

Fast computation of minimal interpolation bases in Popov form for arbitrary shifts

Claude-Pierre Jeannerod
Inria, Université de Lyon
Laboratoire LIP (CNRS, Inria, ENSL, UCBL)

Éric Schost
University of Waterloo
David R. Cheriton School of Computer Science

Vincent Neiger
ENS de Lyon, Université de Lyon
Laboratoire LIP (CNRS, Inria, ENSL, UCBL)

Gilles Villard
CNRS, Université de Lyon
Laboratoire LIP (CNRS, Inria, ENSL, UCBL)

ABSTRACT

We compute minimal bases of solutions for a general interpolation problem, which encompasses Hermite-Padé approximation and constrained multivariate interpolation, and has applications in coding theory and security. The problem asks to find univariate polynomial relations between m vectors of size σ ; these relations should have small degree with respect to an input degree shift. For an arbitrary shift, we propose an algorithm for the computation of an interpolation basis in shifted Popov form with a cost of $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ field operations, where ω is the exponent of matrix multiplication and the notation $\mathcal{O}^\sim(\cdot)$ omits logarithmic terms.

Earlier works, in the case of Hermite-Padé approximation [34] and in the general interpolation case [18], compute non-normalized bases. Since for arbitrary shifts such bases may have size $\Theta(m^2\sigma)$, the cost bound $\mathcal{O}^\sim(m^{\omega-1}\sigma)$ was feasible only with restrictive assumptions on the shift that ensure small output sizes. The question of handling arbitrary shifts with the same complexity bound was left open.

To obtain the target cost for any shift, we strengthen the properties of the output bases, and of those obtained during the course of the algorithm: all the bases are computed in shifted Popov form, whose size is always $\mathcal{O}(m\sigma)$. Then, we design a divide-and-conquer scheme. We recursively reduce the initial interpolation problem to sub-problems with more convenient shifts by first computing information on the degrees of the intermediate bases.

CCS Concepts

•Mathematics of computing → Computations on matrices; Computations on polynomials; Interpolation; •Computing methodologies → Algebraic algorithms;

Keywords

M-Padé approximation, Hermite-Padé approximation, order basis, polynomial matrix, shifted Popov form.

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

ISSAC '16 July 20–22, 2016, Waterloo, ON, Canada

© 2022 ACM. ISBN ?...\$??..??

DOI: ?

1. INTRODUCTION

1.1 Problem and main result

We focus on the following interpolation problem from [31, 2]. For a field \mathbb{K} and some positive integer σ , we have as input m vectors $\mathbf{e}_1, \dots, \mathbf{e}_m$ in $\mathbb{K}^{1 \times \sigma}$, seen as the rows of a matrix $\mathbf{E} \in \mathbb{K}^{m \times \sigma}$. We also have a *multiplication matrix* $\mathbf{J} \in \mathbb{K}^{\sigma \times \sigma}$ which specifies the multiplication of vectors $\mathbf{e} \in \mathbb{K}^{1 \times \sigma}$ by polynomials $p \in \mathbb{K}[X]$ as $p \cdot \mathbf{e} = \mathbf{e}p(\mathbf{J})$. Then, we want to find $\mathbb{K}[X]$ -linear relations between these vectors, that is, some $\mathbf{p} = (p_1, \dots, p_m) \in \mathbb{K}[X]^m$ such that $\mathbf{p} \cdot \mathbf{E} = p_1 \cdot \mathbf{e}_1 + \dots + p_m \cdot \mathbf{e}_m = 0$. Such a \mathbf{p} is called an *interpolant for* (\mathbf{E}, \mathbf{J}) .

Hereafter, the matrix \mathbf{J} is in Jordan canonical form: this assumption is satisfied in many interesting applications, as explained below. The notion of interpolant we consider is directly related to the one introduced in [31, 2]. Suppose that \mathbf{J} has n Jordan blocks of dimensions $\sigma_1 \times \sigma_1, \dots, \sigma_n \times \sigma_n$ and with respective eigenvalues x_1, \dots, x_n ; in particular, $\sigma = \sigma_1 + \dots + \sigma_n$. Then, one may identify \mathbb{K}^σ with

$$\mathfrak{F} = \mathbb{K}[X]/(X^{\sigma_1}) \times \dots \times \mathbb{K}[X]/(X^{\sigma_n}),$$

by mapping any $\mathbf{f} = (f_1, \dots, f_n)$ in \mathfrak{F} to the vector $\mathbf{e} \in \mathbb{K}^\sigma$ made from the concatenation of the coefficient vectors of f_1, \dots, f_n . Over \mathfrak{F} , the $\mathbb{K}[X]$ -module structure on \mathbb{K}^σ given by $p \cdot \mathbf{e} = \mathbf{e}p(\mathbf{J})$ becomes

$$p \cdot \mathbf{f} = (p(X + x_1)f_1 \bmod X^{\sigma_1}, \dots, p(X + x_n)f_n \bmod X^{\sigma_n}).$$

Now, if $(\mathbf{e}_1, \dots, \mathbf{e}_m) \in \mathbb{K}^{m \times \sigma}$ is associated to $(\mathbf{f}_1, \dots, \mathbf{f}_m) \in \mathfrak{F}^m$, with $\mathbf{f}_i = (f_{i,1}, \dots, f_{i,n})$ and $f_{i,j}$ in $\mathbb{K}[X]/(X^{\sigma_j})$ for all i, j , the relation $p_1 \cdot \mathbf{e}_1 + \dots + p_m \cdot \mathbf{e}_m = 0$ means that for all j in $\{1, \dots, n\}$, we have

$$p_1(X + x_j)f_{1,j} + \dots + p_m(X + x_j)f_{m,j} = 0 \bmod X^{\sigma_j};$$

applying a translation by $-x_j$, this is equivalent to

$$p_1f_{1,j}(X - x_j) + \dots + p_mf_{m,j}(X - x_j) = 0 \bmod (X - x_j)^{\sigma_j}.$$

Thus, in terms of vector M-Padé approximation as in [31, 2], (p_1, \dots, p_m) is an interpolant for $(\mathbf{f}_1, \dots, \mathbf{f}_m)$, x_1, \dots, x_n , and $\sigma_1, \dots, \sigma_n$.

The set of all interpolants for (\mathbf{E}, \mathbf{J}) is a free $\mathbb{K}[X]$ -module of rank m . We are interested in computing a basis of this module, represented as a matrix in $\mathbb{K}[X]^{m \times m}$ and called an *interpolation basis for* (\mathbf{E}, \mathbf{J}) . Its rows are interpolants for (\mathbf{E}, \mathbf{J}) , and any interpolant for (\mathbf{E}, \mathbf{J}) can be written as a unique $\mathbb{K}[X]$ -linear combination of its rows.

Besides, we look for interpolants that have some type of minimal degree. Following [31, 34], for a nonzero $\mathbf{p} = [p_1, \dots, p_m] \in \mathbb{K}[X]^{1 \times m}$ and a *shift* $\mathbf{s} = (s_1, \dots, s_m) \in \mathbb{Z}^m$, we define the *s-degree* of \mathbf{p} as $\max_{1 \leq j \leq m} (\deg(p_j) + s_j)$. Up to a change of sign, this notion of *s-degree* is equivalent to the one in [3] and to the notion of *defect* from [1, Definition 3.1].

Then, the *s-row degree* of a matrix $\mathbf{P} \in \mathbb{K}[X]^{k \times m}$ of rank k is the tuple $\text{rdeg}_s(\mathbf{P}) = (d_1, \dots, d_k) \in \mathbb{Z}^k$ with d_i the *s-degree* of the i -th row of \mathbf{P} . The *s-leading matrix* of $\mathbf{P} = [p_{ij}]_{i,j}$ is the matrix in $\mathbb{K}^{k \times m}$ whose entry (i, j) is the coefficient of degree $d_i - s_j$ of p_{ij} . Then, \mathbf{P} is *s-reduced* if its *s-leading matrix* has rank k ; see [3].

Our aim is to compute an *s-minimal* interpolation basis for (\mathbf{E}, \mathbf{J}) , that is, one which is *s-reduced*: equivalently, it is an interpolation basis whose *s-row degree*, once written in nondecreasing order, is lexicographically minimal. This corresponds to Problem 1 below. In particular, an interpolant of minimal degree can be read off from an *s-minimal* interpolation basis for the *uniform shift* $\mathbf{s} = \mathbf{0}$.

PROBLEM 1 (MINIMAL INTERPOLATION BASIS).

Input:

- the base field \mathbb{K} ,
- the dimensions m and σ ,
- a matrix $\mathbf{E} \in \mathbb{K}^{m \times \sigma}$,
- a Jordan matrix $\mathbf{J} \in \mathbb{K}^{\sigma \times \sigma}$,
- a shift $\mathbf{s} \in \mathbb{Z}^m$.

Output: an *s-minimal* interpolation basis for (\mathbf{E}, \mathbf{J}) .

A well-known particular case of this problem is Hermite-Padé approximation, that is, the computation of *order bases* (or σ -bases, or minimal approximant bases), where \mathbf{J} has only eigenvalue 0. Previous work on this case includes [1, 14, 30, 34] with algorithms focusing on \mathbf{J} with n blocks of identical size σ/n . For a shift $\mathbf{s} \in \mathbb{N}^m$ with nonnegative entries, we write $|\mathbf{s}|$ for the sum of its entries. Then, in this context, the cost bound $\mathcal{O}^{\sim}(m^{\omega-1}\sigma)$ has been obtained under each of the following assumptions:

(H_1) $\max(\mathbf{s}) - \min(\mathbf{s}) \in \mathcal{O}(\sigma/m)$ in [34, Theorem 5.3] and more generally $|\mathbf{s} - \min(\mathbf{s})| \in \mathcal{O}(\sigma)$ in [33, Section 4.1];

(H_2) $|\max(\mathbf{s}) - \mathbf{s}| \in \mathcal{O}(\sigma)$ in [34, Theorem 6.14].

These assumptions imply in particular that any *s-minimal* basis has size in $\mathcal{O}(m\sigma)$, where by *size* we mean the number of field elements used to represent the matrix.

An interesting example of a shift not covered by (H_1) or (H_2) is $\mathbf{h} = (0, \sigma, 2\sigma, \dots, (m-1)\sigma)$ which is related to the Hermite form [3, Lemma 2.6]. In general, as detailed in Appendix A, one may assume without loss of generality that $\min(\mathbf{s}) = 0$, $\max(\mathbf{s}) \in \mathcal{O}(m\sigma)$, and $|\mathbf{s}| \in \mathcal{O}(m^2\sigma)$.

There are also applications of Problem 1 to multivariate interpolation, where \mathbf{J} is not nilpotent anymore, and for which we have neither (H_1) nor (H_2), as we will see in Subsection 1.3. It was left as an open problem in [34, Section 7] to obtain algorithms with cost bound $\mathcal{O}^{\sim}(m^{\omega-1}\sigma)$ for such matrices \mathbf{J} and for arbitrary shifts. In this paper, we solve this open problem.

An immediate challenge is that for an arbitrary shift \mathbf{s} , the size of an *s-minimal* interpolation basis may be beyond

our target cost: we show this in Appendix B with an example of Hermite-Padé approximation. Our answer is to compute a basis in *s-Popov form*: among its many interesting features, it can be represented using at most $m(\sigma+1)$ elements from \mathbb{K} , and it is canonical: for every nonsingular $\mathbf{A} \in \mathbb{K}[X]^{m \times m}$ and $\mathbf{s} \in \mathbb{Z}^m$, there is a unique matrix \mathbf{P} in *s-Popov form* which is left-unimodularly equivalent to \mathbf{A} . We use the definition from [2, Section 7], phrased using the notion of pivot [19, Section 6.7.2].

DEFINITION 1.1 (PIVOT OF A ROW). Let $\mathbf{p} = [p_j]_j \in \mathbb{K}[X]^{1 \times m}$ be a nonzero row vector and let $\mathbf{s} \in \mathbb{Z}^m$. The *s-pivot index* of \mathbf{p} is the largest index $j \in \{1, \dots, m\}$ such that $\text{rdeg}_s(\mathbf{p}) = \deg(p_j) + s_j$; then, p_j and $\deg(p_j)$ are called the *s-pivot entry* and the *s-pivot degree* of \mathbf{p} .

DEFINITION 1.2 (POPOV FORM). Let $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$ be nonsingular and let $\mathbf{s} \in \mathbb{Z}^m$. Then, \mathbf{P} is said to be in *s-Popov form* if its *s-pivot entries* are monic and on its diagonal, and if in each column of \mathbf{P} the nonpivot entries have degree less than the pivot entry.

We call *s-Popov interpolation basis* for (\mathbf{E}, \mathbf{J}) the unique interpolation basis for (\mathbf{E}, \mathbf{J}) which is in *s-Popov form*; in particular, it is an *s-minimal* one. For small values of σ , namely $\sigma \in \mathcal{O}(m)$, we gave in [18, Section 7] an algorithm which computes the *s-Popov* interpolation basis in $\mathcal{O}^{\sim}(\sigma^{\omega-1}m)$ operations for an arbitrary \mathbf{s} [18, Theorem 1.4]. Hence, in what follows, we focus on the case $m \in \mathcal{O}(\sigma)$.

We use the convenient assumption that \mathbf{J} is given to us as a list of eigenvalues and block sizes:

$$\mathbf{J} = ((x_1, \sigma_{1,1}), \dots, (x_1, \sigma_{1,r_1}), \dots, (x_t, \sigma_{t,1}), \dots, (x_t, \sigma_{t,r_t})),$$

for some pairwise distinct eigenvalues x_1, \dots, x_t , with $r_1 \geq \dots \geq r_t$ and $\sigma_{i,1} \geq \dots \geq \sigma_{i,r_i}$ for all i ; we say that this representation is *standard*.

THEOREM 1.3. Assuming that $\mathbf{J} \in \mathbb{K}^{\sigma \times \sigma}$ is a Jordan matrix given by a standard representation, there is a deterministic algorithm which solves Problem 1 using

$$\begin{aligned} \mathcal{O}(m^{\omega-1}M(\sigma) \log(\sigma) \log(\sigma/m)^2) & \quad \text{if } \omega > 2, \\ \mathcal{O}(mM(\sigma) \log(\sigma) \log(\sigma/m)^2 \log(m)^3) & \quad \text{if } \omega = 2 \end{aligned}$$

operations in \mathbb{K} and returns the *s-Popov* interpolation basis for (\mathbf{E}, \mathbf{J}) .

In this result, $M(\cdot)$ is such that polynomials of degree at most d in $\mathbb{K}[X]$ can be multiplied using $M(d)$ operations in \mathbb{K} , and $M(\cdot)$ satisfies the super-linearity properties of [13, Chapter 8]. It follows from [8] that $M(d)$ can be taken in $\mathcal{O}(d \log(d) \log(\log(d)))$. The exponent ω is so that we can multiply $m \times m$ matrices in $\mathcal{O}(m^\omega)$ ring operations on any ring, the best known bound being $\omega < 2.38$ [11, 22].

Compared to our work in [18], our algorithm here has two key new features:

- it supports arbitrary shifts with a cost $\mathcal{O}^{\sim}(m^{\omega-1}\sigma)$;
- it computes the basis in *s-Popov form*.

To the best of our knowledge, no algorithm for Problem 1 with cost $\mathcal{O}^{\sim}(m^{\omega-1}\sigma)$ was known previously for *arbitrary* shifts, even for the specific case of order basis computation.

If \mathbf{J} is given as an arbitrary list $((x_1, \sigma_1), \dots, (x_n, \sigma_n))$, we can reorder it (and permute the columns of \mathbf{E} accordingly) to obtain an equivalent standard representation in

time $\mathcal{O}(M(\sigma) \log(\sigma)^3)$ [5, Proposition 12]; if \mathbb{K} is equipped with an order, and if we assume that comparisons take unit time, this can of course be done in time $\mathcal{O}(\sigma \log(\sigma))$.

1.2 Overview of our approach

Several previous algorithms for order basis computation, such as those in [1, 14], follow a divide-and-conquer scheme inspired by the Knuth-Schönhage-Moenck algorithm [20, 29, 23]. This paper builds on our previous work in [18], where we extended this recursive approach to more general interpolation problems. However, the main algorithm in [18] does not handle an arbitrary shift \mathbf{s} with a satisfactory complexity; here, we use it as a black box, after showing how to reduce the problem to a new one with suitable shift.

Let \mathbf{E} , \mathbf{J} , and \mathbf{s} be our input, and write $\mathbf{J}^{(1)}$ and $\mathbf{J}^{(2)}$ for the $\sigma/2 \times \sigma/2$ leading and trailing principal submatrices of \mathbf{J} . First, compute an \mathbf{s} -minimal interpolation basis $\mathbf{P}^{(1)}$ for $\mathbf{J}^{(1)}$ and the first $\sigma/2$ columns of \mathbf{E} ; then, compute the last $\sigma/2$ columns $\mathbf{E}^{(2)}$ of the residual $\mathbf{P}^{(1)} \cdot \mathbf{E}$; then, compute a \mathbf{t} -minimal interpolation basis $\mathbf{P}^{(2)}$ for $(\mathbf{E}^{(2)}, \mathbf{J}^{(2)})$ with $\mathbf{t} = \text{rdeg}_{\mathbf{s}}(\mathbf{P}^{(1)})$; finally, return the matrix product $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$.

This approach allows to solve Problem 1 using $\mathcal{O}(m^\omega \sigma)$ operations in \mathbb{K} . In the case of Hermite-Padé approximation, this is the divide-and-conquer algorithm in [1]. Besides, an \mathbf{s} -minimal basis computed by this method has degree at most σ and thus size in $\mathcal{O}(m^2 \sigma)$, and there are indeed instances of Problem 1 for which this size reaches $\Theta(m^2 \sigma)$. In Appendix B, we show such an instance for the algorithm in [1], in the case of Hermite-Padé approximation.

It is known that the average degree of the rows of any \mathbf{s} -minimal interpolation basis is at most $(\sigma + \xi)/m$, where $\xi = |\mathbf{s} - \min(\mathbf{s})|$ [31, Theorem 4.1]. In [18], focusing on the case where ξ is small compared to σ , and preserving such a property in recursive calls via changes of shifts, we obtained the cost bound

$$\mathcal{O}(m^{\omega-1} M(\sigma) \log(\sigma) \log(\sigma/m) + m^{\omega-1} M(\xi) \log(\xi/m)) \quad (1)$$

to solve Problem 1; this cost is for $\omega > 2$, and a similar one holds for $\omega = 2$, both being in $\mathcal{O}(m^{\omega-1}(\sigma + \xi))$. The fundamental reason for this kind of improvement over $\mathcal{O}(m^\omega \sigma)$, already seen with [34], is that one controls the average row degree of the bases $\mathbf{P}^{(2)}$ and $\mathbf{P}^{(1)}$, and of their product $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$.

This result is $\mathcal{O}(m^{\omega-1} \sigma)$ for ξ in $\mathcal{O}(\sigma)$. The main difficulty to extend it to any shift \mathbf{s} is to control the size of the computed bases: the Hermite-Padé example pointed out above corresponds to $\xi = \Theta(m\sigma)$ and leads to an output of size $\Theta(m^2 \sigma)$ for the algorithm of [18] as well.

The key ingredient to control this size is to work with bases in \mathbf{s} -Popov form: for any \mathbf{s} , the \mathbf{s} -Popov interpolation basis \mathbf{P} for (\mathbf{E}, \mathbf{J}) has average column degree at most σ/m and size at most $m(\sigma + 1)$, as detailed in Section 2.

Now, suppose that we have computed recursively the bases $\mathbf{P}^{(2)}$ and $\mathbf{P}^{(1)}$ in \mathbf{s} - and \mathbf{t} -Popov form; we want to output the \mathbf{s} -Popov form \mathbf{P} of $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$. In general, this product is not normalized and may have size $\Theta(m^2 \sigma)$: its computation is beyond our target cost. Thus, one main idea is that we will *not* rely on polynomial matrix multiplication to combine the bases obtained recursively; instead, we use a minimal interpolation basis computation for a shift that has good properties as explained below.

An important remark is that if we know *a priori* the column degree δ of \mathbf{P} , then the problem becomes easier. This

idea was already used in algorithms for the Hermite form \mathbf{H} of a polynomial matrix [15, 33], which first compute the column degree δ of \mathbf{H} , and then obtain \mathbf{H} as a submatrix of some minimal nullspace basis for a shift involving $-\delta$.

In Section 4, we study the problem of computing the \mathbf{s} -Popov interpolation basis \mathbf{P} for (\mathbf{E}, \mathbf{J}) having its column degree δ as an additional input. We show that this reduces to the computation of a \mathbf{d} -minimal interpolation basis \mathbf{R} with the specific shift $\mathbf{d} = -\delta$. The properties of this shift \mathbf{d} allow us first to compute \mathbf{R} in $\mathcal{O}(m^{\omega-1} \sigma)$ operations using the partial linearization framework from [30, Section 3] and the minimal interpolation basis algorithm in [18, Section 3], and second to easily retrieve \mathbf{P} from \mathbf{R} .

Still, in general we do not know δ . We will thus compute it, relying on a variation of the divide-and-conquer strategy at the beginning of this subsection. We stop the recursion as soon as $\sigma \leq m$, in which case we do not need δ to achieve efficiency: the algorithm from [18, Section 7] computes the \mathbf{s} -Popov interpolation basis in $\mathcal{O}(\sigma^{\omega-1} m)$ operations for any \mathbf{s} [18, Theorem 1.4]. Then, we show in Section 3 that from $\mathbf{P}^{(1)}$ and $\mathbf{P}^{(2)}$ computed recursively *in shifted Popov form*, we can obtain δ for free. Finally, instead of considering $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$, we use the knowledge of δ to compute the basis \mathbf{P} from scratch as explained in the previous paragraph.

This summarizes our main algorithm, which is presented in Section 2.

1.3 Previous work and applications

As a particular case of Problem 1, when all the eigenvalues of \mathbf{J} are zero, we obtain the following complexity result about *order basis* computation [34, Definition 2.2].

THEOREM 1.4. *Let $m, n \in \mathbb{Z}_{>0}$, let $(\sigma_1, \dots, \sigma_n) \in \mathbb{Z}_{>0}^n$, let $\mathbf{s} \in \mathbb{Z}^m$, and let $\mathbf{F} \in \mathbb{K}[X]^{m \times n}$ with its j -th column $\mathbf{F}_{*,j}$ of degree less than σ_j . The unique basis $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$ in \mathbf{s} -Popov form of the $\mathbb{K}[X]$ -module of approximants*

$$\{\mathbf{p} \in \mathbb{K}[X]^{1 \times m} \mid \mathbf{p}\mathbf{F}_{*,j} = 0 \text{ mod } X^{\sigma_j} \text{ for each } j\}$$

can be computed deterministically using

$$\begin{aligned} \mathcal{O}(m^{\omega-1} M(\sigma) \log(\sigma) \log(\sigma/m)^2) & \quad \text{if } \omega > 2, \\ \mathcal{O}(mM(\sigma) \log(\sigma) \log(\sigma/m)^2 \log(m)^3) & \quad \text{if } \omega = 2 \end{aligned}$$

operations in \mathbb{K} , where $\sigma = \sigma_1 + \dots + \sigma_n$.

Previous work on this problem includes [1, 14, 30, 34, 18], mostly with identical orders $\sigma_1 = \dots = \sigma_n$; an interesting particular case is Hermite-Padé approximation with $n = 1$. To simplify matters, for all our comparisons, we consider $\omega > 2$. For order basis computation with $\sigma_1 = \dots = \sigma_n$ and $n \leq m$, the cost bound $\mathcal{O}(m^\omega M(\sigma/m) \log(\sigma/n))$ was achieved in [34] under either of the assumptions (H_1) and (H_2) on the shift. Still, the corresponding algorithm returns a basis \mathbf{P} which is only \mathbf{s} -reduced, and because both the shift \mathbf{s} and the degrees in \mathbf{P} may be unbalanced, one cannot directly rely on the fastest known normalization algorithm [28] to compute the \mathbf{s} -Popov form of \mathbf{P} within the target cost.

Another application of Problem 1 is a multivariate interpolation problem that arises for example in the first step of algorithms for the list-decoding of Parvaresh-Vardy codes [26] and of folded Reed-Solomon codes [16], as well as in robust Private Information Retrieval [12]. The bivariate case corresponds to the interpolation steps of Kötter and Vardy's soft-decoding [21] and Guruswami and Sudan's list-decoding [17] algorithms for Reed-Solomon codes.

Given a set of points in \mathbb{K}^{r+1} and associated multiplicities, this problem asks to find a multivariate polynomial $Q(X, Y_1, \dots, Y_r)$ such that: (a) Q has prescribed exponents for the Y variables, so that the problem can be linearized with respect to Y , leaving us with a linear algebra problem over $\mathbb{K}[X]$; (b) Q vanishes at all the given points with their multiplicities, inducing a structure of $\mathbb{K}[X]$ -module on the set of solutions; (c) Q has some type of minimal weighted degree, which can be seen as the minimality of the shifted degree of the vector over $\mathbb{K}[X]$ that represents Q .

Following the coding theory context [17, 26], given a point $(x, y) \in \mathbb{K} \times \mathbb{K}^r$ and a set of exponents $\mu \subset \mathbb{N}^{r+1}$, we say that the polynomial $Q(X, Y) \in \mathbb{K}[X, Y_1, \dots, Y_r]$ vanishes at (x, y) with multiplicity support μ if the shifted polynomial $Q(X + x, Y + y)$ has no monomial with exponent in μ . We will only consider supports that are *stable under division*, meaning that if $(\gamma_0, \gamma_1, \dots, \gamma_r)$ is in μ , then any $(\gamma'_0, \gamma'_1, \dots, \gamma'_r)$ with $\gamma'_j \leq \gamma_j$ for all j is also in μ .

Now, given a set of exponents $\Gamma \subset \mathbb{N}^r$, we represent $Q(X, Y) = \sum_{\gamma \in \Gamma} p_\gamma Y^\gamma$ as the row $\mathbf{p} = [p_\gamma]_{\gamma \in \Gamma} \in \mathbb{K}[X]^{1 \times m}$ where m is the cardinality of Γ . Again, we assume that the exponent set Γ is stable under division; then, the set of solutions is a free $\mathbb{K}[X]$ -module of rank m . In the mentioned applications, we typically have $\Gamma = \{(\gamma_1, \dots, \gamma_r) \in \mathbb{N}^r \mid \gamma_1 + \dots + \gamma_r \leq \ell\}$ for an integer ℓ called the *list-size parameter*.

Besides, we are given some weights $\mathbf{w} = (w_1, \dots, w_r) \in \mathbb{N}^r$ on the variables $Y = Y_1, \dots, Y_r$, and we are looking for $Q(X, Y)$ which has minimal \mathbf{w} -weighted degree, which is the degree in X of the polynomial

$$\begin{aligned} Q(X, X^{w_1} Y_1, \dots, X^{w_r} Y_r) \\ = \sum_{\gamma \in \Gamma} p_\gamma X^{\gamma_1 w_1 + \dots + \gamma_r w_r} Y_1^{\gamma_1} \dots Y_r^{\gamma_r}. \end{aligned}$$

This is exactly requiring that the \mathbf{s} -degree of $\mathbf{p} = [p_\gamma]_\gamma$ be minimal, for $\mathbf{s} = [\gamma_1 w_1 + \dots + \gamma_r w_r]_\gamma$. We note that it is sometimes important, for example in [12], to return a whole \mathbf{s} -minimal interpolation basis and not only one interpolant of small \mathbf{s} -degree.

PROBLEM 2 (MULTIVARIATE INTERPOLATION).

Input:

- number of Y variables $r > 0$,
- set $\Gamma \subset \mathbb{N}^r$ of cardinality m , stable under division,
- pairwise distinct points $\{(x_k, y_k) \in \mathbb{K} \times \mathbb{K}^r\}_{1 \leq k \leq p}$,
- supports $\{\mu_k \subset \mathbb{N}^{r+1}\}_{1 \leq k \leq p}$, stable under division,
- a shift $\mathbf{s} \in \mathbb{Z}^m$.

Output: a matrix $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$ such that

- the rows of \mathbf{P} form a basis of the $\mathbb{K}[X]$ -module

$$\left\{ \mathbf{p} = [p_\gamma]_{\gamma \in \Gamma} \in \mathbb{K}[X]^{1 \times m} \mid \sum_{\gamma \in \Gamma} p_\gamma(X) Y^\gamma \text{ vanishes at } (x_k, y_k) \text{ with support } \mu_k \text{ for } 1 \leq k \leq p \right\},$$

- \mathbf{P} is \mathbf{s} -reduced.

For more details about the reduction from Problem 2 to Problem 1, explaining how to build the input matrices (\mathbf{E}, \mathbf{J}) with \mathbf{J} a Jordan matrix in standard representation, we refer the reader to [18, Subsection 2.4]. In particular, the dimension σ is the sum of the cardinalities of the multiplicity supports. In the mentioned applications to coding theory, we have $m = \binom{r+\ell}{r}$ where ℓ is the list-size parameter; and σ is the so-called *cost* in the soft-decoding context [21, Section III], that is, the number of linear equations when linearizing the problem over \mathbb{K} . As a consequence of Theorem 1.3, we obtain the following complexity result.

THEOREM 1.5. *Let $\sigma = \sum_{1 \leq k \leq p} \#\mu_k$. There is a deterministic algorithm which solves Problem 2 using*

$$\begin{aligned} \mathcal{O}(m^{\omega-1} \mathbf{M}(\sigma) \log(\sigma) \log(\sigma/m)^2) & \quad \text{if } \omega > 2, \\ \mathcal{O}(m \mathbf{M}(\sigma) \log(\sigma) \log(\sigma/m)^2 \log(m)^3) & \quad \text{if } \omega = 2 \end{aligned}$$

operations in \mathbb{K} , and returns the unique basis of solutions which is in \mathbf{s} -Popov form.

Under the assumption that the x_k are pairwise distinct, the cost bound $\mathcal{O}(m^{\omega-1} \mathbf{M}(\sigma) \log(\sigma)^2)$ was achieved for an arbitrary shift using fast structured linear algebra [9, Theorems 1 and 2], following work by [25, 27, 32]. However, the corresponding algorithm is randomized and returns only one interpolant of small \mathbf{s} -degree. For a broader overview of previous work on this problem, we refer the reader to the introductory sections of [34, 4, 9] and to [18, Section 2].

The term $\mathcal{O}(m^{\omega-1} \mathbf{M}(\xi) \log(\xi/m))$ reported in (1) for the cost of the algorithm of [18] can be neglected if $\xi \in \mathcal{O}(\sigma)$; this is for instance satisfied in the context of bivariate interpolation for soft- or list-decoding of Reed-Solomon codes [18, Sections 2.5 and 2.6]. However, we do not have this bound on ξ in the list-decoding of Parvaresh-Vardy codes and folded Reed-Solomon codes and in Private Information Retrieval. Thus, in these cases our algorithm achieves the best known cost bound, improving upon [7, 6, 10, 12, 18].

2. FAST POPOV INTERPOLATION BASIS

In this section, we present our main result, Algorithm 1. It relies on three subroutines; two of them are from [18], while the third is a key new ingredient, detailed in Section 4.

- **LINEARIZATIONMIB** [18, Algorithm 9] solves the base case $\sigma \leq m$ using linear algebra over \mathbb{K} . The inputs are $\mathbf{E}, \mathbf{J}, \mathbf{s}$, as well as an integer for which we can take the first power of two greater than or equal to σ .
- **COMPUTERESIDUALS** [18, Algorithm 5] (with an additional pre-processing detailed at the end of Section 4) computes the residual $\mathbf{P}^{(1)} \cdot \mathbf{E}$ from the first basis $\mathbf{P}^{(1)}$ obtained recursively.
- **KNOWNMINDEGMIB**, detailed in Section 4, computes the \mathbf{s} -Popov interpolation basis when one knows a *priori* the \mathbf{s} -minimal degree of (\mathbf{E}, \mathbf{J}) (see below).

In what follows, by *\mathbf{s} -minimal degree of (\mathbf{E}, \mathbf{J})* we mean the tuple of degrees of the diagonal entries of the \mathbf{s} -Popov interpolation basis \mathbf{P} for (\mathbf{E}, \mathbf{J}) . Because \mathbf{P} is in \mathbf{s} -Popov form, this is also the column degree of \mathbf{P} , and the sum of these degrees is $\deg(\det(\mathbf{P}))$. As a consequence, using Theorem 4.1 in [31] we obtain the following lemma, which implies in particular that the size of \mathbf{P} is at most $m(\sigma + 1)$.

LEMMA 2.1. Let $\mathbf{E} \in \mathbb{K}^{m \times \sigma}$, $\mathbf{J} \in \mathbb{K}^{\sigma \times \sigma}$, $\mathbf{s} \in \mathbb{Z}^m$, and let $(\delta_1, \dots, \delta_m)$ be the \mathbf{s} -minimal degree of (\mathbf{E}, \mathbf{J}) . Then, we have $\delta_1 + \dots + \delta_m \leq \sigma$.

ALGORITHM 1. POPOVMIB

Input:

- a matrix $\mathbf{E} \in \mathbb{K}^{m \times \sigma}$,
- a Jordan matrix $\mathbf{J} \in \mathbb{K}^{\sigma \times \sigma}$ in standard representation,
- a shift $\mathbf{s} \in \mathbb{Z}^m$.

Output:

- the \mathbf{s} -Popov interpolation basis \mathbf{P} for (\mathbf{E}, \mathbf{J}) ,
- the \mathbf{s} -minimal degree $\boldsymbol{\delta} = (\delta_1, \dots, \delta_m)$ of (\mathbf{E}, \mathbf{J}) .

1. If $\sigma \leq m$, return LINEARIZATIONMIB($\mathbf{E}, \mathbf{J}, \mathbf{s}, 2^{\lceil \log_2(\sigma) \rceil}$)
2. Else
 - a. $\mathbf{E}^{(1)} \leftarrow$ first $\lceil \sigma/2 \rceil$ columns of \mathbf{E}
 - b. $(\mathbf{P}^{(1)}, \boldsymbol{\delta}^{(1)}) \leftarrow$ POPOVMIB($\mathbf{E}^{(1)}, \mathbf{J}^{(1)}, \mathbf{s}$)
 - c. $\mathbf{E}^{(2)} \leftarrow$ last $\lfloor \sigma/2 \rfloor$ columns of $\mathbf{P}^{(1)} \cdot \mathbf{E} =$ COMPUTERESIDUALS($\mathbf{J}, \mathbf{P}^{(1)}, \mathbf{E}$)
 - d. $(\mathbf{P}^{(2)}, \boldsymbol{\delta}^{(2)}) \leftarrow$ POPOVMIB($\mathbf{E}^{(2)}, \mathbf{J}^{(2)}, \mathbf{s} + \boldsymbol{\delta}^{(1)}$)
 - e. $\mathbf{P} \leftarrow$ KNOWNMINDEGMIB($\mathbf{E}, \mathbf{J}, \mathbf{s}, \boldsymbol{\delta}^{(1)} + \boldsymbol{\delta}^{(2)}$)
 - f. Return $(\mathbf{P}, \boldsymbol{\delta}^{(1)} + \boldsymbol{\delta}^{(2)})$

Taking for granted the results in the next sections, we now prove our main theorem.

PROOF OF THEOREM 1.3. For the case $\sigma \leq m$, the correctness and the cost bound of Algorithm 1 both follow from [18, Theorem 1.4]: it uses $\mathcal{O}(\sigma^{\omega-1}m + \sigma^\omega \log(\sigma))$ operations (with an extra $\log(\sigma)$ factor if $\omega = 2$).

Now, we consider the case $\sigma > m$. Using the notation in the algorithm, assume that $\mathbf{P}^{(1)}$ is the \mathbf{s} -Popov interpolation basis for $(\mathbf{E}^{(1)}, \mathbf{J}^{(1)})$, and $\mathbf{P}^{(2)}$ is the \mathbf{t} -Popov interpolation basis for $(\mathbf{E}^{(2)}, \mathbf{J}^{(2)})$, where $\mathbf{t} = \mathbf{s} + \boldsymbol{\delta}^{(1)} = \text{rdeg}_{\mathbf{s}}(\mathbf{P}^{(1)})$, and $\boldsymbol{\delta}^{(1)}$ and $\boldsymbol{\delta}^{(2)}$ are the \mathbf{s} - and $\mathbf{s} + \boldsymbol{\delta}^{(1)}$ -minimal degrees of $(\mathbf{E}^{(1)}, \mathbf{J}^{(1)})$ and $(\mathbf{E}^{(2)}, \mathbf{J}^{(2)})$, respectively.

We claim that $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$ is \mathbf{s} -reduced: this will be proved in Lemma 3.2. Let us then prove that $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$ is an interpolation basis for (\mathbf{E}, \mathbf{J}) . Let $\mathbf{p} \in \mathbb{K}[X]^{1 \times m}$ be an interpolant for (\mathbf{E}, \mathbf{J}) . Since \mathbf{J} is upper triangular, \mathbf{p} is in particular an interpolant for $(\mathbf{E}^{(1)}, \mathbf{J}^{(1)})$, so there exists $\mathbf{v} \in \mathbb{K}[X]^{1 \times m}$ such that $\mathbf{p} = \mathbf{v}\mathbf{P}^{(1)}$. Besides, we have $\mathbf{P}^{(1)} \cdot \mathbf{E} = [0 | \mathbf{E}^{(2)}]$, so that $0 = \mathbf{p} \cdot \mathbf{E} = \mathbf{v}\mathbf{P}^{(1)} \cdot \mathbf{E} = [0 | \mathbf{v} \cdot \mathbf{E}^{(2)}]$, and thus $\mathbf{v} \cdot \mathbf{E}^{(2)} = 0$. Then, there exists $\mathbf{w} \in \mathbb{K}[X]^{1 \times m}$ such that $\mathbf{v} = \mathbf{w}\mathbf{P}^{(2)}$, which gives $\mathbf{p} = \mathbf{w}\mathbf{P}^{(2)}\mathbf{P}^{(1)}$.

In particular, the \mathbf{s} -Popov interpolation basis for (\mathbf{E}, \mathbf{J}) is the \mathbf{s} -Popov form of $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$. Thus, Lemma 3.2 combined with Lemma 3.3 will show that the \mathbf{s} -minimal degree of (\mathbf{E}, \mathbf{J}) is $\boldsymbol{\delta}^{(1)} + \boldsymbol{\delta}^{(2)}$. As a result, Proposition 4.3 states that Step 2.e correctly computes the \mathbf{s} -Popov interpolation basis for (\mathbf{E}, \mathbf{J}) .

Concerning the cost bound, the recursion stops when $\sigma \leq m$, and thus the algorithm uses $\mathcal{O}(m^\omega \log(m))$ operations (with an extra $\log(m)$ factor if $\omega = 2$). The depth of the

recursion is $\mathcal{O}(\log(\sigma/m))$; we have two recursive calls in dimensions $m \times \sigma/2$, and two calls to subroutines with cost bounds given in Corollary 4.5 and Proposition 4.3, respectively. The conclusion follows from the super-linearity properties of $\mathbf{M}(\cdot)$. \square

3. OBTAINING THE MINIMAL DEGREE FROM RECURSIVE CALLS

In this section, we show that the \mathbf{s} -minimal degree of (\mathbf{E}, \mathbf{J}) can be deduced for free from two bases computed recursively as in Algorithm 1. To do this, we actually prove a slightly more general result about the degrees of the \mathbf{s} -pivot entries of so-called weak Popov matrix forms [24].

DEFINITION 3.1 (WEAK POPOV FORM, PIVOT DEGREE). Let $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$ be nonsingular and let $\mathbf{s} \in \mathbb{Z}^m$. Then, \mathbf{P} is said to be in \mathbf{s} -weak Popov form if the \mathbf{s} -pivot indices of its rows are pairwise distinct; \mathbf{P} is said to be in \mathbf{s} -diagonal weak Popov form if its \mathbf{s} -pivot entries are on its diagonal.

If \mathbf{P} is in \mathbf{s} -weak Popov form, the \mathbf{s} -pivot degree of \mathbf{P} is the tuple $(\delta_1, \dots, \delta_m)$ where for $j \in \{1, \dots, m\}$, δ_j is the \mathbf{s} -pivot degree of the row of \mathbf{P} which has \mathbf{s} -pivot index j .

We recall from Section 1 that for $\mathbf{P} \in \mathbb{K}[X]^{k \times m}$, its \mathbf{s} -leading matrix $\text{lm}_{\mathbf{s}}(\mathbf{P}) \in \mathbb{K}^{k \times m}$ is formed by the coefficients of degree 0 of $\mathbf{X}^{-\mathbf{d}}\mathbf{P}\mathbf{X}^{\mathbf{s}}$, where $\mathbf{d} = \text{rdeg}_{\mathbf{s}}(\mathbf{P})$ and $\mathbf{X}^{\mathbf{s}}$ stands for the diagonal matrix with entries X^{s_1}, \dots, X^{s_m} . Then, a nonsingular $\mathbf{P} \in \mathbb{K}[X]^{m \times m}$ is in \mathbf{s} -diagonal weak Popov form with \mathbf{s} -pivot degree $\boldsymbol{\delta}$ if and only if $\text{lm}_{\mathbf{s}}(\mathbf{P})$ is lower triangular and invertible and $\text{rdeg}_{\mathbf{s}}(\mathbf{P}) = \mathbf{s} + \boldsymbol{\delta}$.

For example, at all stages of the algorithms in [31, 1, 18] for Problem 1 (as well as [14] if avoiding row permutations at the base case of the recursion), the computed bases are in shifted diagonal weak Popov form. This is due to the compatibility of this form with matrix multiplication, as stated in the next lemma.

LEMMA 3.2. Let $\mathbf{s} \in \mathbb{Z}^m$, $\mathbf{P}^{(1)} \in \mathbb{K}[X]^{m \times m}$ in \mathbf{s} -diagonal weak Popov form with \mathbf{s} -pivot degree $\boldsymbol{\delta}^{(1)}$, $\mathbf{t} = \mathbf{s} + \boldsymbol{\delta}^{(1)} = \text{rdeg}_{\mathbf{s}}(\mathbf{P}^{(1)})$, and $\mathbf{P}^{(2)} \in \mathbb{K}[X]^{m \times m}$ in \mathbf{t} -diagonal weak Popov form with \mathbf{t} -pivot degree $\boldsymbol{\delta}^{(2)}$. Then, $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$ is in \mathbf{s} -diagonal weak Popov form with \mathbf{s} -pivot degree $\boldsymbol{\delta}^{(1)} + \boldsymbol{\delta}^{(2)}$.

PROOF. By the predictable-degree property [19, Theorem 6.3-13] we have $\text{rdeg}_{\mathbf{s}}(\mathbf{P}^{(2)}\mathbf{P}^{(1)}) = \text{rdeg}_{\mathbf{t}}(\mathbf{P}^{(2)}) = \mathbf{t} + \boldsymbol{\delta}^{(2)} = \mathbf{s} + \boldsymbol{\delta}^{(1)} + \boldsymbol{\delta}^{(2)}$. The result follows since $\text{lm}_{\mathbf{s}}(\mathbf{P}^{(2)}\mathbf{P}^{(1)}) = \text{lm}_{\mathbf{t}}(\mathbf{P}^{(2)})\text{lm}_{\mathbf{s}}(\mathbf{P}^{(1)})$ is lower triangular and invertible. \square

For matrices in \mathbf{s} -Popov form, the \mathbf{s} -pivot degree coincides with the column degree: in particular, the \mathbf{s} -minimal degree of (\mathbf{E}, \mathbf{J}) is the \mathbf{s} -pivot degree of the \mathbf{s} -Popov interpolation basis for (\mathbf{E}, \mathbf{J}) . With the notation of Algorithm 1, the previous lemma proves that the \mathbf{s} -pivot degree of $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$ is $\boldsymbol{\delta}^{(1)} + \boldsymbol{\delta}^{(2)}$. In the rest of this section, we prove that the \mathbf{s} -Popov form of $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$ has the same \mathbf{s} -pivot degree as $\mathbf{P}^{(2)}\mathbf{P}^{(1)}$. Consequently, the \mathbf{s} -minimal degree of (\mathbf{E}, \mathbf{J}) is $\boldsymbol{\delta}^{(1)} + \boldsymbol{\delta}^{(2)}$ and thus can be found from $\mathbf{P}^{(2)}$ and $\mathbf{P}^{(1)}$ without computing their product.

It is known that left-unimodularly equivalent \mathbf{s} -reduced matrices have the same \mathbf{s} -row degree up to permutation [19, Lemma 6.3-14]. Here, we prove that the \mathbf{s} -pivot degree is invariant among left-unimodularly equivalent matrices in \mathbf{s} -weak Popov form.

Thus every row of \mathbf{R} has \mathbf{d} -degree at least 0, and the predictable degree property [19, Theorem 6.3.13] shows that \mathbf{U} is a constant matrix, and therefore unimodular. Then, $\bar{\mathbf{P}}$ is an interpolation basis for $(\mathcal{E} \cdot \mathbf{E}, \mathbf{J})$, and since it is in \mathbf{d} -Popov form, by Lemma 4.1 we obtain that $\bar{\mathbf{P}} = \text{lm}_{\mathbf{d}}(\mathbf{R})^{-1} \mathbf{R}$. The conclusion follows. \square

Then, it remains to prove that such a basis \mathbf{R} can be computed efficiently using the algorithm MINIMALINTERPOLATIONBASIS in [18]; this leads to Algorithm 2.

ALGORITHM 2. KNOWNMINDEGMIB
Input:

- a matrix $\mathbf{E} \in \mathbb{K}^{m \times \sigma}$ with $\sigma \geq m > 0$,
- a Jordan matrix $\mathbf{J} \in \mathbb{K}^{\sigma \times \sigma}$ in standard representation,
- a shift $\mathbf{s} \in \mathbb{Z}^m$,
- $\boldsymbol{\delta} = (\delta_1, \dots, \delta_m) \in \mathbb{N}^m$ the \mathbf{s} -minimal degree of (\mathbf{E}, \mathbf{J}) .

Output: the \mathbf{s} -Popov interpolation basis \mathbf{P} for (\mathbf{E}, \mathbf{J}) .

1. $\delta \leftarrow \lceil \sigma/m \rceil$, $\alpha_i \leftarrow \lfloor \delta_i/\delta \rfloor + 1$ for $1 \leq i \leq m$,
 $\bar{m} \leftarrow \alpha_1 + \dots + \alpha_m$
2. Let $\bar{\boldsymbol{\delta}} \in \mathbb{N}^{\bar{m}}$ as in (2) and $\mathbf{d} \leftarrow -\bar{\boldsymbol{\delta}} \in \mathbb{N}^{\bar{m}}$
3. Let $\mathcal{E} \in \mathbb{K}[X]^{\bar{m} \times m}$ as in (3) and $\bar{\mathbf{E}} \leftarrow \mathcal{E} \cdot \mathbf{E}$
4. $\mathbf{R} \leftarrow \text{MINIMALINTERPOLATIONBASIS}(\bar{\mathbf{E}}, \mathbf{J}, \mathbf{d} + (\delta, \dots, \delta))$
5. $\bar{\mathbf{P}} \leftarrow \text{lm}_{\mathbf{d}}(\mathbf{R})^{-1} \mathbf{R}$
6. Return the submatrix of $\bar{\mathbf{P}} \mathcal{E}$ formed by the rows at indices $\alpha_1 + \dots + \alpha_i$ for $1 \leq i \leq m$

PROPOSITION 4.3. Assuming that $\mathbf{J} \in \mathbb{K}^{\sigma \times \sigma}$ is a Jordan matrix given by a standard representation, and assuming we have the \mathbf{s} -minimal degree of (\mathbf{E}, \mathbf{J}) as an additional input, there is a deterministic algorithm KNOWNMINDEGMIB which solves Problem 1 using

$$\begin{aligned} \mathcal{O}(m^{\omega-1} \mathbf{M}(\sigma) \log(\sigma) \log(\sigma/m)) & \quad \text{if } \omega > 2, \\ \mathcal{O}(m \mathbf{M}(\sigma) \log(\sigma) \log(\sigma/m) \log(m)^3) & \quad \text{if } \omega = 2 \end{aligned}$$

operations in \mathbb{K} .

PROOF. We focus on the case $\sigma \geq m$; otherwise, a better cost bound can be achieved even without knowing $\boldsymbol{\delta}$ [18, Theorem 1.4]. The correctness of Algorithm 2 follows from Lemma 4.2. We remark that it uses $\mathbf{d} + (\delta, \dots, \delta)$ rather than \mathbf{d} because the minimal interpolation basis algorithm in [18] requires the input shift to have non-negative entries. Since adding a constant to every entry of \mathbf{d} does not change the notion of \mathbf{d} -reducedness, the basis \mathbf{R} obtained at Step 4 is a \mathbf{d} -minimal interpolation basis for (\mathbf{E}, \mathbf{J}) .

Concerning the cost bound, we will show that it is dominated by the time spent in Step 4. First, we prove that $|\mathbf{d} - \min(\mathbf{d})| \in \mathcal{O}(\sigma)$, so that the cost of Step 4 follows from [18, Theorem 1.5]. We have $\alpha_i = 1 + \lfloor \delta_i/\delta \rfloor \leq 1 + m\delta_i/\sigma$ for all i . Thus, $\bar{m} = \alpha_1 + \dots + \alpha_m \leq m + \sum_{1 \leq i \leq m} m\delta_i/\sigma \leq 2m$ thanks to Lemma 2.1. Then, since all entries of \mathbf{d} are in $\{-\delta, \dots, 0\}$, we obtain $|\mathbf{d} - \min(\mathbf{d})| \leq \bar{m}\delta \leq 2m(1 + \sigma/m) \leq 4\sigma$.

Step 3 can be done in $\mathcal{O}(m \mathbf{M}(\sigma) \log(\sigma))$ operations according to Lemma 4.4 below.

Lemma 4.1 proves that the sum of the column degrees of \mathbf{R} is $|\bar{\boldsymbol{\delta}}| = |\boldsymbol{\delta}| \leq \sigma$. Then, the product in Step 5 can be done in $\mathcal{O}(m^{\omega-1} \sigma)$ operations, by first linearizing the columns of \mathbf{R} into a $\bar{m} \times \bar{m} + |\bar{\boldsymbol{\delta}}|$ matrix over \mathbb{K} , then left-multiplying this matrix by $\text{lm}_{\mathbf{d}}(\mathbf{R})^{-1}$ (itself computed using $\mathcal{O}(m^{\omega})$ operations), and finally performing the inverse linearization.

Because of the degrees in $\bar{\mathbf{P}}$ and the definition of \mathcal{E} , the output in Step 6 can be formed without using any arithmetic operation. \square

The efficient computation of $\mathcal{E} \cdot \mathbf{E}$ can be done with the algorithm for computing residuals in [18, Section 6].

LEMMA 4.4. The product $\mathcal{E} \cdot \mathbf{E}$ at Step 3 of Algorithm 2 can be computed using $\mathcal{O}(m \mathbf{M}(\sigma) \log(\sigma))$ operations in \mathbb{K} .

PROOF. The product $\mathcal{E} \cdot \mathbf{E}$ has \bar{m} rows, with $\bar{m} \leq 2m$ as above. Besides, by definition of \mathcal{E} , each row of $\mathcal{E} \cdot \mathbf{E}$ is a product of the form $X^{i\delta} \cdot \mathbf{E}_{j,*}$, where $0 \leq i \leq m$, $1 \leq j \leq m$, and $\mathbf{E}_{j,*}$ denotes the row j of \mathbf{E} . In particular, $i\delta \leq 2\sigma$: then, according to [18, Proposition 6.1], each of these \bar{m} products can be performed using $\mathcal{O}(\mathbf{M}(\sigma) \log(\sigma))$ operations in \mathbb{K} . \square

This lemma and the partial linearization technique can also be used to compute the residual at Step 2.c of Algorithm 1, that is, a product of the form $\bar{\mathbf{P}} \cdot \mathbf{E}$ with the sum of the column degrees of $\bar{\mathbf{P}}$ bounded by σ . First, we expand the high-degree columns of $\bar{\mathbf{P}}$ to obtain $\bar{\mathbf{P}} \in \mathbb{K}[X]^{m \times \bar{m}}$ of degree less than $\lceil \sigma/m \rceil$ such that $\bar{\mathbf{P}} = \bar{\mathbf{P}} \mathcal{E}$; then, we compute $\bar{\mathbf{E}} = \mathcal{E} \cdot \mathbf{E}$; and finally we rely on the algorithm in [18, Proposition 6.1] to compute $\bar{\mathbf{P}} \cdot \bar{\mathbf{E}} = \bar{\mathbf{P}} \cdot \mathbf{E}$ efficiently.

COROLLARY 4.5. Let $\mathbf{E} \in \mathbb{K}^{m \times \sigma}$ with $\sigma \geq m$, and let $\mathbf{J} \in \mathbb{K}^{\sigma \times \sigma}$ be a Jordan matrix given by a standard representation. Let $\bar{\mathbf{P}} \in \mathbb{K}[X]^{m \times \bar{m}}$ with column degree $(\delta_1, \dots, \delta_m)$ such that $\delta_1 + \dots + \delta_m \leq \sigma$. Then, the product $\bar{\mathbf{P}} \cdot \mathbf{E}$ can be computed using $\mathcal{O}(m^{\omega-1} \mathbf{M}(\sigma) \log(\sigma))$ operations in \mathbb{K} .

Acknowledgments. We thank B. Beckermann and G. Labahn for their valuable comments, as well as an anonymous referee for suggesting a shorter proof of Lemma 3.3. C.-P. Jeannerod and G. Villard were partly supported by the ANR project HPAC (ANR 11 BS02 013). V. Neiger was supported by the international mobility grants Expl'ra Doc from Région Rhône-Alpes, PALSE, and Mitacs Globalink - Inria. É. Schost was supported by NSERC.

5. REFERENCES

- [1] B. Beckermann and G. Labahn. A uniform approach for the fast computation of matrix-type Padé approximants. *SIAM J. Matrix Anal. Appl.*, 15(3):804–823, 1994.
- [2] B. Beckermann and G. Labahn. Fraction-free computation of matrix rational interpolants and matrix gcds. *SIAM J. Matrix Anal. Appl.*, 22(1):114–144, 2000.
- [3] B. Beckermann, G. Labahn, and G. Villard. Normal forms for general polynomial matrices. *J. Symbolic Comput.*, 41(6):708–737, 2006.
- [4] P. Beelen and K. Brander. Key equations for list decoding of Reed-Solomon codes and how to solve them. *J. Symbolic Comput.*, 45(7):773–786, 2010.
- [5] A. Bostan, C.-P. Jeannerod, and É. Schost. Solving structured linear systems with large displacement rank. *Theor. Comput. Sci.*, 407(1-3):155–181, 2008.
- [6] K. Brander. *Interpolation and List Decoding of Algebraic Codes*. PhD thesis, Technical University of Denmark, 2010.

