to work in an extension of $\mathbb{K}$, which is currently beyond the scope of our work.

In our experiments, to create the input $\mathbf{A}$, we started from a random $m \times m$ matrix of degree $d/3$ (which is reduced with high probability), and we left-multiplied it by a lower unit triangular matrix and then by an upper one, both chosen at random of degree $d/3$. The following table shows timings for both steps, with the first step either based on Newton iteration or on high-order lifting; the displayed total time is when using the faster of the two. We conclude that for reduction, as opposed to the above observations for system solving, it is crucial to rely on high-order lifting. Indeed, it improves over Newton iteration already for dimension 8, and the gap becomes quite significant when the dimension grows.

m	d	Newton	high-order	reconstruct	total
4	24574	1.251	1.688	8.772	10.02
8	6142	2.617	2.244	8.851	11.09
16	1534	4.457	3.044	8.506	11.55
32	382	11.147	4.858	7.977	12.83
64	94	30.62	5.509	5.833	11.34

5 Applications to bivariate resultants

We conclude this paper with algorithms originating from Villard's recent breakthrough on computing the determinant of structured polynomial matrices [49]. Fix a field $\mathbb{K}$ and consider the two following questions: computing the resultant of two polynomials F, G in $\mathbb{K}[x, z]$ with respect to z , and computing the characteristic polynomial of an element A in $\mathbb{K}[z]/(P)$, for some P in $\mathbb{K}[z]$.

The second problem is a particular case of the former, since the characteristic polynomial of A modulo P is the resultant of $x - A(z)$ and $P(z)$ with respect to z , up to a nonzero constant. Let n be an upper bound on the degree in z of the polynomials we consider, and d be a bound on their degree in x (so in the second problem, $d = 1$). Villard proved that for generic inputs, both problems can be solved in $O(n^{2-1/\omega}d) \subset O(n^{1.58}d)$ operations in $\mathbb{K}$. For the first problem, the best previous bound is $O(n^2d)$, obtained either by evaluation/interpolation techniques or Reischert's algorithm [37]. For the second problem, the previous record was $O(n^{\omega_2/2})$, where ω_2 is the exponent of matrix multiplication in size $(s, s) \times (s, s^2)$, with $\omega_2/2 \leq 1.63$ [31]. Note that these bounds apply to all inputs.

We show how the work we presented above allows us to put Villard's ideas to practice, and outperform the previous state of the art for large input sizes. This is however not straightforward: in both cases, this required modifications of Villard's original designs (for the second case, using an algorithm from [34]).

5.1. Overview of the approach. In [49], Villard designed the following algorithm to find the determinant of a matrix $\mathbf{P}$ over $\mathbb{K}[x]$.

Algorithm 6: DETERMINANT($\mathbf{P}, m$)

Input: nonsingular $\mathbf{P}$ in $\mathbb{K}[x]^{v \times v}$; parameter $m \in \{1, \dots, v\}$

Output: $\det(\mathbf{P})$

1. compute $\tilde{\mathbf{H}} = \mathbf{H} \bmod x^{2\lceil v/m \rceil d + 1}$, where d is the degree of $\mathbf{P}$ and $\mathbf{H}$ is the $m \times m$ top-right quadrant of $\mathbf{P}^{-1} \in \mathbb{K}(x)^{v \times v}$
2. from $\tilde{\mathbf{H}}$, find a minimal left fraction description $(\mathbf{Q}, \mathbf{R})$ of $\mathbf{H}$
3. return $\det(\mathbf{Q})$

The parameter m is chosen so as to minimize the theoretical cost. The correctness of the algorithm follows from the next properties, which do not hold for an arbitrary nonsingular $\mathbf{P}$: the matrix $\mathbf{H}$ is strictly proper and admits a left fraction description $\mathbf{H} = \mathbf{Q}^{-1}\mathbf{R}$ such that $\det(\mathbf{P}) = \det(\mathbf{Q})$, for $\mathbf{Q}$ and $\mathbf{R}$ in $\mathbb{K}[x]^{m \times m}$ of degree at most $\lceil v/m \rceil d$ (see Section 4.1 for definitions). In [49], they are proved to hold for generic instances of the problems discussed here.

Once sufficiently many terms of the expansion of $\mathbf{H}$ have been obtained in Step 1, the denominator $\mathbf{Q}$ is recovered by an approximant basis algorithm and its determinant is computed by a general algorithm in Steps 2 and 3, which cost $O(m^\omega(vd/m))$.

While the algorithm applies to any nonsingular matrix $\mathbf{P}$ satisfying the properties above, in general it does not improve over previously known methods (see Section 4.4). Indeed, the fastest known algorithm for obtaining $\tilde{\mathbf{H}}$ costs $O(v^\omega d)$ operations via high-order lifting (see for example [16, Thm. 1]).

However, sometimes $\mathbf{P}$ has some structure which helps to speed up the first step. Villard pointed out that when $\mathbf{P}$ is the Sylvester matrix of two bivariate polynomials, then $\mathbf{P}^{-1}$ is a Toeplitz-like matrix which can be described succinctly as $\mathbf{P}^{-1} = \mathbf{L}_1\mathbf{U}_1 + \mathbf{L}_2\mathbf{U}_2$; here, $\mathbf{L}_1, \mathbf{L}_2$ (resp. $\mathbf{U}_1, \mathbf{U}_2$) are lower (resp. upper) triangular Toeplitz matrices with entries in $\mathbb{K}(x)$. Hence, we start by computing the first columns $\mathbf{c}_1$ of $\mathbf{L}_1$ and $\mathbf{c}_2$ of $\mathbf{L}_2$ as well as the first rows $\mathbf{r}_1$ of $\mathbf{U}_1$ and $\mathbf{r}_2$ of $\mathbf{U}_2$, all of them modulo $x^{2\lceil v/m \rceil d + 1}$; then, $\tilde{\mathbf{H}}$ is directly obtained via the above formula for $\mathbf{P}^{-1}$, using $O(mvd)$ operations. Computing these rows and columns is done by solving systems with matrices $\mathbf{P}$ and $\mathbf{P}^T$ and very simple right-hand sides [49, Prop. 5.1], with power series coefficients, in time $O(v^2d/m)$.

Altogether, taking $m = v^{1/\omega}$ minimizes the cost, yielding the runtime $O(v^{2-1/\omega}d)$. In the case of bivariate resultants described above, the Sylvester matrix of F and G has size $v = 2n$, hence the cost bound $O(n^{2-1/\omega}d)$.

5.2. Resultant of generic bivariate polynomials. We implemented the algorithm described in the previous section to compute the resultant of generic F, G in $\mathbb{K}[x, z]$; first experiments showed that obtaining $\mathbf{c}_1, \mathbf{c}_2, \mathbf{r}_1, \mathbf{r}_2$ was a bottleneck. These vectors have power series entries and are solutions of linear systems whose matrix is the Sylvester matrix of F and G or its transpose: they were obtained via Hensel lifting techniques, following [11, Ch. 15.4].

To get better performance, we designed a minor variant of Villard's algorithm: instead of computing the power series expansion of $\mathbf{H}$ modulo x^δ , where $\delta = 2\lceil v/m \rceil d + 1$, we compute values of $\mathbf{H}$ at δ points. We choose these points in geometric progression and use the interpolant basis algorithm of Section 3.2 to recover $\mathbf{Q}$ and $\mathbf{N}$, as detailed in Section 4.1. The value of $\mathbf{H}$ at $x = \alpha$ is computed following the same approach as above, but over $\mathbb{K}$ instead of $\mathbb{K}[[x]]$. In particular, our implementation directly relies on NTL's extended GCD algorithm over $\mathbb{K} = \mathbb{F}_p$ to compute the vectors $\mathbf{c}_1, \mathbf{c}_2, \mathbf{r}_1, \mathbf{r}_2$.

The next table compares our implementation to the direct approach via evaluation/interpolation; note that the latter approach, while straightforward conceptually, is the state of the art.

$n = d$	Direct	Algo. 6	$n = d$	Direct	Algo. 6
100	1.75	3.48	600	797	653
200	17.4	29.3	700	1343	1081
300	72.3	106	800	2121	1388
400	182	182	900	3203	1760

For these running times, input polynomials were chosen at random with partial degree n both in x and in z ; such polynomials have total degree $2n$, and their resultant has degree $2n^2$. The largest examples have quite significant sizes, but such degrees are not unheard-of in applications, as for instance in the genus-2 point counting algorithms of [1, 12–14]. Overall, with $n = d$, we observe a crossover point around $n = 400$. Besides, in a close match with the analysis above, the parameter m was set to $\lceil n^{0.4} \rceil$ since this gave us the best runtimes. As an example, for $d = 300$, the cost of each individual steps were 65s for computing structured inversions, and 40s for obtaining $\mathbf{Q}$ and its determinant, which is a good balance.

5.3. Characteristic polynomial. We consider the computation of the characteristic polynomial of an element A in $\mathbb{K}[z]/(P)$, for some monic P in $\mathbb{K}[z]$ of degree n . The algorithm we implemented, and which we sketch below, is from [34] and assumes that A and P are generic.

As explained previously, this problem is a particular case of a bivariate resultant, but we rely on another point of view that allows for a better asymptotic cost. Indeed, the characteristic polynomial of A modulo P is by definition the characteristic polynomial of the matrix $\mathbf{M}$ of multiplication by A modulo P . In other words, it is the determinant of the degree-1 matrix $\mathbf{P} = x\mathbf{I} - \mathbf{M} \in \mathbb{K}[x]^{n \times n}$.

The genericity assumption ensures that $\mathbf{M}$ is invertible, hence the power series expansion of $\mathbf{P}^{-1}$ is $\sum_{k \geq 0} -\mathbf{M}^{-k-1}x^k$. Here, we use the top-left $m \times m$ quadrant $\mathbf{H}$ of $\mathbf{P}^{-1}$; it has entries $h_{i,j} \in \mathbb{K}[[x]]$, where

$$h_{i,j,k} := \text{coeff}(h_{i,j}, x^k) = \text{coeff}(-z^j A^{-k-1} \bmod P, z^i).$$

for $0 \leq i, j < m$ and for all $k \geq 0$.

A direct implementation of this idea does not improve on the runtime given in Section 5.1, since it computes $A^{-k-1} \bmod P$ for all $0 \leq k < \delta = 2\lceil n/m \rceil$ and therefore costs $\Omega(n^2/m)$. It turns out that baby-steps giant-steps techniques allow one to compute $h_{i,j,k}$ for $0 \leq i, j < m$ and $0 \leq k < \delta$ in $O(\delta^{(\omega-1)/2}n + mn)$ operations in $\mathbb{K}$. Taking $m = \lceil n^{1/3} \rceil$ minimizes the overall cost, resulting in the runtime $O(n^{(\omega+2)/3}) \subset O(n^{1.46})$.

The following table compares our implementation to NTL's built-in characteristic polynomial algorithm, with random inputs A and P . For such inputs, NTL uses Shoup's algorithm for power projection [39], which runs in time $O(n^{(\omega+1)/2})$.

n	m	NTL	new	n	m	NTL	new
5000	5	0.143	0.225	60000	10	8.45	8.34
20000	8	1.43	1.62	80000	10	16.6	12.1
40000	8	4.69	4.42	100000	10	23.1	17.4

Acknowledgements. Hyun was supported by a Mitacs Globalink award; Neiger was supported by CNRS/INS2I's program for young researchers; Schost was supported by an NSERC Discovery Grant.

References

- [1] S. Ablard, P. Gaudry, and P.-J. Spaenlehauer. 2018. Counting points on genus-3 hyperelliptic curves with explicit real multiplication. In *ANTS-XIII*.
- [2] B. Beckermann and G. Labahn. 1994. A Uniform Approach for the Fast Computation of Matrix-Type Pad Approximants. *SIAM J. Matrix Anal. Appl.* 15, 3 (1994), 804–823.
- [3] B. Beckermann and G. Labahn. 2000. Fraction-free computation of matrix rational interpolant and matrix gcds. *SIAM J. Matrix Anal. Appl.* 22, 1 (2000), 114–144.
- [4] B. Beckermann, G. Labahn, and G. Villard. 1999. Shifted Normal Forms of Polynomial Matrices. In *ISSAC'99*. ACM, 189–196.
- [5] A. Bostan, G. Lecerf, and . Schost. 2003. Tellegen's Principle Into Practice. In *ISSAC'03*. ACM, 37–44.
- [6] A. Bostan and . Schost. 2005. Polynomial evaluation and interpolation on special sets of points. *J. Complexity* 21, 4 (2005), 420–446.
- [7] D. Coppersmith. 1994. Solving homogeneous linear equations over GF(2) via block Wiedemann algorithm. *Math. Comp.* 62, 205 (1994), 333–350.
- [8] D. Coppersmith and S. Winograd. 1990. Matrix multiplication via arithmetic progressions. *J. Symbolic Comput.* 9, 3 (1990), 251–280.
- [9] J. Dixon. 1982. Exact solution of linear equations using p -adic expansions. *Numer. Math.* 40 (1982), 137–141.
- [10] J. Doliskani, P. Giorgi, R. Lebreton, and . Schost. 2018. Simultaneous conversions with the residue number system using linear algebra. *ACM Trans. Math. Software* 44, 3 (2018), 1–21.
- [11] J. von zur Gathen and J. Gerhard. 2013. *Modern Computer Algebra (third edition)*. Cambridge University Press.
- [12] P. Gaudry and R. Harley. 2000. Counting points on hyperelliptic curves over finite fields. In *ANTS-IV (LNCS)*, Vol. 1838. Springer-Verlag, 313–332.
- [13] P. Gaudry and . Schost. 2004. Construction of secure random curves of genus 2 over prime fields. In *Eurocrypt'04 (LNCS)*, Vol. 3027. Springer, 239–256.
- [14] P. Gaudry and . Schost. 2012. Point-counting in genus 2 over prime fields. *J. Symbolic Comput.* 47, 4 (2012), 368400.
- [15] P. Giorgi, C.-P. Jeannerod, and G. Villard. 2003. On the complexity of polynomial matrix computations. In *ISSAC'03*. ACM, 135–142.
- [16] S. Gupta, S. Sarkar, A. Storjohann, and J. Valeriote. 2012. Triangular x -basis decompositions and derandomization of linear algebra algorithms over $K[x]$. *J. Symbolic Comput.* 47, 4 (2012), 422–453.
- [17] G. Hanrot, M. Quercia, and P. Zimmermann. 2004. The Middle Product Algorithm, I. *Appl. Algebra Engrg. Comm. Comp.* 14, 6 (2004), 415–438.
- [18] W. Hart, F. Johansson, and S. Pancratz. 2015. FLINT: Fast Library for Number Theory. Version 2.5.2, <http://flintlib.org>.
- [19] D. Harvey. 2009. A cache-friendly truncated FFT. *Theoretical Computer Science* 410, 27–29 (2009), 2649–2658.
- [20] D. Harvey. 2014. Faster arithmetic for number-theoretic transforms. *Journal of Symbolic Computation* 60 (2014), 113–119.
- [21] Seung Gyu Hyun, Romain Lebreton, and ric Schost. 2017. Algorithms for structured linear systems solving and their implementation. In *ISSAC'17*. ACM Press, 205–212.
- [22] S. G. Hyun, V. Neiger, H. Rahkooy, and . Schost. 2017. Block-Krylov techniques in the context of sparse-FGLM algorithms. (2017). arXiv:1712.04177
- [23] C.-P. Jeannerod, V. Neiger, . Schost, and G. Villard. 2016. Fast computation of minimal interpolation bases in Popov form for arbitrary shifts. In *ISSAC'16*. ACM, 295–302.
- [24] C.-P. Jeannerod, V. Neiger, . Schost, and G. Villard. 2017. Computing minimal interpolation bases. *J. Symbolic Comput.* 83 (2017), 272–314.
- [25] C.-P. Jeannerod, V. Neiger, and G. Villard. 2018. Fast computation of approximant bases in canonical form. *Journal of Symbolic Computation* (2018). to appear.
- [26] E. Kaltofen. 1995. Analysis of Coppersmith's block Wiedemann algorithm for the parallel solution of sparse linear systems. *Math. Comp.* 64, 210 (1995), 777–806.
- [27] M. Kaminski, D.G. Kirkpatrick, and N.H. Bshouty. 1988. Addition requirements for matrix and transposed matrix products. *J. Algorithms* 9, 3 (1988), 354–364.
- [28] G. Labahn, D. K. Choi, and S. Cabay. 1990. The inverses of block Hankel and block Toeplitz matrices. *SIAM J. Comput.* 19, 1 (1990), 98–123.
- [29] G. Labahn, V. Neiger, and W. Zhou. 2017. Fast, deterministic computation of the Hermite normal form and determinant of a polynomial matrix. *J. Complexity* (in press) (2017).
- [30] F. Le Gall. 2014. Powers of Tensors and Fast Matrix Multiplication. In *ISSAC'14*. ACM, 296–303.
- [31] F. Le Gall and F. Urrutia. 2018. Improved rectangular matrix multiplication using powers of the Coppersmith-Winograd tensor. In *SODA '18*. SIAM, 1029–1046.
- [32] P. L. Montgomery. 2005. Five, six, and seven-term Karatsuba-like formulae. *IEEE Trans. Comput.* 54, 3 (2005), 362–369.
- [33] Vincent Neiger, Johan Rosenkilde, and Grigory Solomatov. 2018. Computing Popov and Hermite Forms of Rectangular Polynomial Matrices. In *ISSAC'18*. ACM, New York, NY, USA, 295–302.
- [34] V. Neiger, B. Salvy, . Schost, and G. Villard. 2019. In preparation.
- [35] J. S. R. Nielsen and P. Beelen. 2015. Sub-quadratic decoding of one-point Hermitian codes. *IEEE Trans. Inf. Theory* 61, 6 (June 2015), 3225–3240.
- [36] Victor Pan. 1988. Computing the determinant and the characteristic polynomial of a matrix via solving linear systems of equations. *Inform. Process. Lett.* 28, 2 (1988), 71–75.
- [37] D. Reischert. 1997. Asymptotically fast computation of subresultants. In *ISSAC'97*. ACM, 233–240.
- [38] G. Schulz. 1933. Iterative Berechnung der reziproken Matrix. *Zeitschrift fr Angewandte Mathematik und Mechanik* 13, 1 (1933), 57–59.
- [39] V. Shoup. 1994. Fast construction of irreducible polynomials over finite fields. *J. Symbolic Comput.* 17, 5 (1994), 371–391.
- [40] V. Shoup. 2018. NTL: A Library for doing Number Theory, version 11.3.2. <http://www.shoup.net>.
- [41] A. Storjohann. 2003. High-order lifting and integrality certification. *J. Symbolic Comput.* 36, 3–4 (2003), 613–648.
- [42] A. Storjohann. 2006. Notes on computing minimal approximant bases. In *Challenges in Symbolic Computation Software (Dagstuhl Seminar Proceedings)*.
- [43] V. Strassen. 1969. Gaussian elimination is not optimal. *Numer. Math.* 13 (1969), 354–356.
- [44] The CADO-NFS Development Team. 2017. CADO-NFS, An implementation of the number field sieve algorithm. <http://cado-nfs.gforge.inria.fr/> Release 2.3.0.
- [45] The FFLAS-FFPACK Group. 2019. FFLAS-FFPACK: Finite Field Linear Algebra Subroutines / Package, version 2.3.2. <http://github.com/linbox-team/fflas-ffpack>.
- [46] The LinBox Group. 2018. LinBox: Linear algebra over black-box matrices, version 1.5.3. <https://github.com/linbox-team/linbox/>.
- [47] M. Van Barel and A. Bultheel. 1992. A general module theoretic framework for vector M-Pad and matrix rational interpolation. *Numer. Algorithms* 3 (1992), 451–462.
- [48] G. Villard. 1997. Further analysis of Coppersmith's block Wiedemann algorithm for the solution of sparse linear systems. In *ISSAC'97*. ACM, 32–39.
- [49] G. Villard. 2018. On computing the resultant of generic bivariate polynomials. In *ISSAC'18*. ACM, 391–398.
- [50] A. Waksman. 1970. On Winograd's Algorithm for Inner Products. *IEEE Transactions On Computers* C-19 (1970), 360–361.
- [51] S. Winograd. 1968. A New Algorithm for Inner Product. *IEEE Trans. Comput.* C-17, 7 (1968), 693–694.
- [52] S. Winograd. 1971. On multiplication of 2×2 matrices. *Linear Algebra Appl.* 4 (1971), 381–388.
- [53] W. Zhou and G. Labahn. 2012. Efficient Algorithms for Order Basis Computation. *J. Symbolic Comput.* 47, 7 (2012), 793–819.
- [54] W. Zhou, G. Labahn, and A. Storjohann. 2012. Computing Minimal Nullspace Bases. In *ISSAC'12*. ACM, 366–373.