



Solving structured linear systems with large displacement rank

Alin Bostan^a, Claude-Pierre Jeannerod^b, Éric Schost^{c,*}

^a Algorithms Project, INRIA Paris-Rocquencourt, 78153 Le Chesnay Cedex, France

^b Arénair Project, INRIA Rhône-Alpes, Laboratoire LIP (CNRS, ENSL, INRIA, UCBL), ENS Lyon, France

^c Computer Science Department, The University of Western Ontario, London, Ontario, Canada

ARTICLE INFO

Article history:

Received 23 October 2007

Received in revised form 14 May 2008

Accepted 16 May 2008

Communicated by V. Pan

Keywords:

Structured linear algebra

Dense linear algebra

ABSTRACT

Linear systems with structures such as Toeplitz, Vandermonde or Cauchy-likeness can be solved in $O(\alpha^2 n)$ operations, where n is the matrix size, α is its displacement rank, and O denotes the omission of logarithmic factors. We show that for such matrices, this cost can be reduced to $O(\alpha^{\omega-1} n)$, where ω is a feasible exponent for matrix multiplication over the base field. The best known estimate for ω is $\omega < 2.38$, resulting in costs of order $O(\alpha^{1.38} n)$. We present consequences for Hermite–Padé approximation and bivariate interpolation.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Structured linear algebra techniques are a versatile set of tools. They enable one to deal at once with matrices with features such as Toeplitz, Vandermonde or Cauchy-likeness, which arise in various problems, from interpolation to reconstruction of rational or algebraic functions, etc.

Following [26], the usual way of measuring to what extent a matrix possesses one such structure is through its *displacement rank*, that is, the rank of its image through a suitable *displacement operator*. For P and Q in respectively $\mathbb{K}^{n \times n}$ and $\mathbb{K}^{m \times m}$, where $\mathbb{K}$ is our base field, we will use the displacement operator

$$\Delta[P, Q] : \mathbb{K}^{n \times m} \rightarrow \mathbb{K}^{n \times m} \\ A \mapsto A - PAQ.$$

Two matrices (Y, Z) in $\mathbb{K}^{n \times \alpha} \times \mathbb{K}^{m \times \alpha}$ will be called a P, Q -generator of length α for A if $\Delta[P, Q](A) = YZ^t$. The main idea behind algorithms for structured matrices is to use such generators as a compact data structure, in cases when $\Delta[P, Q](A)$ has low rank. Even though these definitions hold for rectangular A , in most of this article, we have $n = m$.

Usual choices for P or Q are either diagonal matrices or cyclic down-shift matrices of size n , defined for φ in $\mathbb{K}$ by

$$\mathbb{Z}_{n, \varphi} = \begin{bmatrix} 0 & & & \varphi \\ 1 & 0 & & \\ & \ddots & \ddots & \\ & & 1 & 0 \end{bmatrix} \in \mathbb{K}^{n \times n}.$$

The Toeplitz structure corresponds to $P = \mathbb{Z}_{n, 0}$ and $Q = \mathbb{Z}_{m, 0}^t$, so that $\Delta[\mathbb{Z}_{n, 0}, \mathbb{Z}_{m, 0}^t](A)$ equals

$A - (A \text{ shifted down and right by one unit}).$

* Corresponding author. Tel.: +1 519 661 2111.

E-mail addresses: Alin.Bostan@inria.fr (A. Bostan), claud-pierre.jeannerod@ens-lyon.fr (C.-P. Jeannerod), eschost@uwo.ca, Eric.Schost@polytechnique.fr (É. Schost).

The Vandermonde structure is obtained by taking P diagonal and Q a cyclic right-shift matrix $\mathbb{Z}_{m,\varphi}^t$. For the Cauchy structure, both P and Q are diagonal. In particular, in all these cases, storing P and Q requires no more than $O(n)$ elements.

In this paper, we consider the following task:

LinearSystem(P, Q, α): Given a P, Q -generator of length α for a matrix $A \in \mathbb{K}^{n \times n}$, with $\alpha \leq n$, and given $v \in \mathbb{K}^n$, find a solution to the equation $Au = v$, or determine that none exists.

This problem makes sense only when the operator $\Delta[P, Q]$ is invertible: this will be the case in our three cases of focus, Toeplitz-like, Vandermonde-like and Cauchy-like matrices (however, non-invertible operators can be dealt with if some extra information about A is known, such as suitable columns, see [46] and [45, Section 4.5]). Previous work then yielded the following kind of results: for these three structures, one can solve the problem LinearSystem using $\tilde{O}(\alpha^2 n)$ operations in $\mathbb{K}$, where the $\tilde{O}$ notation hides logarithmic factors.

When α is constant, such estimates are optimal up to logarithmic factors. However, there are several situations where α is not bounded *a priori* (see examples below). In the extreme case of very loosely structured matrices, when α goes up to $\alpha \simeq n$, the cost above becomes $\tilde{O}(n^3)$.

On the other side of the spectrum, we find *dense* linear algebra methods. Let $\omega \leq 3$ be such that $n \times n$ matrices over $\mathbb{K}$ can be multiplied in $O(n^\omega)$ operations (the current record estimate is $\omega < 2.38$ [13]). As in many other references, we will assume that $\omega > 2$ (in a potential situation with $\omega = 2$, logarithmic factors would appear in some estimates). Linear systems of size n can then be solved in time $O(n^\omega)$, using for example LSP factorization [25]; with $\omega < 3$, this is better than the above $\tilde{O}(n^3)$ estimate.

Our contribution bridges a gap between the approaches of structured and dense linear algebra by providing algorithms of cost $\tilde{O}(\alpha^{\omega-1} n)$ in the case of Toeplitz-like, Vandermonde-like and Cauchy-like matrices.

Model of computation. Our underlying computational model is the *algebraic RAM* over the field $\mathbb{K}$; a complete definition is given by Kaltofen in [27]. Concretely, the cost estimates count two kinds of operations:

- Operations in $\mathbb{K}$ (sums, products, equality tests, inversions) have unit cost; the generation of a random element in $\mathbb{K}$ has unit cost as well.
- We use integer arithmetic for handling indices in arrays. With a view towards simplicity, we see all these operations at unit cost.

Regarding the last point, in all our algorithms, the indices we will work with will be polynomial in n (when dealing with matrices of size n); hence, if desired, the translation to actual binary cost estimates would only induce a polylogarithmic overhead in n .

Our algorithms rely on *polynomial multiplication*; we will thus denote by $M : \mathbb{N}_{>0} \rightarrow \mathbb{R}_{>0}$ a function such that polynomials in $\mathbb{K}[x]$ of degree less than d can be multiplied using at most $M(d)$ operations. Using Fast Fourier Transform algorithms, $M(d)$ can be taken in $O(d \log(d) \log \log(d))$, see Section 2 for more details.

Our algorithms are *probabilistic*. For simplicity, we say that an algorithm has type $P(s, d)$ if it chooses $r \leq s$ elements in $\mathbb{K}$ and, if these elements are chosen uniformly at random in a finite subset S of $\mathbb{K}$, the probability of success is at least $1 - d/|S|$. In particular, $\mathbb{K}$ should have cardinality more than d to guarantee that this probability can be made positive.

Main results. The first result covers matrices having Toeplitz-like structure, with $P = \mathbb{Z}_{n,0}$ and $Q = \mathbb{Z}_{n,0}^t$. We obtain a complexity in $\tilde{O}(\alpha^{\omega-1} n) \subset \tilde{O}(\alpha^{1.38} n)$, to be compared with an optimal cost of $O(\alpha n)$. For α constant, our result is quasi-linear in n ; when $\alpha \simeq n$, we recover the $O(n^\omega)$ behaviour of dense methods, up to logarithmic factors.

Theorem 1. *The problem LinearSystem($\mathbb{Z}_{n,0}, \mathbb{Z}_{n,0}^t, \alpha$) can be solved in time $O(\alpha^{\omega-1} M(n) \log^2(n))$, by a probabilistic algorithm of type $P(3n - 2, n^2 + n)$.*

A fundamental application is the solution of approximation problems: given a *master polynomial* M and polynomials $f_1, \dots, f_s$, one seeks a combination of the f_i , with polynomial coefficients of prescribed degrees, which vanishes modulo M . This includes in particular Padé and Hermite–Padé approximation (taking $M = x^n$), with applications to e.g. recovering the minimal polynomial of an algebraic power series f (taking $f_i = f^{i-1}$).

Corollary 1. *Let $M \in \mathbb{K}[x]$ be of degree n , let $f_1, \dots, f_s \in \mathbb{K}[x]$ be of degrees less than n , and let $v_1, \dots, v_s \in \mathbb{N}$ be such that $\sum_{i \leq s} v_i = n + 1$. One can find $g_1, \dots, g_s \in \mathbb{K}[x]$, not all zero, of respective degrees less than $v_1, \dots, v_s$, such that*

$$g_1 f_1 + \dots + g_s f_s = 0 \text{ mod } M,$$

in time $O(s^{\omega-1} M(n) \log^2(n))$. The algorithm is probabilistic of type $P(3n - 2, n^2 + n)$.

Our second result addresses the Vandermonde-like case, where $P = \mathbb{D}(x)$ is the diagonal matrix defined by a vector $x = [x_1, \dots, x_n]^t \in \mathbb{K}^n$ and Q is $\mathbb{Z}_{n,\psi}^t$. We assume that x has the form

$$x = \begin{bmatrix} x_1 \\ \vdots \\ x_s \end{bmatrix}, \quad \text{with } x_i = \begin{bmatrix} x_{i,1} \\ \vdots \\ x_{i,v_i} \end{bmatrix} \in \mathbb{K}^{v_i}, \quad (1)$$

and with the following conditions: for all j , all entries of $\mathbf{x}_j$ are pairwise distinct; for $j < s$, $\mathbf{x}_{j+1}$ is a *prefix* of $\mathbf{x}_j$, in the sense that $0 < v_{j+1} \leq v_j$ and $x_{j+1,k} = x_{j,k}$ for $k = 1, \dots, v_{j+1}$. We will call s the *multiplicity* of $\mathbf{x}$ (see Section 2). Note also that, given $\mathbf{x}$ as in Eq. (1), we can deduce $v_1, \dots, v_s$ in time $O(n)$.

Assuming that Eq. (1) is satisfied is a mild assumption: any vector $\mathbf{x}$ can be put into this form after permuting its entries. However, finding the permutation has a cost: this is particularly easy if there is an order on $\mathbb{K}$, using sorting algorithms; without this assumption, though, the question seems harder: we mention an algorithm of complexity $O(M(n) \log^3(n))$ in [Appendix](#). Hence, we rather stick to the above assumption for simplicity.

We make a second assumption, which is necessary to ensure that our problem is well posed:


$$\psi x_i^n \neq 1 \quad \text{for } i \leq n. \quad (2)$$

Theorem 2. *Given $\mathbf{x}$ and ψ as in Eqs. (1) and (2), one can solve the problem $\text{LinearSystem}(\mathbb{D}(\mathbf{x}), \mathbb{Z}_{n,\psi}^t, \alpha)$ in time $O(\alpha^{\omega-1} M(n) \log^2(n))$ by a probabilistic algorithm of type $P(4n - 2, 4n^2 + n)$.*

An application of the latter theorem is polynomial interpolation. The approach applies to any number of variables, but we discuss only the bivariate case for simplicity. Consider n pairwise distinct interpolation points in $\mathbb{K}^2$; without loss of generality, we assume that they are given through the following lists, where the points with the same abscissas are grouped together:

$$\begin{aligned} P_1 &= [p_{1,1} = (x_1, y_{1,1}), \dots, p_{1,v_1} = (x_1, y_{1,v_1})], \\ &\vdots \\ P_s &= [p_{s,1} = (x_s, y_{s,1}), \dots, p_{s,v_s} = (x_s, y_{s,v_s})], \end{aligned} \quad (3)$$

with $v_1 \geq \dots \geq v_s > 0$ and $n = v_1 + \dots + v_s$. The following figure illustrates the case $s = 3$, $n = 7$ and $(v_1, v_2, v_3) = (3, 2, 2)$.



In general, it is difficult to state *a priori* that a multivariate interpolation problem is well defined. Here, however, given values $[v_{i,j}] \in \mathbb{K}^n$, with $1 \leq i \leq s$ and $1 \leq j \leq v_i$, Theorem 1 in [35] (see also [19]) implies that there exists a unique $F \in \mathbb{K}[x, y]$ of the form

$$F = \sum_{0 \leq i < s, 0 \leq j < v_{i+1}} f_{i,j} x^i y^j$$

such that $F(p_{i,j}) = v_{i,j}$ for all i, j (hence the monomial support of the polynomial F corresponds to the index set of the sample points). In our previous example, F thus has the form

$$F = f_{0,0} + f_{0,1}y + f_{0,2}y^2 + f_{1,0}x + f_{1,1}xy + f_{2,0}x^2 + f_{2,1}x^2y.$$

Finding the coefficients $f_{i,j}$ of F is a linear problem with Vandermonde-like structure; in Section 5.4 we will prove the following corollary of [Theorem 2](#).

Corollary 2. *Given the lists of points $P_1, \dots, P_s$ and the lists of values $[F(p) : p \in P_1], \dots, [F(p) : p \in P_s]$, the coefficients $f_{i,j}$ of F can be computed in time $O(v_1^{\omega-1} M(n) \log^2(n))$. The algorithm is probabilistic of type $P(4n - 2, 4n^2 + n)$.*

Suppose for instance that $v_1 = s$, $v_2 = s - 1, \dots, v_s = 1$, so that we are interpolating on the simplex of monomials of degree less than s ; here, $n = s(s+1)/2$. Then, our algorithm has subquadratic complexity $\tilde{O}(n^{(\omega+1)/2}) \subset \tilde{O}(n^{1.69})$.

Our third result deals with the Cauchy-like case, where $\mathbf{P} = \mathbb{D}(\mathbf{x})$ and $\mathbf{Q} = \mathbb{D}(\mathbf{y})$ are diagonal matrices defined by some vectors $\mathbf{x}$ and $\mathbf{y}$ in $\mathbb{K}^n$. We assume as before that $\mathbf{x}$ satisfies (1) and, in addition, make a similar assumption on $\mathbf{y}$:

$$\mathbf{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_t \end{bmatrix}, \quad \text{with } \mathbf{y}_i = \begin{bmatrix} y_{i,1} \\ \vdots \\ y_{i,\delta_i} \end{bmatrix} \in \mathbb{K}^{\delta_i}, \quad (4)$$

and with the following conditions: for all k , all entries of $\mathbf{y}_k$ are pairwise distinct; for $k < t$, $\mathbf{y}_{k+1}$ is a *prefix* of $\mathbf{y}_k$. Again, given $\mathbf{x}$ as in Eq. (1) and $\mathbf{y}$ as in Eq. (4), one can deduce all the v_j and δ_k in time $O(n)$.

For our Cauchy-like problem to be well posed, we now replace the assumption in Eq. (2) with the following one:

$$x_i y_j \neq 1 \quad \text{for } 1 \leq i, j \leq n. \quad (5)$$

The complexity result we obtain in this case is then essentially the same as for the Toeplitz-like and Vandermonde-like cases:

Theorem 3. Given x and y as in Eqs. (1), (4) and (5), one can solve the problem $\text{LinearSystem}(\mathbb{D}(x), \mathbb{D}(y), \alpha)$ in time $O(\alpha^{\omega-1} M(n) \log^2(n))$, by a probabilistic algorithm of type $P(5n - 2, 7n^2 + n)$.

Previous work on structured systems. The notions of displacement rank and displacement operator originate from the work of Kailath, Kung and Morf [26]. Since then, the literature has vastly developed; see [45] for a list of references.

The basis for Theorem 1 is the seminal algorithm of Bitmead and Anderson [7] and Morf [37,38], which nontrivially adapts Strassen's divide-and-conquer approach [53] to structured computations. This algorithm requires several invertibility conditions to hold. Kaltofen [28,29] extended this idea to arbitrary matrices (see also [6, p. 204] for some related ideas), obtaining a cost of $O(\alpha^2 M(n) \log(n))$; for small α , this is better than our result in Theorem 1 by a factor of $\log(n)$.

We follow these previous approaches; our main technical contribution is an algorithm for the fast multiplication of a Toeplitz-like matrix (given by a generator of length α) by α vectors: we replace the separate computation of these products (which has cost quadratic in α) by simultaneous computations, where using polynomial matrix multiplication yields a subquadratic cost in α .

To prove Theorem 2, we transform a Vandermonde-like system into a Toeplitz-like one, following Pan's idea [41,42]. We use a transformation from [24], generalizing it by taking into account the possibility of repetitions in the diagonal component of the operator. Similarly, the proof of Theorem 3 uses a reduction from Cauchy-like systems to Vandermonde-like ones.

In both the Vandermonde and Cauchy cases, a direct approach (not relying on the reduction to Toeplitz systems) is possible, see [48,44,47,12] and [45, Chapter 5]. As in the Toeplitz case, the bottleneck of these algorithms is the multiplication of a structured matrix (given by a generator of length α) by α vectors. However, following this approach, we are not able to obtain algorithms whose costs would match in all cases those of Theorems 2 and 3.

Applications. An important example of Toeplitz-like system solving is the approximation problem of Corollary 1. In the particular case of Hermite–Padé approximation, with $M = x^n$, a central reference is Beckermann–Labahn's algorithm [2], that has complexity $O(s^\omega M(n) \log(n))$ for computing a σ -basis of order n of the input system (and thus a solution to the approximation problem); see [23]. In generic cases, an unpublished result of Lecerf reduces the cost to $O(s^{\omega-1} M(n) \log(n))$ and Storjohann [52] subsequently obtained a deterministic algorithm of similar complexity, applying in all cases. However, to our knowledge, these results do not extend to an arbitrary choice of M . Following notably [1,54], Beckermann and Labahn study that general case in [3] under the angle of fraction-free algorithms, with however a complexity more than linear in n .

Another example of a Toeplitz-like system that occurs frequently is when the matrix is block-Toeplitz, a block-size equal to α giving a displacement rank in $O(\alpha)$. Although Theorem 1 applies to any such system, a deterministic cost of $O(\alpha^{\omega-1} M(n) \log(n))$ can be obtained in the particular case where the matrix is invertible. As described for example in [16, 17], this cost follows from combining an inversion formula of [33] with σ -basis computations as in [23].

Multivariate polynomial interpolation has been studied extensively (see [20] for a survey and [4,10,56] for algorithms relevant to sparse techniques). However, to our knowledge, previous references either do not cover the problems we deal with, or have higher complexity (typically, quadratic). Regarding the converse evaluation problem, let us mention the subquadratic complexity result of [40], which however deals with more general situations than ours.

Practical issues. Few practical algorithms are currently known for matrix multiplication with complexity better than cubic (see [53,34,32] for an exponent 2.776). However, even when using algorithms of cubic complexity, the re-introduction of dense matrix arithmetic in our algorithms means that we can rely on extremely optimized implementations of matrix multiplication, such as the ones relying on BLAS libraries for finite field arithmetic [15].

Hence, besides theoretical estimates, the re-introduction of dense matrix multiplication in algorithms for structured matrices may lead to practical improvements. However, extra work may be required: useful refinements would for instance consist in removing the undesired logarithmic factors that appear in Theorems 1–3.

Organization of the paper. After introducing basic notation and results in Section 2, we present in Section 3 the bases of our technical improvement, which can be stated in terms of polynomial operations only. These results are then applied, first to the Toeplitz case in Section 4, then to the Vandermonde case in Section 5 and the Cauchy case in Section 6. The Appendix presents an algorithm for putting an arbitrary vector into the form of Eq. (1), without relying on sorting algorithms.

2. Notation and preliminaries

We gather here all needed notation for vectors and matrices as well as some basic identities and complexity results on structured and dense matrices.

Matrices, vectors and polynomials. In what follows, we consider matrices and vectors over a field $\mathbb{K}$; vectors in $\mathbb{K}^n$ are identified with column-matrices in $\mathbb{K}^{n \times 1}$.

Matrices (resp. vectors) are written in upper-case (resp. lower-case) sans-serif font. If A (resp. $B, C, \dots$) is a matrix, a_i (resp. $b_i, c_i, \dots$) is its i th column. If x (resp. $y, z, \dots$) is a vector, its i th entry is written x_i (resp. $y_i, z_i, \dots$). Special matrices (diagonal, Vandermonde, Toeplitz, ...) will be written with blackboard bold letters ($\mathbb{D}, \mathbb{V}, \dots$).

For n a positive integer and $i \leq n$, we let $e_{n,i}$ be the i th unit vector in $\mathbb{K}^n$, so that $e_{n,i}$ is zero, except at the i th entry, which is 1. The identity matrix in $\mathbb{K}^{n \times n}$, whose i th column equals $e_{n,i}$, will be written $\mathbb{I}_n$.

Regarding univariate polynomials, we use the following notation:

- The degree of $F \in \mathbb{K}[x]$ is written $\deg(F)$.
- For $n \in \mathbb{N}$, we let $\mathbb{K}[x]_n$ be the n -dimensional vector space of polynomials of degree less than n .
- For $n \in \mathbb{N}$ and $F = f_0 + \dots + f_{n-1}x^{n-1} \in \mathbb{K}[x]_n$, we will write $\text{Rev}_n(F)$ for the reversal $f_{n-1} + \dots + f_0x^{n-1}$ of F .
- For $F \in \mathbb{K}[x]$ and $G \in \mathbb{K}[x]$ non-zero, $F \text{ div } G$ and $F \text{ mod } G$ are the quotient and the remainder in the division of F by G .
Notation such as $AB \text{ mod } C$ or $AB \text{ div } C$ must be interpreted as $(AB) \text{ mod } C$ or $(AB) \text{ div } C$, respectively.

We say that a vector is *repetition-free* when its entries are pairwise distinct. Besides, we use the following notation for vectors and polynomials derived from a vector $\mathbf{a} = [a_1, \dots, a_n]^t$ in $\mathbb{K}^n$:

- $\text{Flip}(\mathbf{a})$ is the vector $[0, a_n, \dots, a_2]^t \in \mathbb{K}^n$.
- $\text{Pol}(\mathbf{a})$ is the polynomial $\sum_{i=0}^{n-1} a_{i+1}x^i \in \mathbb{K}[x]_n$.
- If F is a function on $\mathbb{K}$, then $F(\mathbf{a})$ is the vector $[F(a_1), \dots, F(a_n)]^t \in \mathbb{K}^n$.
- Conversely, if $\mathbf{a}$ is repetition-free and $\mathbf{y}$ is in $\mathbb{K}^n$, then $\text{Interp}(\mathbf{a}, \mathbf{y})$ is the unique polynomial F in $\mathbb{K}[x]_n$ such that $y_i = F(a_i)$ for $1 \leq i \leq n$.

Families of structured matrices. Several families of structured matrices will be used along this paper. First, we associate matrices to a vector $\mathbf{x} = [x_1, \dots, x_n]^t$ in $\mathbb{K}^n$:

- $\mathbb{D}(\mathbf{x}) \in \mathbb{K}^{n \times n}$ is the diagonal matrix whose i th diagonal entry equals x_i .
- $\mathbb{L}(\mathbf{x})$ is the lower-triangular Toeplitz matrix with first column $\mathbf{x}$:

$$\mathbb{L}(\mathbf{x}) = \begin{bmatrix} x_1 & & & \\ x_2 & x_1 & & \\ \vdots & \vdots & \ddots & \\ x_n & x_{n-1} & \dots & x_1 \end{bmatrix} \in \mathbb{K}^{n \times n}.$$

- $\mathbb{U}(\mathbf{x}) = \mathbb{L}(\mathbf{x})^t$ is the upper-triangular Toeplitz matrix with first row $\mathbf{x}^t$:

$$\mathbb{U}(\mathbf{x}) = \begin{bmatrix} x_1 & x_2 & \dots & x_n \\ & x_1 & \dots & x_{n-1} \\ & & \ddots & \vdots \\ & & & x_1 \end{bmatrix} \in \mathbb{K}^{n \times n}.$$

- For φ in $\mathbb{K}$, we denote by $\mathbb{T}(\mathbf{x}, \varphi)$ the φ -circulant matrix with first column $\mathbf{x}$; that is, $\mathbb{T}(\mathbf{x}, \varphi) = \mathbb{L}(\mathbf{x}) + \varphi \mathbb{U}(\text{Flip}(\mathbf{x}))$:

$$\mathbb{T}(\mathbf{x}, \varphi) = \begin{bmatrix} x_1 & \varphi x_n & \dots & \varphi x_2 \\ x_2 & x_1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \varphi x_n \\ x_n & x_{n-1} & \dots & x_1 \end{bmatrix} \in \mathbb{K}^{n \times n}.$$

- For $m \in \mathbb{N}$, the $n \times m$ Vandermonde matrix $\mathbb{V}(\mathbf{x}, m)$ is given by:

$$\mathbb{V}(\mathbf{x}, m) = \begin{bmatrix} 1 & x_1 & \dots & x_1^{m-1} \\ \vdots & \vdots & & \vdots \\ 1 & x_n & \dots & x_n^{m-1} \end{bmatrix} \in \mathbb{K}^{n \times m}.$$

- Given $\mathbf{y} = [y_1, \dots, y_m]^t$ in $\mathbb{K}^m$ such that $x_i y_j \neq 1$ for all i, j , the $n \times m$ Cauchy matrix $\mathbb{C}(\mathbf{x}, \mathbf{y})$ is defined by

$$\mathbb{C}(\mathbf{x}, \mathbf{y}) = \begin{bmatrix} \frac{1}{1-x_1 y_1} & \dots & \frac{1}{1-x_1 y_m} \\ \vdots & & \vdots \\ \frac{1}{1-x_n y_1} & \dots & \frac{1}{1-x_n y_m} \end{bmatrix} \in \mathbb{K}^{n \times m}.$$

We also associate families of matrices to univariate polynomials:

- Given $F = f_0 + \cdots + f_{n-1}x^{n-1} \in \mathbb{K}[x]_n$ and $m \in \mathbb{N}$, we denote by $\mathbb{M}(F, m)$ the matrix of the map $\mathbb{K}[x]_m \rightarrow \mathbb{K}[x]_{n+m-1}$ of multiplication by F :

$$\mathbb{M}(F, m) = \begin{bmatrix} f_0 & & & \\ \vdots & \ddots & & \\ \vdots & & \ddots & f_0 \\ f_{n-1} & & \ddots & \vdots \\ & & \ddots & \vdots \\ & & & f_{n-1} \end{bmatrix} \in \mathbb{K}^{(n+m-1) \times m}.$$

- Given the monic polynomial $F = f_0 + \cdots + f_{n-1}x^{n-1} + x^n$, we denote by $\mathbb{X}(F)$ the matrix of multiplication by x in $\mathbb{K}[x]/\langle F \rangle$:

$$\mathbb{X}(F) = \begin{bmatrix} & & -f_0 \\ & & -f_1 \\ & \ddots & \vdots \\ & & 1 & -f_{n-1} \end{bmatrix} \in \mathbb{K}^{n \times n}.$$

Structured matrix identities. The families of structured matrices shown above satisfy many identities. We present now the ones we will need below; we start by showing how to rewrite products (diagonal matrix) $\times$ (Vandermonde matrix) using polynomial multiplication matrices.

Lemma 1. Let $\mathbf{x}$ be repetition-free in $\mathbb{K}^n$, let $\mathbf{y}$ be in $\mathbb{K}^n$, and let $P \in \mathbb{K}[x]_n$ be defined by $P = \text{Interp}(\mathbf{x}, \mathbf{y})$. Then, for m in $\mathbb{N}$, we have

$$\mathbb{D}(\mathbf{y}) \mathbb{V}(\mathbf{x}, m) = \mathbb{V}(\mathbf{x}, n+m-1) \mathbb{M}(P, m).$$

Proof. Let $\mathbf{f}$ be in $\mathbb{K}^m$ and let $F = \text{Pol}(\mathbf{f}) \in \mathbb{K}[x]_m$. Then, both vectors $\mathbb{D}(\mathbf{y}) \mathbb{V}(\mathbf{x}, m) \mathbf{f}$ and $\mathbb{V}(\mathbf{x}, n+m-1) \mathbb{M}(P, m) \mathbf{f}$ have entries $y_i F(x_i)$. $\square$

The above identity is used in Sections 5.1 and 6.1 to solve multiplication problems associated to, respectively, Vandermonde and Cauchy-like matrices. In fact, for Cauchy-like matrices, we will use it in combination with another identity, shown in Lemma 2 below: it factors $n \times m$ Cauchy matrices by means of Vandermonde and transposed Vandermonde matrices; when $n = m$, this is [18, Proposition 3.2] (see also [45, Exercise 3.10(b)]).

Lemma 2. Let $\mathbf{x}$ in $\mathbb{K}^n$ and $\mathbf{y}$ in $\mathbb{K}^m$ be such that $x_i y_j \neq 1$ for all i, j . Let $F \in \mathbb{K}[x]_{m+1}$ be given by $F = \prod_{j \leq m} (1 - xy_j)$ and let $\mathbf{f}$ in $\mathbb{K}^m$ be such that $\text{Pol}(\mathbf{f}) = F \bmod x^m$. Then

$$\mathbb{C}(\mathbf{x}, \mathbf{y}) = \mathbb{D}(F(\mathbf{x}))^{-1} \mathbb{V}(\mathbf{x}, m) \mathbb{L}(\mathbf{f}) \mathbb{V}(\mathbf{y}, m)^t.$$

Proof. Since $F(x_i) \neq 0$ for all i , we have $(1 - x_i y_j)^{-1} = F(x_i)^{-1} G_j(x_i)$, with $G_j = F/(1 - xy_j)$. Therefore, $\mathbb{C}(\mathbf{x}, \mathbf{y}) = \mathbb{D}(F(\mathbf{x}))^{-1} [G_1(\mathbf{x}) \cdots G_m(\mathbf{x})]$. Now, $G_j \in \mathbb{K}[x]_m$ and thus

$$G_j = (F \bmod x^m)(1 + y_j x + \cdots + y_j^{m-1} x^{m-1}) \bmod x^m.$$

Writing $\mathbf{g}_j$ for the vector in $\mathbb{K}^m$ such that $\text{Pol}(\mathbf{g}_j) = G_j$, we obtain

$$G_j(\mathbf{x}) = \mathbb{V}(\mathbf{x}, m) \mathbf{g}_j = \mathbb{V}(\mathbf{x}, m) \mathbb{L}(\mathbf{f}) [1 \ y_j \ \cdots \ y_j^{m-1}]^t$$

for $1 \leq j \leq m$, which concludes the proof. $\square$

Our final expression, already seen e.g. in [24], shows that some companion matrices and diagonal matrices are similar. We will use it in Section 5.2 for reducing Vandermonde-like systems to Toeplitz-like ones, and then again in Section 6.2 for reducing Cauchy-like systems to Vandermonde-like ones.

Lemma 3. Let $\mathbf{x}$ be a repetition-free vector in $\mathbb{K}^n$ and let F be the monic polynomial defined by $F = \prod_{i \leq n} (x - x_i)$. Then we have the equality

$$\mathbb{D}(\mathbf{x}) = \mathbb{V}(\mathbf{x}, n) \mathbb{X}(F) \mathbb{V}(\mathbf{x}, n)^{-1}.$$

Proof. For $\mathbf{y}$ in $\mathbb{K}^n$, we claim that the vectors $\mathbb{D}(\mathbf{x}) \mathbf{y}$ and $\mathbb{V}(\mathbf{x}, n) \mathbb{X}(F) \mathbb{V}(\mathbf{x}, n)^{-1} \mathbf{y}$ have entries $x_i y_i$. This is readily seen for the former one. Defining $P = \text{Interp}(\mathbf{x}, \mathbf{y}) \in \mathbb{K}[x]_n$, computing the latter vector amounts to evaluate the polynomial $(xP \bmod F)$ at $\mathbf{x}$, yielding the values $x_i y_i$ as well. $\square$

Polynomial expressions. Multiplication by the various families of structured matrices seen before can be reinterpreted in terms of polynomial operations. This is foremost the case for multiplication by lower- or upper-triangular Toeplitz matrices, as seen in the following well-known lemma.

Lemma 4. Let y and z be in $\mathbb{K}^n$ and let $u = \mathbb{L}(y)z$ and $v = \mathbb{U}(y)z$. We have

$$\begin{aligned}\text{Pol}(u) &= \text{Pol}(y) \text{Pol}(z) \bmod x^n, \\ \text{Pol}(v) &= \text{Rev}_n(\text{Pol}(y)) \text{Pol}(z) \text{div } x^{n-1}.\end{aligned}$$

We gather here some further expressions of a similar spirit, that are needed later. First, we show how to rewrite transposed polynomial multiplication in terms of plain multiplication; the following result is from [9, Section 4.1].

Lemma 5. Let P be in $\mathbb{K}[x]_m$, f be in $\mathbb{K}^{n+m-1}$ and let $u = \mathbb{M}(P, n)^t f \in \mathbb{K}^n$. Then we have

$$\text{Pol}(u) = (\text{Rev}_m(P) \text{Pol}(f) \text{div } x^{m-1}) \bmod x^n.$$

The next lemma describes a more complex operation, needed in Section 5.1 for handling Vandermonde-like matrices.

Lemma 6. Let z be in $\mathbb{K}^n$, φ be in $\mathbb{K}$, P be in $\mathbb{K}[x]_m$, f be in $\mathbb{K}^{n+m-1}$, and define $g \in \mathbb{K}^n$ by

$$g = \mathbb{T}(z, \varphi) \mathbb{M}(P, n)^t f.$$

Define further

- $z' = \text{Flip}(z)$ in $\mathbb{K}^n$,
- $Z = \text{Pol}(z)$ and $Z' = \text{Pol}(z')$ in $\mathbb{K}[x]_n$,
- $F = \text{Pol}(f)$ in $\mathbb{K}[x]_{n+m-1}$.

Then we have the equality

$$\text{Pol}(g) = Z \left(\text{Rev}_m(P) F \text{div } x^{m-1} \right) \bmod x^n + \varphi \text{Rev}_n \left(Z' (P \text{Rev}_{n+m-1}(F) \text{div } x^{m-1}) \bmod x^n \right).$$

Proof. Since by definition $\mathbb{T}(z, \varphi)$ equals $\mathbb{L}(z) + \varphi \mathbb{U}(z')$, we see that $\text{Pol}(g)$ equals $\text{Pol}(g') + \varphi \text{Pol}(g'')$, where $g' = \mathbb{L}(z) \mathbb{M}(P, n)^t f$ and $g'' = \mathbb{U}(z') \mathbb{M}(P, n)^t f$. Thus, by applying successively Lemma 4 and Lemma 5 to g' , we first obtain

$$\text{Pol}(g') = Z \left(\text{Rev}_m(P) F \text{div } x^{m-1} \right) \bmod x^n.$$

To deal with g'' , notice that $\mathbb{U}(z') = \mathbb{J} \mathbb{L}(z') \mathbb{J}$, where the *reversal* matrix $\mathbb{J}$ is zero, except on the anti-diagonal, whose entries are 1's. Then, doing as before,

$$\text{Pol}(g'') = \text{Rev}_n \left(Z' \text{Rev}_n \left((\text{Rev}_m(P) F \text{div } x^{m-1}) \bmod x^n \right) \bmod x^n \right);$$

observe the presence of the two extra Rev_n operations, due to the conjugation by the $\mathbb{J}$ matrix. To get rid of the second reversal, one checks that

$$\text{Rev}_n \left((\text{Rev}_m(P) F \text{div } x^{m-1}) \bmod x^n \right) = (P \text{Rev}_{n+m-1}(F) \text{div } x^{m-1}) \bmod x^n,$$

which gives the required result. $\square$

Algorithms for univariate polynomials. As mentioned before, M denotes a function from $\mathbb{N}_{>0}$ to $\mathbb{R}_{>0}$ such that over any ring, polynomials of degree less than d can be multiplied using at most $M(d)$ ring operations.

Following [21, Chapter 8], we make the assumption that the function $d \mapsto M(d)/d$ is non-decreasing; this implies in particular that the super-linearity condition $M(d) + M(d') \leq M(d + d')$ holds for all d, d' . Using the results of [49,11], one can take $M(d) \in O(d \log(d) \log \log(d))$.

Several fast algorithms are based on fast multiplication. Let thus $x = [x_1, \dots, x_n]^t$ be in $\mathbb{K}^n$; we will use the following well-known results.

Construction from roots. The polynomials $\prod_{i \leq n} (x - x_i)$ and $\prod_{i \leq n} (1 - x_i x)$ can be computed in $O(M(n) \log(n))$ operations [21, Lemma 10.4].

Evaluation. Given $f \in \mathbb{K}^n$, one can compute $\mathbb{V}(x, n) f$ (equivalently, evaluate any polynomial $F \in \mathbb{K}[x]_n$ at x) in $O(M(n) \log(n))$ operations [21, Corollary 10.8].

Transposed evaluation. Given $f \in \mathbb{K}^n$, one can compute the vector $\mathbb{V}(x, n)^t f$ in $O(M(n) \log(n))$ operations [22, Theorem 10.4].

Interpolation. If x is repetition-free and f is in $\mathbb{K}^n$, one can compute the vector $\mathbb{V}(x, n)^{-1} f$ (equivalently, interpolate any polynomial $F \in \mathbb{K}[x]_n$ at x) in $O(M(n) \log(n))$ operations [21, Corollary 10.12].

Transposed interpolation. If x is repetition-free and f is in $\mathbb{K}^n$, one can compute the vector $\mathbb{V}(x, n)^{-t} f$ in $O(M(n) \log(n))$ operations [30, Section 5] (a similar result can be found in [41, Section 10]).

Algorithms for dense matrices. As said earlier, we let $2 < \omega \leq 3$ be such that $n \times n$ matrices over $\mathbb{K}$ can be multiplied in $O(n^\omega)$ operations in $\mathbb{K}$.

Dense matrix operations will be used for several purposes. First, since we require that our fast polynomial multiplication algorithms apply over any ring, we deduce that polynomial matrices over $\mathbb{K}$ having degree less than d and size n can be multiplied in $O(n^\omega M(d))$ operations. In Section 3, we will actually use the following two results, which deal more precisely with some specific rectangular polynomial matrix products.

Lemma 7. Let A, B, C be matrices of respective sizes $(n \times p)$, $(p \times n)$ and $(n \times p)$, with entries in $\mathbb{K}[x]_d$. Then one can compute the product ABC using $O(p^{\omega-1} n M(d))$ operations.

Proof. Suppose first that $p \leq n$. Up to bordering the matrices by less than p zero rows or columns, we can suppose that p divides n . Let then $\ell = n/p$. We can rewrite the product ABC as a product of matrices of size $(\ell \times 1) \times (1 \times \ell) \times (\ell \times 1)$, having blocks of size $p \times p$ as entries. We then compute this product as $A(BC)$; the cost is $O(p^\omega \ell M(d))$, which is $O(p^{\omega-1} n M(d))$.

Suppose now that $n \leq p$; as above, we suppose that n divides p and let $\ell = p/n$. Then, the product ABC is rewritten as a product of matrices of size $(1 \times \ell) \times (\ell \times 1) \times (1 \times \ell)$, having blocks of size $n \times n$ as entries. We then compute this product as $(AB)C$; the cost is $O(n^\omega \ell M(d))$, that is, in $O(n^{\omega-1} p M(d))$. This is also in $O(p^{\omega-1} n M(d))$, since $0 \leq n \leq p$ and $\omega \geq 2$ imply $n^{\omega-1} p \leq p^{\omega-1} n$. $\square$

Lemma 8. Let A and B be matrices of respective sizes $(n \times p)$ and $(p \times m)$, with entries in $\mathbb{K}[x]_d$. If $n \leq p$, then one can compute the product AB using $O(p^{\omega-1} \max\{n, m\} M(d))$ operations.

Proof. Suppose first that $n \leq m$. As in the previous lemma, up to bordering A by less than n zero columns and B by less than n zero columns and rows, we can suppose that n divides p and m . Let $k = p/n$ and $\ell = m/n$. Then, the product AB is rewritten as a product of matrices of size $(1 \times k) \times (k \times \ell)$, having blocks of size $n \times n$ as entries. The cost for computing this product is thus $O(n^\omega k \ell M(d))$ operations, which is $O(n^{\omega-2} p m M(d))$. Since $0 \leq n \leq p$ and $\omega - 2 \geq 0$, this is in $O(p^{\omega-1} m M(d))$, as required.

Assuming $m \leq n$, we obtain a similar estimate of $O(m^{\omega-2} p n M(d))$. Since then $0 \leq m \leq p$ and $\omega - 2 \geq 0$, this is in $O(p^{\omega-1} n M(d))$, as required. $\square$

Besides fast (polynomial) matrix multiplication in the above two particular cases, we shall also use fast elimination on dense matrices. We thus give the following result on isolating a basis of the row-span of a matrix. This is a simple consequence of Proposition 2.15 in [51].

Lemma 9. Let A be in $\mathbb{K}^{n \times p}$. One can compute in time $O(r^{\omega-2} np)$ a quadruple (r, J, G, P) such that r is the rank of A , J is a subset of $\{1, \dots, n\}$ of length r , G is a matrix in $\mathbb{K}^{(n-r) \times r}$, P is a permutation matrix of order n , and

$$EA = \begin{bmatrix} A' \\ 0 \end{bmatrix}, \quad \text{where } E = \begin{bmatrix} \mathbb{I}_r & 0 \\ G & \mathbb{I}_{n-r} \end{bmatrix} P,$$

and where $A' \in \mathbb{K}^{r \times p}$ consists of the rows of A indexed by J .

Proof. Using Algorithm 2.14 in [51], one can compute in time $O(r^{\omega-2} np)$ the rank r of A along with a permutation matrix P , and a matrix U of the form

$$U = \begin{bmatrix} F & 0 \\ G & \mathbb{I}_{n-r} \end{bmatrix}, \quad \text{with } F \in \mathbb{K}^{r \times r},$$

such that the last $n - r$ rows of the product UPA are zero. Then let $J = \{i_1, \dots, i_r\}$, where i_j is the index of the non-zero entry of the j th row of P . Define E as above, replacing F by the identity $\mathbb{I}_r$ in U ; then, the first r rows of PA , and thus of EA , are the rows of A indexed by J . $\square$

The above lemma will be used in Sections 5.3 and 6.3 for introducing zeros in some generators of Vandermonde and Cauchy-like matrices, and thus handling the high multiplicities that may arise in those cases. Another application of this lemma is the fast computation of generators of minimal length (or, equivalently, of maximal rank), an operation already needed for the Toeplitz case. The proof of the result below, which is Remark 4.6.7 in [45], shows this.

Lemma 10. Let $P, Q \in \mathbb{K}^{n \times n}$. Given a P, Q -generator of length $\alpha \leq n$ for $A \in \mathbb{K}^{n \times n}$, one can compute a P, Q -generator for A of minimal length in time $O(\alpha^{\omega-1} n)$.

Proof. Denoting by (Y, Z) the given generator, we have $A - PAQ = YZ^t$. Let V be an invertible matrix in $\mathbb{K}^{\alpha \times \alpha}$ such that $YV = [Y' \ 0]$, with $Y' \in \mathbb{K}^{n \times r}$ and $r = \text{rank}(Y)$. Let also $Z' \in \mathbb{K}^{n \times r}$ consist of the first r columns of ZV^{-t} . We have $YZ^t = Y'Z'^t$ and thus (Y', Z') is a P, Q -generator for A of length r .

Although Y' has full rank, Z' may have rank less than r . Therefore, let further V' be an invertible matrix in $\mathbb{K}^{r \times r}$ such that $Z'V' = [H \ 0]$, with $H \in \mathbb{K}^{n \times r'}$ and $r' = \text{rank}(Z')$. Defining G as the matrix in $\mathbb{K}^{n \times r'}$ that consists of the first r' columns of $Y'V'^{-t}$, we obtain the equality $Y'Z'^t = GH^t$. Since now both G and H have full rank, they form a P, Q -generator for A of minimal length.

By Lemma 9, the cost of computing r and V is in $O(\alpha^{\omega-1} n)$. One then obtains Y' in time $O(\alpha^{\omega-1} n)$ by multiplication with V . To get Z' , one can first compute V^{-t} in time $O(\alpha^\omega)$ and then deduce ZV^{-t} in time $O(\alpha^{\omega-1} n)$. Similarly, deducing (G, H) from (Y', Z') can be done in time $O(\alpha^{\omega-1} n)$, as required. $\square$

Partitions and multiplicities. Our algorithms for Vandermonde and Cauchy-like systems may require to perform rearrangements of the diagonal matrices appearing in the corresponding displacement operators. We detail this process here, pointing out in particular that performing the required permutations is inexpensive.

We start by integer partitions. A partition ν of $n \in \mathbb{N}_{>0}$ is a sequence of positive integers $\nu_1 \geq \dots \geq \nu_s$ such that $\nu_1 + \dots + \nu_s = n$. It will be useful to associate to ν the sequence $\nu_0^* = 0, \nu_1^* = \nu_1, \dots, \nu_i^* = \nu_1 + \dots + \nu_i, \dots$

The *conjugate* of ν is another partition $\mu = \mu_1 \geq \dots \geq \mu_t$ of n , where μ_j is the number of elements i in $\{1, \dots, s\}$ such that $\nu_i \geq j$; in particular, the conjugate of μ is ν and we have $t = \nu_1$ and $s = \mu_1$. Remark that given ν , one deduces μ in n integer additions: after initializing μ at $0, \dots, 0$ (ν_1 repetitions), we obtain μ by incrementing μ_j , for $i = 1, \dots, s$ and $j = 1, \dots, \nu_i$.

Next, we discuss partitions of vectors. Let thus x be in $\mathbb{K}^n$. We suppose that x has the form of Eq. (1):

$$x = \begin{bmatrix} x_1 \\ \vdots \\ x_s \end{bmatrix}, \quad \text{with } x_i = \begin{bmatrix} x_{i,1} \\ \vdots \\ x_{i,\nu_i} \end{bmatrix} \in \mathbb{K}^{\nu_i},$$

where each x_j is a repetition-free vector of size ν_j and where for $j < s$, x_{j+1} is a prefix of x_j . Since $\nu = \nu_1 \geq \dots \geq \nu_s$ is a partition of n , we let $\mu = \mu_1 \geq \dots \geq \mu_t$ be the conjugate partition. The number of occurrences of $x_{1,j}$ in x is thus μ_j ; we call it its *multiplicity*, and call $\mu_1 = s$ the multiplicity of x .

It follows that, after permuting its entries, x can be rewritten as

$$x' = \begin{bmatrix} x_{\sigma(1)} \\ \vdots \\ x_{\sigma(n)} \end{bmatrix} = \begin{bmatrix} x'_1 \\ \vdots \\ x'_t \end{bmatrix}, \quad (6)$$

where x'_j is a vector consisting of μ_j repetitions of $x_{1,j}$. For example, if $x = [1, 2, 4, 1, 2, 1, 2]^t$, we have

$$x_1 = \begin{bmatrix} 1 \\ 2 \\ 4 \end{bmatrix}, x_2 = \begin{bmatrix} 1 \\ 2 \end{bmatrix}, x_3 = \begin{bmatrix} 1 \\ 2 \end{bmatrix} \quad \text{and} \quad x'_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, x'_2 = \begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix}, x'_3 = [4],$$

with $s = t = 3$, $(\nu_1, \nu_2, \nu_3) = (3, 2, 2)$, $(\mu_1, \mu_2, \mu_3) = (3, 3, 1)$, $\sigma = [1, 4, 6, 2, 5, 7, 3]$ and $\sigma^{-1} = [1, 4, 7, 2, 5, 3, 6]$.

Lemma 11. Given x as in Eq. (1), one can deduce a permutation σ as in Eq. (6) in time $O(n)$; the converse operation can be done within the same complexity.

Proof. The partition ν is obtained by scanning x for occurrences of $x_{1,1}$, in time $O(n)$. One obtains μ as the conjugate of ν for a similar cost; the permutation σ can then be made explicit as

$$\sigma(1), \dots, \sigma(n) = 1, 1 + \nu_1^*, \dots, 1 + \nu_{\mu_1-1}^*, \dots, \nu_1, \nu_1 + \nu_1^*, \dots, \nu_1 + \nu_{\mu_t-1}^*,$$

the first block giving the indices of $x_{1,1} = x_{2,1} = \dots = x_{s,1}$ in x and so on. Conversely, given x' , μ is obtained by scanning x' for indices where two consecutive entries differ. Then, one recovers ν and σ^{-1} in the same way as before, changing the roles of ν and μ . In any case, the cost fits into the required bound. $\square$

By analogy with the square-free decomposition of polynomials, we call the representation in Eq. (6) a *repetition-free decomposition* of x .

3. Polynomial operations

We discuss now two problems involving polynomials. Both boil down to suitably using polynomial matrix multiplication to speed up the simultaneous computation of several bilinear or trilinear expressions; as it turns out, these questions are the main ingredients of the algorithms in the following sections.

3.1. First problem

In the following, some integers n and $\alpha \leq n$ are fixed. Let

$$(Y_i)_{i \leq \alpha}, \quad (Z_i)_{i \leq \alpha} \quad \text{and} \quad (F_j)_{j \leq \alpha}$$

be in $\mathbb{K}[x]_n$. The next proposition will be used in Section 4, as the key to the proof of Theorem 1 for Toeplitz-like systems. The cost reported in this proposition is actually the bottleneck of the main algorithm: any reduction on the following estimate would entail a reduction of the overall cost.

Proposition 1. *One can compute the polynomials*

$$G_j = \sum_{i=1}^{\alpha} Y_i (Z_i F_j \bmod x^n), \quad j = 1, \dots, \alpha$$

in time $O(\alpha^{\omega-1} M(n) \log(n))$.

The direct approach, used in previous work, consists in computing all polynomials G_j independently, for a cost of $O(\alpha^2 M(n))$ operations. Our contribution shows how to compute the polynomials G_j simultaneously using polynomial matrix multiplication.

Proof. Up to replacing n with $\bar{n} = 2^{\lceil \log(n) \rceil}$ and F_j with $x^{\bar{n}-n} F_j$, we can (and will) suppose that n is a power of 2.

We first show how to rewrite truncated products using non-truncated ones, using ideas reminiscent of *short products* [39]. Let $k \geq 1$ be a power of 2 and let ℓ be in $\mathbb{N}$. For $P = p_0 + p_1 x + \dots$, we define $P^{(\ell,k)} \in \mathbb{K}[x]$ as follows:

$$P^{(\ell,1)} = p_\ell \quad \text{and} \quad P^{(\ell,k)} = \sum_{i=\ell k}^{\ell k + k/2 - 1} p_i x^{i-\ell k} \quad \text{for } k \geq 2.$$

In all cases, $P^{(\ell,k)}$ is a polynomial of degree less than $k/2$. Using this subdivision enables us to rewrite a truncated product $PQ \bmod x^n$ as a sum of non-truncated ones.

Lemma 12. *For P and Q in $\mathbb{K}[x]$ and m a power of 2,*

$$PQ \bmod x^m = \sum_{k=1,2,4,\dots,m} x^{m-k} \sum_{\ell=0}^{m/k-1} P^{(\ell,k)} Q^{(m/k-1-\ell,k)},$$

where the sum is taken on all $k \leq m$ that are powers of 2.

Proof. We proceed by induction on $m \geq 1$, for m a power of 2. If $m = 1$ the result is clear, so, assuming that the property holds at index $m/2$, we prove it at index m . Let us write

$$P \bmod x^m = P_0 + x^{m/2} P_1 \quad \text{and} \quad Q \bmod x^m = Q_0 + x^{m/2} Q_1,$$

with P_0, P_1, Q_0, Q_1 of degree less than $m/2$. Then we have

$$P_0^{(\ell,k)} = P^{(\ell,k)} \quad \text{and} \quad P_1^{(\ell,k)} = P^{(\ell+m/2k,k)}$$

for any $k \geq 1$ and $\ell \geq 0$ such that $\ell k + k/2 \leq m/2$. Analogous equalities hold for Q, Q_0 and Q_1 . Now, by definition, the following equality holds:

$$PQ \bmod x^m = P_0 Q_0 + x^{m/2} (P_0 Q_1 + P_1 Q_0 \bmod x^{m/2}). \quad (7)$$

Observe first that $P_0 Q_0$ equals $P^{(0,m)} Q^{(0,m)}$, which corresponds to the term $k = m$ on the right-hand side of the formula we wish to establish. Next, the induction assumption shows that $P_0 Q_1 \bmod x^{m/2}$ is given by

$$\sum_{k=1,2,\dots,m/2} x^{m/2-k} \sum_{\ell=0}^{m/2k-1} P_0^{(\ell,k)} Q_1^{(m/2k-1-\ell,k)},$$

which we rewrite as

$$\sum_{k=1,2,\dots,m/2} x^{m/2-k} \sum_{\ell=0}^{m/2k-1} P^{(\ell,k)} Q^{(m/k-1-\ell,k)}.$$

Similarly, $P_1 Q_0 \bmod x^{m/2}$ equals

$$\sum_{k=1,2,\dots,m/2} x^{m/2-k} \sum_{\ell=m/2k}^{m/k-1} P^{(\ell,k)} Q^{(m/k-1-\ell,k)}.$$

Putting these equalities in Eq. (7) ends the proof. $\square$

We can now prove the proposition. By Lemma 12, we have for all i and j

$$Y_i (Z_i F_j \bmod x^n) = \sum_{k=1,2,4,\dots,n} x^{n-k} \sum_{\ell=0}^{n/k-1} Y_i Z_i^{(\ell,k)} F_j^{(n/k-1-\ell,k)}.$$

Thus for $j \leq \alpha$, we have $G_j = \sum_{k=1,2,4,\dots,n} x^{n-k} G_{j,k}$, with

$$G_{j,k} = \sum_{i=1}^{\alpha} \sum_{\ell=0}^{n/k-1} Y_i Z_i^{(\ell,k)} F_j^{(n/k-1-\ell,k)}.$$

We next show how to compute all polynomials $G_{1,k}, \dots, G_{\alpha,k}$, for a fixed k .

Lemma 13. Let $k \leq n$ be a power of 2. Then one can compute $G_{1,k}, \dots, G_{\alpha,k}$ in time $O(\alpha^{\omega-1} M(n))$.

Proof. Let $k' = n/k$, and let Z and F be the $(\alpha \times k')$ and $(k' \times \alpha)$ polynomial matrices

$$Z = \begin{bmatrix} Z_1^{(0,k)} & \dots & Z_1^{(k'-1,k)} \\ \vdots & & \vdots \\ Z_\alpha^{(0,k)} & \dots & Z_\alpha^{(k'-1,k)} \end{bmatrix}, \quad F = \begin{bmatrix} F_1^{(k'-1,k)} & \dots & F_\alpha^{(k'-1,k)} \\ \vdots & & \vdots \\ F_1^{(0,k)} & \dots & F_\alpha^{(0,k)} \end{bmatrix}.$$

Then we have the equality

$$[G_{1,k} \quad \dots \quad G_{\alpha,k}] = [Y_1 \quad \dots \quad Y_\alpha] Z F.$$

All entries of Z and F have degree less than $k/2$, whereas the polynomials Y_i have degree less than n . To balance the degrees, for $i \leq \alpha$, we write $Y_i = \sum_{\ell=0}^{k'-1} Y_{i,\ell} x^{k\ell}$, with $Y_{i,\ell}$ of degree less than k . We then define the $(k' \times \alpha)$ matrix

$$Y = \begin{bmatrix} Y_{1,0} & \dots & Y_{\alpha,0} \\ \vdots & & \vdots \\ Y_{1,k'-1} & \dots & Y_{\alpha,k'-1} \end{bmatrix}$$

with polynomial entries of degree less than k , such that

$$[Y_1 \quad \dots \quad Y_\alpha] = [1 \quad x^k \quad x^{2k} \quad \dots \quad x^{(k'-1)k}] Y. \quad (8)$$

Using Lemma 7, we can compute the product $Y Z F$ in $O(\alpha^{\omega-1} k' M(k))$ operations. Since $k' M(k) \leq M(n)$, this cost is in $O(\alpha^{\omega-1} M(n))$. Finally, by Eq. (8), we deduce $G_{1,k}, \dots, G_{\alpha,k}$ from the product $Y Z F$ in time $O(\alpha k' k)$, which is in $O(\alpha n)$. $\square$

To conclude the proof of Proposition 1, we apply Lemma 13 to $k = 1, 2, 4, \dots, n$, for a total cost of $O(\alpha^{\omega-1} M(n) \log(n))$. The cost of deducing $G_1, \dots, G_\alpha$ is a negligible $O(\alpha n \log(n))$. $\square$

3.2. Second problem

As above, some integers $n \in \mathbb{N}$ and $\alpha \leq n$ are fixed. Let also s and t be positive integers bounded by α and let $v_1, \dots, v_s$ and $\delta_1, \dots, \delta_t$ be positive integers such that

$$n = v_1 + \dots + v_s \quad \text{and} \quad n = \delta_1 + \dots + \delta_t.$$

Let then

$$(Q_{i,j})_{i \leq \alpha, j \leq s} \quad \text{and} \quad (R_{i,k})_{i \leq \alpha, k \leq t}$$

be in $\mathbb{K}[x]$, with $\deg(Q_{i,j}) < v_j$ and $\deg(R_{i,k}) < \delta_k$. The following result is used in Section 6 to reduce the solution of Cauchy-like systems to that of Vandermonde-like systems.

Proposition 2. One can compute the polynomials

$$P_{j,k} = \sum_{i=1}^{\alpha} Q_{i,j} R_{i,k}, \quad j = 1, \dots, s, \quad k = 1, \dots, t$$

in time $O(\alpha^{\omega-1} M(n) \log(n))$.

Remark that computing all polynomials $P_{j,k}$ independently leads to a cost in $O(\alpha^2 M(n))$, which is quadratic in α . As we did in the previous subsection, we are going to use polynomial matrix multiplication to compute the polynomials $P_{j,k}$ simultaneously and obtain a cost subquadratic in α .

Proof. Let $L = \{0, 1, \dots, \lfloor \log(n) \rfloor\}$. For $\beta, \gamma \in L$, define

$$S_\beta = \{j : 2^\beta \leq v_j < 2^{\beta+1}\} \quad \text{and} \quad T_\gamma = \{k : 2^\gamma \leq \delta_k < 2^{\gamma+1}\};$$

we then let $q_\beta = |S_\beta|$ and $r_\gamma = |T_\gamma|$. Rewrite the equality $n = \sum_{j=1}^s v_j$ as

$$n = \sum_{\beta \in L} \sum_{j \in S_\beta} v_j.$$

Since for $j \in S_\beta$, all v_j are at least 2^β , we deduce the following inequality, together with its analogue for the subsets T_γ :

$$\sum_{\beta \in L} 2^\beta q_\beta \leq n \quad \text{and} \quad \sum_{\gamma \in L} 2^\gamma r_\gamma \leq n. \quad (9)$$

Given the integers $v_1, \dots, v_s$ and $\delta_1, \dots, \delta_t$, all subsets S_β and T_γ can be obtained in time $O(n)$. Then, for a given pair (β, γ) , the polynomials $\{P_{j,k} : j \in S_\beta, k \in T_\gamma\}$ are the entries of the matrix product $Q R$, with

$$Q = \begin{bmatrix} \vdots & & \vdots \\ Q_{1,j} & \cdots & Q_{\alpha,j} \\ \vdots & & \vdots \end{bmatrix}_{j \in S_\beta}, \quad R = \begin{bmatrix} \cdots & R_{1,k} & \cdots \\ \vdots & & \\ \cdots & R_{\alpha,k} & \cdots \end{bmatrix}_{k \in T_\gamma}.$$

The polynomial matrix Q has dimensions $(q_\beta \times \alpha)$ and degree less than $2^{\beta+1}$; R has dimensions $(\alpha \times r_\gamma)$ and degree less than $2^{\gamma+1}$.

Let us first deal with the case where $\beta \leq \gamma$. Then, the polynomials in R have potentially larger degrees than the ones in Q . To balance these degrees, we write $R = \sum_{\ell=0}^{2^{\gamma-\beta}-1} R_\ell x^{2^{\beta+1}\ell}$, where R_ℓ is an $(\alpha \times r_\gamma)$ polynomial matrix with entries of degree less than $2^{\beta+1}$.

From the knowledge of the products $\{Q R_\ell : 0 \leq \ell < 2^{\gamma-\beta}\}$, one can deduce $Q R$ using $O(\alpha 2^\gamma r_\gamma)$ extra additions; summing over all β, γ with $\beta \leq \gamma$, and using Eq. (9), this amounts to $O(\alpha n \log(n))$ additions. Hence, we focus on computing the former products. Concatenating all matrices R_ℓ , let us define

$$R' = [R_0 \quad \cdots \quad R_{2^{\gamma-\beta}-1}];$$

it is then enough to compute $Q R'$, since the $Q R_\ell$ can be read off this product.

The matrix R' has dimensions $(\alpha \times r'_\gamma)$, with $r'_\gamma = 2^{\gamma-\beta} r_\gamma$, and entries of degree less than $2^{\beta+1}$. Since $q_\beta \leq \alpha$, Lemma 8 shows that the cost of computing the product $Q R'$ is in

$$O(\alpha^{\omega-1} \max\{q_\beta, r'_\gamma\} M(2^{\beta+1})).$$

Recall that by assumption on the function M , we have $M(2^{\beta+1}) \leq 2^{\beta+1} M(n)/n$. Bounding the max by a sum and substituting r'_γ by its value, we deduce the upper bound

$$O\left(\alpha^{\omega-1} 2^\beta q_\beta \frac{M(n)}{n} + \alpha^{\omega-1} 2^\gamma r_\gamma \frac{M(n)}{n}\right).$$

Summing over all β, γ with $\beta \leq \gamma$, and using Eq. (9), we obtain an upper bound in $O(\alpha^{\omega-1} M(n) \log(n))$. The case where $\beta > \gamma$ is handled similarly. $\square$

We will need in Section 5 the following corollary of the previous proposition, to reduce the solution of Vandermonde-like systems to that of Toeplitz-like systems. Let $s \leq \alpha$ and $v_1, \dots, v_s$ be as above and let

$$(Z_i)_{i \leq \alpha}, \quad (Y_{i,j})_{i \leq \alpha, j \leq s}, \quad \text{and} \quad (F_j)_{j \leq s}$$

be in $\mathbb{K}[x]$, with $\deg(Z_i) < n$, $\deg(Y_{i,j}) < v_j$ and $\deg(F_j) < n + v_j$.

Proposition 3. *One can compute the polynomials*

$$P_j = \sum_{i=1}^{\alpha} Z_i (Y_{i,j} F_j \operatorname{div} x^{v_j-1}), \quad j = 1, \dots, s$$

in time $O(\alpha^{\omega-1} M(n) \log(n))$.

The remarks we did before still apply: a direct approach consists in computing all polynomials P_j independently, for a cost in $O(\alpha^2 M(n))$. Using polynomial matrix multiplication to share computations, we get a cost subquadratic in α .

Proof. For $i \leq \alpha$ and $j \leq s$, let $G_{i,j} = Y_{i,j} F_j \operatorname{mod} x^{v_j-1}$. Then, for $j \leq s$, we can define

$$Q_j = \sum_{i=1}^{\alpha} Z_i G_{i,j} \quad \text{and} \quad R_j = \sum_{i=1}^{\alpha} Z_i Y_{i,j}.$$

It follows that P_j is given by $(R_j F_j - Q_j) \operatorname{div} x^{v_j-1}$.

For given i and j , the polynomial $G_{i,j}$ can be computed in $M(v_j)$ operations; hence, the whole time for computing all these polynomials is at most $\alpha M(n)$. Applying Proposition 2 with $t = 1$ and $\delta_1 = n$, we deduce that we can compute all polynomials Q_j and R_j using $O(\alpha^{\omega-1} M(n) \log(n))$ operations. Finally, we recover the polynomials P_j in time $O(s M(n)) \subset O(\alpha M(n))$. $\square$

4. The Toeplitz case

The operator associated with the Toeplitz structure is

$$\Delta[\mathbb{Z}_{n,0}, \mathbb{Z}_{n,0}^t](A) = A - \mathbb{Z}_{n,0} A \mathbb{Z}_{n,0}^t, \quad A \in \mathbb{K}^{n \times n}. \quad (10)$$

This operator is invertible: given (Y, Z) in $\mathbb{K}^{n \times \alpha} \times \mathbb{K}^{n \times \alpha}$, there exists a unique A such that $\Delta[\mathbb{Z}_{n,0}, \mathbb{Z}_{n,0}^t](A) = YZ^t$; it is given by the representation

$$A = \sum_{i=1}^{\alpha} \mathbb{L}(y_i) \mathbb{U}(z_i),$$

called in [28,29] a ΣLU representation of length α for A . Using Lemma 4, this representation allows one to compute a matrix-vector product Au in $O(\alpha M(n))$ operations. Our problem in this section is the converse one: given v in $\mathbb{K}^n$, find u such that $Au = v$ (or conclude that no such vector exists).

For large α , we improve previous algorithms, reducing their complexity from $O(\alpha^2 M(n) \log(n))$ to $O(\alpha^{\omega-1} M(n) \log^2(n))$. The structure of our algorithm is similar to those initiated by Bitmead and Anderson [7] and Morf [37,38]; the key difference consists in using the results of the previous section to perform efficiently the following operation: given Y, Z as above and $u_1, \dots, u_{\alpha}$ in $\mathbb{K}^n$, compute the α products $v_j = Au_j$.

4.1. Preliminaries

In addition to the operator in Eq. (10) we will use the operator

$$\Delta[\mathbb{Z}_{n,0}^t, \mathbb{Z}_{m,0}](A) = A - \mathbb{Z}_{n,0}^t A \mathbb{Z}_{m,0}, \quad A \in \mathbb{K}^{n \times m}.$$

Regardless of dimensions, the operators $\Delta[\mathbb{Z}_{n,0}^t, \mathbb{Z}_{m,0}]$ and $\Delta[\mathbb{Z}_{n,0}, \mathbb{Z}_{m,0}^t]$ are called respectively ϕ_- and ϕ_+ in [43,28,29]; their generators are called ϕ_- -generators and ϕ_+ -generators. From now on, we use this simplifying notation.

We give in this subsection some useful results on generators for submatrices, sums, products, ... Our contribution is Proposition 4, which is a faster version of [43, Proposition A.3] for generating matrix products; as in [28,29] we extend the result to rectangular matrices. Proofs not given here can be found in e.g. [7,38,43,28,45].

A first key feature of the operators ϕ_+ and ϕ_- is that when A is invertible, the ranks of $\phi_+(A)$ and $\phi_-(A^{-1})$ coincide. Secondly, when A is square, the ranks of $\phi_+(A)$ and $\phi_-(A)$ differ by at most 2.

The next lemma gives the complexity of converting from ϕ_- to ϕ_+ -generators; the same holds for converting back.

Lemma 14. *Given a ϕ_- -generator of length α for the matrix $A \in \mathbb{K}^{n \times n}$, one can compute a ϕ_+ -generator of length $\alpha + 2$ for A in time $O(\alpha M(n))$.*

Assuming that $n = m$, partition A into blocks as

$$A = \begin{bmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{bmatrix}, \quad (11)$$

with $A_{i,j} \in \mathbb{K}^{n_i \times n_j}$, and $n_1 + n_2 = n$. Then the rank of $\phi_+(A_{1,1})$ is at most that of $\phi_+(A)$; if $A_{1,1}$ is invertible and has its upper-left entry non-zero then the same bound holds for the Schur complement $A_{2,2} - A_{2,1}A_{1,1}^{-1}A_{1,2}$.

Next, we consider the cost of deducing generators for the blocks $A_{i,j}$ from generators for A , and conversely.

Lemma 15. *Given a ϕ_+ -generator of length α for A , one can find ϕ_+ -generators of length $O(\alpha)$ for all $A_{i,j}$ in time $O(\alpha M(n))$. Conversely, given ϕ_+ -generators of length at most α for all $A_{i,j}$, one can find a ϕ_+ -generator of length $O(\alpha)$ for A in time $O(\alpha M(n))$.*

Adding matrices given by their generators is straightforward (even though the generators thus obtained may not be minimal).

Lemma 16. *If (T, U) and (Y, Z) are ϕ_+ -generators for some matrices A and B of the same dimensions, then $([T \ Y], [U \ Z])$ is a ϕ_+ -generator for $A + B$.*

We conclude with the key novelty of this subsection, which concerns the complexity of computing a generator for the product of two matrices.

Proposition 4. *Let (T, U) and (Y, Z) be ϕ_+ -generators of lengths α and β for some matrices $A \in \mathbb{K}^{n \times m}$ and $B \in \mathbb{K}^{m \times p}$. Then one can find a ϕ_+ -generator of length $\alpha + \beta + 1$ for the product AB in time $O(\gamma^{\omega-1} M(q) \log(q))$, with $\gamma = \max\{\alpha, \beta\}$ and $q = \max\{n, m, p\}$.*

Proof. Let $V = B^t U$ and $W = \mathbb{Z}_{n,0} A \mathbb{Z}_{m,0}^t Y$; let also a (resp. b) be the lower shift of the last column of A (resp. B^t). Then, following [43, Proposition A.3], the proof of [28, Proposition 2] shows that $[T \ W \ a]$ and $[V \ Z \ -b]$ form a ϕ_+ -generator of length $\alpha + \beta + 1$ for AB . Hence, it suffices to bound the cost of computing V, W, a and b .

Let us detail the computation of $V = B^t U$ when $m \geq p$. We reduce to the square case by defining $B' = [0 \ B] \in \mathbb{K}^{m \times m}$ and $V' = B'^t U \in \mathbb{K}^{m \times \alpha}$. Since V can be read off V' , we focus on computing the latter matrix.

Remark that $\phi_+(B') = YZ'^t$ with $Z'^t = [0 \ Z^t] \in \mathbb{K}^{\beta \times m}$. Since B' is square, the Gohberg–Semencul formula then shows that

$$B' = \sum_{i=1}^{\beta} \mathbb{L}(y_i) \mathbb{U}(z'_i),$$

where y_i is the i th column of Y and z'_i the i th column of Z' . Thus, recalling that $\mathbb{J}$ denotes the reversal matrix of order m , the transpose of B' is given by

$$B'^t = \sum_{i=1}^{\beta} \mathbb{L}(z'_i) \mathbb{U}(y_i) = \sum_{i=1}^{\beta} \mathbb{J} \mathbb{U}(z'_i) \mathbb{L}(y_i) \mathbb{J}.$$

Now, let u_j (resp. v'_j) be the j th column of U (resp. V'). The previous formula for B'^t and the equation $V' = B'^t U$ thus give

$$\mathbb{J} v'_j = \sum_{i=1}^{\beta} \mathbb{U}(z'_i) \mathbb{L}(y_i) \mathbb{J} u_j.$$

In polynomial terms, in view of Lemma 4, this reads

$$\text{Rev}_m(\text{Pol}(v'_j)) = \left[\sum_{i=1}^{\beta} Z'_i (Y_i U_j \bmod x^m) \right] \text{div } x^{m-1},$$

with $Z'_i = \text{Rev}_m(\text{Pol}(z'_i))$, $Y_i = \text{Pol}(y_i)$ and $U_j = \text{Rev}_m(\text{Pol}(u_j))$. By Proposition 1, we can compute the polynomials

$$\sum_{i=1}^{\beta} Z'_i (Y_i U_j \bmod x^m), \quad j = 1, \dots, \alpha$$

in time $O(\max\{\alpha, \beta\}^{\omega-1} M(m) \log(m))$; the vectors v'_j are then deduced by coefficient extraction. The case $p > m$ is treated similarly, padding B with $p - m$ zero rows, and giving a cost of $O(\max\{\alpha, \beta\}^{\omega-1} M(p) \log(p))$. Hence, in any case, V can be obtained in time $O(\gamma^{\omega-1} M(q) \log(q))$, as claimed. The computation of W is done similarly too, by multiplying A on the right by $Z_{m,0}^t Y$, and has a similar cost estimate. Computing a and b is faster: it suffices to multiply A and B^t by a single vector, so the cost is merely $O(\gamma M(q))$. $\square$

4.2. Solving Toeplitz-like linear systems

We now prove Theorem 1. Let $(T, U, w) \in \mathbb{K}^{n \times \alpha} \times \mathbb{K}^{n \times \alpha} \times \mathbb{K}^n$ be the input of problem $\text{LinearSystem}(\mathbb{Z}_{n,0}, \mathbb{Z}_{n,0}^t, \alpha)$. As in [28,29] we will reduce by randomization to a similar problem but with a regularized input (Y, Z, v) .

Let B be (implicitly) defined by $\phi_+(B) = TU^t$, so that the system we want to solve is $Bt = w$. Define $A = \mathbb{U}(y) B \mathbb{L}(z)$ and $v = \mathbb{U}(y) w$, where y, z are random vectors in $\mathbb{K}^n$ with first entry equal to 1. Then, a vector t satisfies $Bt = w$ if and only if the vector $u = \mathbb{L}(z)^{-1} t$ satisfies $Au = v$.

Using the proof of Proposition 4 and the simple structure of the matrices $\mathbb{U}(y)$ and $\mathbb{L}(z)$, one may check that a ϕ_+ -generator (Y, Z) of length $\alpha + O(1)$ for A can be computed in time $O(\alpha M(n))$. Besides, the vector t can be recovered from u in time $M(n)$. Hence, we can focus on solving the latter problem $Au = v$.

By Theorem 2 in [31], there exists a non-zero polynomial Γ of $2n - 2$ variables and degree $n^2 + n$, such that if $\Gamma(y_2, \dots, y_n, z_2, \dots, z_n) \neq 0$, the matrix A has generic rank profile.

Supposing that this condition is satisfied, let r be the rank of A and let $A_r \in \mathbb{K}^{r \times r}$ be the largest non-singular leading principal submatrix of A . The next proposition bounds the cost of finding a ϕ_- -generator of length α for A_r^{-1} . Then, using a third random vector of size n , Theorem 4 in [31] (see also [28, Proposition 3]) shows how to find a random solution to the system $Au = v$, if one exists, in $O(\alpha M(n))$ operations. This concludes the cost analysis of Theorem 1.

The probability analysis follows from the previous discussion: the Zippel–Schwartz lemma [14,55,50] shows that this algorithm has type $P(3n - 2, n^2 + n)$.

Proposition 5. Assume that A has generic rank profile. Given a ϕ_+ -generator of length α for $A \in \mathbb{K}^{n \times n}$, one can compute its rank r as well as a ϕ_- -generator of length at most α for A_r^{-1} in time $O(\alpha^{\omega-1} M(n) \log^2(n))$.

Proof. We use Kaltofen’s Leading Principal Inverse algorithm [28,29], which follows the ideas of Morf [37,38] and Bitmead–Anderson [7]; with Lemma 10, this algorithm becomes deterministic, as noted in [47, §7]. The proof of Theorem 3 in [28] shows that its cost is $T(\alpha, n) = O(\alpha^{\omega})$ if $n \leq \alpha$ and otherwise

$$T(\alpha, n) = T(\alpha, \lceil n/2 \rceil) + T(\alpha, \lfloor n/2 \rfloor) + T_1(\alpha, n) + T_2(\alpha, n) + O(\alpha^{\omega-1} n + \alpha M(n)). \quad (12)$$

Here the term in $O(\alpha^{\omega-1}n + \alpha M(n))$ bounds the cost of some conversions between ϕ_+ - and ϕ_- -generators (Lemma 14) and the cost of finding generators of minimal length (Lemma 10); the terms $T_1(\alpha, n)$ and $T_2(\alpha, n)$ are the costs of two tasks we shall describe now, after recalling some notation from [28].

With $n_1 = \lceil n/2 \rceil$, partition A as in Eq. (11) and A_r as

$$A_r = \begin{bmatrix} A_{1,1} & A'_{1,2} \\ A'_{2,1} & A'_{2,2} \end{bmatrix}.$$

Assume that $A_{1,1}$ is non-singular (else, the cost is smaller) and let

$$\Delta = A_{2,2} - A_{2,1} A_{1,1}^{-1} A_{1,2} \quad \text{and} \quad \Delta' = A'_{2,2} - A'_{2,1} A_{1,1}^{-1} A'_{1,2}.$$

Given ϕ_+ -generators of length $O(\alpha)$ for A and $A_{1,1}^{-1}$, the first task is to compute a ϕ_+ -generator for Δ . Using Lemmas 15 and 16 and Proposition 4, its cost is thus $T_1(\alpha, n) = O(\alpha^{\omega-1}M(n) \log(n))$.

Given ϕ_+ -generators of length $O(\alpha)$ for A , $A_{1,1}^{-1}$ and Δ'^{-1} , the second task consists in computing a ϕ_+ -generator for A_r^{-1} . Recall first that (see for example Theorem 5.2.3 in [45])

$$A_r^{-1} = \begin{bmatrix} B'_{1,1} & B'_{1,2} \\ B'_{2,1} & \Delta'^{-1} \end{bmatrix} \quad \text{with} \quad \begin{aligned} B'_{1,2} &= -A_{1,1}^{-1} A'_{1,2} \Delta'^{-1} \\ B'_{2,1} &= -\Delta'^{-1} A'_{2,1} A_{1,1}^{-1} \\ B'_{1,1} &= A_{1,1}^{-1} - B'_{1,2} A'_{2,1} A_{1,1}^{-1}. \end{aligned}$$

Using again Lemmas 15 and 16 and Proposition 4, we get the same estimate as before: $T_2(\alpha, n) = O(\alpha^{\omega-1}M(n) \log(n))$. In view of the recurrence relation in Eq. (12), this implies that $T(\alpha, n) = O(\alpha^{\omega-1}M(n) \log^2(n))$. $\square$

4.3. Application: Padé-type approximation

We conclude by proving Corollary 1 on polynomial approximation. Let $M \in \mathbb{K}[x]$ be of degree n , let $f_1, \dots, f_s \in \mathbb{K}[x]$ be of degrees less than n and let $v_1, \dots, v_s$ be positive integers such that $\sum_{i \leq s} v_i = n + 1$. We look for approximants $g_1, \dots, g_s \in \mathbb{K}[x]$, not all zero, with $\deg(g_i) < v_i$ and such that $g_1 f_1 + \dots + g_s f_s = 0 \pmod{M}$.

Write $M = \sum_{i=0}^n m_i x^i$, with $m_n = 1$ and let $\mathbb{X} = \mathbb{X}(M) \in \mathbb{K}^{n \times n}$ be the matrix of multiplication by x modulo M . For $i \leq s$, define A_i as the Krylov matrix

$$A_i = [f_i \ \mathbb{X} f_i \ \dots \ \mathbb{X}^{v_i-1} f_i] \in \mathbb{K}^{n \times v_i},$$

where $f_i = [f_{i,0} \ \dots \ f_{i,n-1}]^t$ is the vector of coefficients of f_i . Let finally $A = [A_1 \ \dots \ A_s] \in \mathbb{K}^{n \times (n+1)}$ and $A' \in \mathbb{K}^{(n+1) \times (n+1)}$ be the matrix obtained by padding A with an $(n+1)$ st row full of 1's.

Since the right null space of A is nontrivial, the square system $A'u = [0 \ \dots \ 0 \ 1]^t$ admits a solution, and any such solution solves our problem. The following lemma shows the Toeplitz-like structure of the matrix A' ; combining it with Theorem 1 proves Corollary 1.

Lemma 17. *Given $M, f_1, \dots, f_s$ and $v_1, \dots, v_s$ as above, one can compute a ϕ_+ -generator of length $s + 2$ for the matrix A' in time $O(sM(n))$.*

Proof. Remark that $\mathbb{X} = \mathbb{Z}_{n,0} - m \mathbf{e}_{n,n}^t$, with

$$m = \begin{bmatrix} m_0 \\ \vdots \\ m_{n-1} \end{bmatrix} \in \mathbb{K}^n \quad \text{and} \quad \mathbf{e}_{n,n} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix} \in \mathbb{K}^n;$$

since A has dimensions $n \times (n+1)$, it follows that

$$\phi_+(A) = A - \mathbb{X} A \mathbb{Z}_{n+1,0}^t - m \mathbf{e}_{n,n}^t A \mathbb{Z}_{n+1,0}^t.$$

Let $\mathbf{a} = A^t \mathbf{e}_{n,n} \in \mathbb{K}^{n+1}$ be the transpose of the last row of A and let $\mathbf{b} = \mathbb{Z}_{n+1,0} \mathbf{a}$. Then, the previous formula becomes $\phi_+(A) = A - \mathbb{X} A \mathbb{Z}_{n+1,0}^t - m \mathbf{b}^t$.

Taking $f_0 = 0$ and $v_0 = 0$, we can write $A - \mathbb{X} A \mathbb{Z}_{n+1,0}^t$ as $Y Z^t$, with Y and Z of dimensions $n \times s$ and $(n+1) \times s$ given as follows: the i th column y_i of Y is $f_i - \mathbb{X}^{v_i-1} f_{i-1}$; the i th column z_i of Z is zero, except for a 1 at row $1 + v_1 + \dots + v_{i-1}$. Since $\mathbb{X}^{v_i-1} f_{i-1}$ is the coefficient vector of $x^{v_i-1} f_{i-1} \pmod{M}$, it can be computed in time $O(M(n))$, so Y and Z can be computed in time $O(sM(n))$.

To determine a ϕ_+ -generator of A , it remains to determine the vector $\mathbf{a}$, from which $\mathbf{b}$ follows easily. For $i \leq s$, let $\mathbf{a}_i \in \mathbb{K}^{1 \times v_i}$ be the last row of A_i , so that

$$\mathbf{a}_i = [\text{coeff}(f_i, x^{n-1}) \ \dots \ \text{coeff}(x^{v_i-1} f_i \pmod{M}, x^{n-1})].$$

Define next

$$m_i = \begin{bmatrix} m_n \\ \vdots \\ m_{n-v_i+1} \end{bmatrix} \in \mathbb{K}^{v_i} \quad \text{and} \quad f'_i = \begin{bmatrix} f_{i,n-1} \\ \vdots \\ f_{i,n-v_i} \end{bmatrix} \in \mathbb{K}^{v_i}.$$

Noticing that $\mathbb{L}(m_i) a_i^t = f'_i$, we see that the entries of a_i can be computed in time $O(M(v_i))$. Since $\sum_{i \leq s} v_i = n+1$, the vector $a = [a_1 \dots a_s]^t \in \mathbb{K}^{n+1}$ can thus be computed in time $O(M(n))$.

In conclusion, one can compute $G = [Y \ -m] \in \mathbb{K}^{n \times (s+1)}$ and $H = [Z \ b] \in \mathbb{K}^{(n+1) \times (s+1)}$ in time $O(sM(n))$, such that (G, H) is a ϕ_+ -generator of length $s+1$ for A . One then obtains a ϕ_+ -generator of length $s+2$ for A' by adjoining a last row of zeros to G and the columns

$$\begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix} \in \mathbb{K}^{n+1} \quad \text{and} \quad \begin{bmatrix} 1 \\ \vdots \\ 1 \\ 1 \end{bmatrix} - b \in \mathbb{K}^{n+1}$$

to respectively G and H . $\square$

5. The Vandermonde case

In this section, $x \in \mathbb{K}^n$ and $\psi \in \mathbb{K}$ are as in Eq. (1) and (2). The operator associated with the Vandermonde structure is

$$\Delta[\mathbb{D}(x), \mathbb{Z}_{n,\psi}^t](A) = A - \mathbb{D}(x) A \mathbb{Z}_{n,\psi}^t.$$

With our choice of ψ , Theorem 4.3.2 in [45] shows that this operator is invertible. Moreover, given Y, Z in $\mathbb{K}^{n \times \alpha}$, Example 4.4.6 in [45] shows that the unique matrix $A \in \mathbb{K}^{n \times n}$ such that $\Delta[\mathbb{D}(x), \mathbb{Z}_{n,\psi}^t](A) = YZ^t$ is

$$A = \mathbb{D}((1 - \psi x^n)^{-1}) \sum_{i=1}^{\alpha} \mathbb{D}(y_i) \mathbb{V}(x, n) \mathbb{T}(z_i, \psi)^t. \quad (13)$$

Following Pan's idea [41,42], we will prove Theorem 2 on the complexity of solving Vandermonde-like systems of the form $Au = v$ by turning them into Toeplitz-like ones.

We will use the same kind of reduction as in [24]. However, that approach requires the entries of x to be pairwise distinct, i.e., that $\mathbb{V}(x, n)$ be invertible; else, the preprocessing step in [24, Section 2] fails. Similarly, the reduction in [45, Example 4.8.4] does not solve the problem when $\mathbb{V}(x, n)$ is singular (in the application of Section 5.4, this invertibility assumption does not hold).

The partition of x of Eq. (1) will be used to solve this problem in cases when x has low multiplicity: this is exposed in Sections 5.1 and 5.2. When the multiplicity becomes too large, some extra work is needed, presented in Section 5.3.

5.1. A multiplication problem

We start by solving a preliminary subproblem. Let Y and Z be in $\mathbb{K}^{n \times \alpha}$, and let A be the unique $n \times n$ matrix such that $\Delta[\mathbb{D}(x), \mathbb{Z}_{n,\psi}^t](A) = YZ^t$. Splitting A along its rows according to the given partition of x , we thus write

$$x = \begin{bmatrix} x_1 \\ \vdots \\ x_s \end{bmatrix}, \quad \text{with } x_j \in \mathbb{K}^{v_j} \quad \text{and} \quad A = \begin{bmatrix} A_1 \\ \vdots \\ A_s \end{bmatrix}, \quad \text{with } A_j \in \mathbb{K}^{v_j \times n}, \quad (14)$$

where all x_j are repetition-free. Given vectors $w_1, \dots, w_s$, with w_j in $\mathbb{K}^{v_j}$, we study in this subsection the cost of computing all products $A_j^t w_j \in \mathbb{K}^n$, using the results of Section 3.2. This will be the key in our reduction of Vandermonde-like systems to Toeplitz-like ones.

Proposition 6. *Given x, ψ, Y, Z and $w_1, \dots, w_s$ as above, and assuming that $s \leq \alpha$, one can compute all products $A_j^t w_j$ in time $O(\alpha^{\omega-1} M(n) \log(n))$.*

Proof. Let y_i and z_i be the columns of Y and Z . We adapt the partition of x and A to the vectors y_i , writing

$$y_i = \begin{bmatrix} y_{i,1} \\ \vdots \\ y_{i,s} \end{bmatrix}, \quad \text{with } y_{i,j} \in \mathbb{K}^{v_j}.$$

Since A is given by Eq. (13), its submatrices A_j are given by $A_j = \mathbb{D}((1 - \psi x_j^n)^{-1}) B_j$, with

$$B_j = \sum_{i=1}^{\alpha} \mathbb{D}(y_{i,j}) \mathbb{V}(x_j, n) \mathbb{T}(z_i, \psi)^t$$

and thus

$$B_j^t = \sum_{i=1}^{\alpha} \mathbb{T}(z_i, \psi) \mathbb{V}(x_j, n)^t \mathbb{D}(y_{i,j}). \quad (15)$$

For $j \leq s$, let $f_j = \mathbb{D}((1 - \psi x_j^n)^{-1}) w_j$. The vectors f_j can be deduced from x , ψ and the w_j in $O(n \log(n))$ operations. Since we have $A_j^t w_j = B_j^t f_j$, we are thus left with computing all the products $B_j^t f_j$.

Recalling that $x_j \in \mathbb{K}^{v_j}$ is repetition-free, we let $Y_{i,j} = \text{Interp}(x_j, y_{i,j}) \in \mathbb{K}[x]_{v_j}$. Applying Lemma 1 to the right-hand side of Eq. (15) then gives

$$B_j^t = \sum_{i=1}^{\alpha} \mathbb{T}(z_i, \psi) \mathbb{M}(Y_{i,j}, n)^t \mathbb{V}(x_j, n + v_j - 1)^t.$$

We can now factor out the rightmost transposed Vandermonde matrices, which do not depend on the summation index i . Defining

$$f'_j = \mathbb{V}(x_j, n + v_j - 1)^t f_j,$$

we deduce that

$$B_j^t f_j = \sum_{i=1}^{\alpha} \mathbb{T}(z_i, \psi) \mathbb{M}(Y_{i,j}, n)^t f'_j.$$

By fast application of a transposed Vandermonde matrix, each vector f'_j can be computed in time $O(M(n) \log(n))$; hence, the total time for their computation is $O(\alpha M(n) \log(n))$.

For $i \leq \alpha$ and $j \leq s$, define the vector $g_{i,j}$ in $\mathbb{K}^n$ by

$$g_{i,j} = \mathbb{T}(z_i, \psi) \mathbb{M}(Y_{i,j}, n)^t f'_j,$$

so that $B_j^t f_j = \sum_{i=1}^{\alpha} g_{i,j}$. We will obtain the vectors $g_{i,j}$ by means of Lemma 6. To do so, define the polynomials

$$Y'_{i,j} = \text{Rev}_{v_j}(Y_{i,j}), \quad F_j = \text{Pol}(f'_j) \quad \text{and} \quad F'_j = \text{Rev}_{n+v_j-1}(F_j).$$

Let also $Z_i = \text{Pol}(z_i)$ and $Z'_i = \text{Pol}(\text{Flip}(z_i))$. Lemma 6 then gives

$$\text{Pol}(g_{i,j}) = Z_i (Y'_{i,j} F_j \text{div } x^{v_j-1}) \bmod x^n + \psi \text{Rev}_n(Z'_i (Y_{i,j} F'_j \text{div } x^{v_j-1}) \bmod x^n).$$

Summing over i eventually yields our output

$$\begin{aligned} \text{Pol}(B_j^t f_j) &= \sum_{i=1}^{\alpha} \text{Pol}(g_{i,j}) = \sum_{i=1}^{\alpha} Z_i (Y'_{i,j} F_j \text{div } x^{v_j-1}) \bmod x^n \\ &\quad + \psi \text{Rev}_n \left(\sum_{i=1}^{\alpha} Z'_i (Y_{i,j} F'_j \text{div } x^{v_j-1}) \bmod x^n \right). \end{aligned} \quad (16)$$

It remains to perform the cost estimate, using the results recalled in Section 2.

- Using fast interpolation, we compute each $Y_{i,j}$ in time $O(M(v_j) \log(v_j))$ and thus all $Y_{i,j}$ and $Y'_{i,j}$ in time $O(\alpha M(n) \log(n))$.
- Applying Proposition 3 to both summands in Eq. (16) shows that all polynomials $\text{Pol}(B_j^t f_j)$ can be computed in time $O(\alpha^{\omega-1} M(n) \log(n))$, which concludes the proof. $\square$

5.2. The case of low multiplicities

In this subsection, we reduce the resolution of Vandermonde-like systems to that of Toeplitz-like systems. We adapt the reduction of [24], allowing now for repetitions in x . For the moment, we work in the case where the multiplicity of x is bounded by α .

Proposition 7. *Let x and ψ be as in Eqs. (1) and (2). If the multiplicity s of x satisfies $s \leq \alpha$, then the problem $\text{LinearSystem}(\mathbb{D}(x), \mathbb{Z}_{n,\psi}^t, \alpha)$ can be solved in time $O(\alpha^{\omega-1} M(n) \log^2(n))$. The algorithm is probabilistic of type $P(3n-2, n^2+n)$.*

Proof. Given Y and Z in $\mathbb{K}^{n \times \alpha}$ and v in $\mathbb{K}^n$, we are looking for solutions u to the system $Au = v$, where A is such that $\Delta[\mathbb{D}(x), \mathbb{Z}_{n,\psi}^t](A) = YZ^t$.

For $j \leq s$, let M_j be the monic polynomial $M_j = \prod_{a \in x_j} (x - a)$. Since x_j is repetition-free, Lemma 3 shows that

$$\mathbb{D}(x_j) = \mathbb{V}(x_j, v_j) \mathbb{X}(M_j) \mathbb{V}(x_j, v_j)^{-1}.$$

It follows that

$$\mathbb{D}(x) = \mathbb{V} \mathbb{X} \mathbb{V}^{-1}, \quad (17)$$

where $\mathbb{V}$ and $\mathbb{X}$ are block-diagonal with respective blocks $\mathbb{V}(x_j, v_j)$ and $\mathbb{X}(M_j)$.

For $j \leq s$, let $m_j \in \mathbb{K}^{v_j}$ be the coefficient vector of $-M_j \bmod x^{v_j}$, so that $\mathbb{X}(M_j) = \mathbb{Z}_{v_j,0} + m_j e_{v_j, v_j}^t$. Hence, writing $v_j^* = v_1 + \dots + v_j$ as in Section 2,

$$\mathbb{X} = \mathbb{Z}_{n,0} + \sum_{j=1}^s g_j e_{n, v_j^*}^t, \quad (18)$$

where for $j \leq s$, $g_j \in \mathbb{K}^n$ is obtained by padding m_j with v_{j-1}^* zeros on the top and, if $j \neq s$, with -1 followed by $n - v_j^* - 1$ zeros on the bottom.

Defining $B = \mathbb{V}^{-1}A$ and $v' = \mathbb{V}^{-1}v$, solving $Au = v$ amounts to solve $Bu = v'$. To do so in the claimed complexity, we exhibit the Toeplitz-like structure of B and bound the cost of computing v' and a generator for B . Pre-multiplying by $\mathbb{V}^{-1}$ the relation

$$A - \mathbb{D}(x) A \mathbb{Z}_{n,\psi}^t = YZ^t,$$

we get

$$B - \mathbb{V}^{-1} \mathbb{D}(x) \mathbb{V} B \mathbb{Z}_{n,\psi}^t = \mathbb{V}^{-1} YZ^t;$$

using Eq. (17), we rewrite this as

$$B - \mathbb{X} B \mathbb{Z}_{n,\psi}^t = Y' Z^t, \quad \text{with } Y' = \mathbb{V}^{-1} Y.$$

Then, from (18) and the relation $\mathbb{Z}_{n,\psi} = \mathbb{Z}_{n,0} + \psi e_{n,1} e_{n,n}^t$, we deduce that

$$B - \mathbb{Z}_{n,0} B \mathbb{Z}_{n,0}^t = \psi \mathbb{Z}_{n,0} B e_{n,n} e_{n,1}^t + \left(\sum_{j=1}^s g_j e_{n, v_j^*}^t \right) B \mathbb{Z}_{n,\psi}^t + Y' Z^t.$$

Define the vectors $f_1 = \psi \mathbb{Z}_{n,0} B e_{n,n}$ and, for $j \leq s$, $h_j = B^t e_{n, v_j^*}$ and $h'_j = \mathbb{Z}_{n,\psi} h_j$. The above formula then becomes

$$\Delta[\mathbb{Z}_{n,0}, \mathbb{Z}_{n,0}^t](B) = f_1 e_{n,1}^t + G H^t + Y' Z^t,$$

where G (resp. H') has columns g_j (resp. h'_j). The matrices $[f_1 \ G \ Y']$ and $[e_{n,1} \ H' \ Z]$ thus form a $\mathbb{Z}_{n,0}, \mathbb{Z}_{n,0}^t$ -generator of length $\alpha + s + 1 \leq 2\alpha + 1$ for B . Once this generator and v' are known, the system $Bu = v'$ can be solved within the prescribed complexity by the probabilistic algorithm of Theorem 1.

It remains to estimate the cost of computing v' , f_1 , Y' , G and H' . We will do so using Proposition 6 as well as the results recalled in Section 2 on the complexity of polynomial operations.

As a first step, though, we detail further the structure of the vectors h_j . For $j \leq s$, one has $h_j = B^t e_{n, v_j^*} = A^t \mathbb{V}^{-t} e_{n, v_j^*}$. In view of the block structures of A^t and $\mathbb{V}^{-1}$, this can be rewritten as $h_j = A_j^t w_j$, where $w_j = \mathbb{V}(x_j, v_j)^{-t} e_{v_j, v_j}$ is the last column of $\mathbb{V}(x_j, v_j)^{-t}$.

- All polynomials M_j (and thus all vectors g_j and the matrix G) can be constructed from their roots in $O(M(n) \log(n))$ operations.
- In view of Eq. (13), one can multiply A by the vector $e_{n,n}$ in time $O(\alpha M(n) \log(n))$, see e.g. [24, Section 2].
- Since multiplication by $\mathbb{V}(x_j, v_j)^{-1}$ has cost $O(M(v_j) \log(v_j))$, multiplication by $\mathbb{V}^{-1}$ has cost $O(M(n) \log(n))$. This implies that $f_1 = \psi \mathbb{Z}_{n,0} \mathbb{V}^{-1} A e_{n,n}$ can be deduced from $A e_{n,n}$ in time $O(M(n) \log(n))$. Similarly, one can compute $Y' = \mathbb{V}^{-1} Y$ in time $O(\alpha M(n) \log(n))$ and $v' = \mathbb{V}^{-1} v$ in time $O(M(n) \log(n))$.
- Computing the last column of $\mathbb{V}(x_j, v_j)^{-t}$ takes time $O(M(v_j) \log(v_j))$, which induces a cost of $O(M(n) \log(n))$ for finding all vectors w_j .
- Knowing all w_j , Proposition 6 shows that all vectors h_j can be computed in time $O(\alpha^{\omega-1} M(n) \log(n))$. Deducing the vectors h'_j takes time $O(\alpha n)$, which concludes the proof. $\square$

5.3. The case of high multiplicities

We conclude the proof of [Theorem 2](#) by considering the case of high multiplicities ($s > \alpha$), reducing it to the case of low multiplicities ($s \leq \alpha$) seen in [Section 5.2](#). Our reduction has cost $O(\alpha^{\omega-1} n)$, which fits in the requested bound. It however introduces an extra probabilistic aspect; combined with the one of [Proposition 7](#), it yields the overall probability estimate of [Theorem 2](#).

Proposition 8. *Let $\mathbf{x}$ and ψ be as in Eqs. (1) and (2). If $\mathbf{x}$ has multiplicity larger than α , one can reduce the problem $\text{LinearSystem}(\mathbb{D}(\mathbf{x}), \mathbb{Z}_{n,\psi}^t, \alpha)$ to the problem $\text{LinearSystem}(\mathbb{D}(\mathbf{y}), \mathbb{Z}_{n,\psi}^t, \alpha)$, where $\mathbf{y} \in \mathbb{K}^n$ has multiplicity at most α and satisfies the constraints of Eqs. (1) and (2). The reduction can be done in time $O(\alpha^{\omega-1} n)$ by a probabilistic algorithm of type $P(n, 3n^2)$.*

Proof. Given $\mathbf{Y}$ and $\mathbf{Z}$ in $\mathbb{K}^{n \times \alpha}$ and $\mathbf{v}$ in $\mathbb{K}^n$, we are looking for solutions $\mathbf{u}$ to the system $\mathbf{A}\mathbf{u} = \mathbf{v}$, where $\mathbf{A}$ is such that $\Delta[\mathbb{D}(\mathbf{x}), \mathbb{Z}_{n,\psi}^t](\mathbf{A}) = \mathbf{Y}\mathbf{Z}^t$. We assume that $\mathbf{Y}$ and $\mathbf{Z}$ have full rank; if this is not the case, we can replace $(\mathbf{Y}, \mathbf{Z})$ by a minimal-length generator, whose two matrices then have full rank.

We start by reordering the entries of $\mathbf{x}$ to obtain a repetition-free decomposition as in [Eq. \(6\)](#) of [Section 2](#):

$$\mathbf{x}' = \begin{bmatrix} x_{\sigma(1)} \\ \vdots \\ x_{\sigma(n)} \end{bmatrix} = \begin{bmatrix} x'_1 \\ \vdots \\ x'_t \end{bmatrix},$$

where x'_i is a vector consisting of μ_i repetitions of the same element ξ_i , so that $n = \mu_1 + \dots + \mu_t$, with $\xi_i \neq \xi_j$ for $i \neq j$ and $\mu_1 \geq \dots \geq \mu_t > 0$.

Let $\mathbf{v}'$, $\mathbf{A}'$ and $\mathbf{Y}'$ be obtained by applying the same reordering to the entries of $\mathbf{v}$ and to the rows of $\mathbf{A}$ and $\mathbf{Y}$; hence, $(\mathbf{Y}', \mathbf{Z})$ is a $\mathbb{D}(\mathbf{x}')$, $\mathbb{Z}_{n,\psi}^t$ -generator for $\mathbf{A}'$. Since the solution sets of $\mathbf{A}\mathbf{u} = \mathbf{v}$ and $\mathbf{A}'\mathbf{u} = \mathbf{v}'$ are the same, we focus on the latter problem.

By construction, the matrices $\mathbf{A}'$ and $\mathbf{Y}'$ admit the following decompositions:

$$\mathbf{A}' = \begin{bmatrix} \mathbf{A}'_1 \\ \vdots \\ \mathbf{A}'_t \end{bmatrix} \quad \text{and} \quad \mathbf{Y}' = \begin{bmatrix} \mathbf{Y}'_1 \\ \vdots \\ \mathbf{Y}'_t \end{bmatrix},$$

with $\mathbf{A}'_i$ in $\mathbb{K}^{\mu_i \times n}$ and $\mathbf{Y}'_i$ in $\mathbb{K}^{\mu_i \times \alpha}$. Remark then that $\mathbf{A}'_i = \mathbb{D}(\mathbf{x}'_i) \mathbf{A}'_i \mathbb{Z}_{n,\psi}^t = \mathbf{Y}'_i \mathbf{Z}^t$. Now, $\mathbb{D}(\mathbf{x}'_i)$ equals $\xi_i \mathbb{I}_{\mu_i}$ and, since $\psi \xi_i^n \neq 1$ for all i , all matrices $\mathbb{I}_n - \xi_i \mathbb{Z}_{n,\psi}^t$ are invertible. We thus obtain the equalities

$$\mathbf{A}'_i = \mathbf{Y}'_i \mathbf{Z}^t (\mathbb{I}_n - \xi_i \mathbb{Z}_{n,\psi}^t)^{-1} \quad \text{for } 1 \leq i \leq t. \quad (19)$$

We will use dense matrix methods to reduce the number of non-zero entries in $\mathbf{A}'$ and $\mathbf{Y}'$, while maintaining a Vandermonde-like structure. Let thus τ be such that $\mu_\tau > \alpha \geq \mu_{\tau+1}$.

- For $1 \leq i \leq \tau$, we have $\mu_i > \alpha$. We let $J_i \subset \{1, \dots, \mu_i\}$ and $\mathbf{E}_i \in \mathbb{K}^{\mu_i \times \mu_i}$ be the index set and the matrix obtained by applying [Lemma 9](#) to $\mathbf{Y}'_i$. Let also $r_i = |J_i| = \text{rank}(\mathbf{Y}'_i)$, so that $r_i \leq \alpha$. [Lemma 9](#) and [Eq. \(19\)](#) then give

$$\mathbf{E}_i \mathbf{Y}'_i = \begin{bmatrix} \mathbf{Y}''_i \\ 0 \end{bmatrix} \quad \text{and} \quad \mathbf{E}_i \mathbf{A}'_i = \begin{bmatrix} \mathbf{A}''_i \\ 0 \end{bmatrix},$$

where $\mathbf{Y}''_i$ and $\mathbf{A}''_i$ have respective sizes $(r_i \times \alpha)$ and $(r_i \times n)$, and $\mathbf{Y}''_i$ consists of the rows of $\mathbf{Y}'_i$ indexed by J_i .

- For $\tau < i \leq t$, we have $\mu_i \leq \alpha$. We let $\mathbf{E}_i = \mathbb{I}_{\mu_i}$, $\mathbf{Y}''_i = \mathbf{Y}'_i$ and $\mathbf{A}''_i = \mathbf{A}'_i$.

Let $\mathbf{E}$ be the block-diagonal matrix having $\mathbf{E}_1, \dots, \mathbf{E}_t$ on the diagonal. Hence, $\mathbf{Y}'' = \mathbf{E} \mathbf{Y}'$ consists of the matrices $\mathbf{Y}''_i$, interleaved by blocks of zeros when $i \leq \tau$; the same holds for $\mathbf{A}'' = \mathbf{E} \mathbf{A}'$. Besides, since each $\mathbb{D}(\mathbf{x}'_i)$ is a homothety matrix, it commutes with $\mathbf{E}_i$; we deduce that $(\mathbf{Y}'', \mathbf{Z})$ is a $\mathbb{D}(\mathbf{x}'')$, $\mathbb{Z}_{n,\psi}^t$ -generator for $\mathbf{A}''$. Finally, since $\mathbf{E}$ is invertible, the solution sets of $\mathbf{A}'\mathbf{u} = \mathbf{v}'$ and $\mathbf{A}''\mathbf{u} = \mathbf{v}''$ coincide, where we wrote $\mathbf{v}'' = \mathbf{E} \mathbf{v}'$.

We solve the latter problem, by exhibiting the Vandermonde-like structure of $\mathbf{A}''$ for a modified displacement operator with lower multiplicity. Define a vector $\mathbf{x}'' \in \mathbb{K}^n$ by replacing, for $i \leq \tau$, the last $\mu_i - r_i$ entries of $\mathbf{x}'_i$ by new values taken from $\mathbb{K}$. Due to the presence of corresponding blocks of zeros in $\mathbf{Y}''$ and $\mathbf{A}''$, the matrices $\mathbf{Y}''$ and $\mathbf{Z}$ now form a $\mathbb{D}(\mathbf{x}'')$, $\mathbb{Z}_{n,\psi}^t$ -generator for $\mathbf{A}''$.

With $r = \sum_{i \leq \tau} (\mu_i - r_i)$, suppose that the new values $y_1, \dots, y_r$ inserted in $\mathbf{x}''$ are pairwise distinct and that none of them belongs to $\mathbf{x}'$ or satisfies $\psi y_i^n = 1$. Then, $\mathbf{x}''$ satisfies the constraint of [Eq. \(2\)](#) and has multiplicity at most α (since all r_i are at most α). It remains to reorder the entries of $\mathbf{x}''$ to obtain a vector $\mathbf{y}$ that also satisfies [Eq. \(1\)](#), and let Σ be the permutation matrix such that $\mathbf{y} = \Sigma \mathbf{x}''$. Defining $\mathbf{U} = \Sigma \mathbf{Y}''$, it follows that $(\mathbf{U}, \mathbf{Z})$ is a $\mathbb{D}(\mathbf{y})$, $\mathbb{Z}_{n,\psi}^t$ -generator for the matrix $\mathbf{B} = \Sigma \mathbf{A}''$. To conclude, let $\mathbf{w} = \Sigma \mathbf{v}''$; then, the solution sets of $\mathbf{A}''\mathbf{u} = \mathbf{v}''$ and $\mathbf{B}\mathbf{u} = \mathbf{w}$ coincide.

We have thus reduced solving the system $\mathbf{A}\mathbf{u} = \mathbf{v}$ to solving $\mathbf{B}\mathbf{u} = \mathbf{w}$, while providing a generator for $\mathbf{B}$ with respect to the Vandermonde-like structure $\mathbb{D}(\mathbf{y})$, $\mathbb{Z}_{n,\psi}^t$. It only remains to perform the complexity analysis.

- By Lemma 10, the cost of making the input generator minimal is $O(\alpha^{\omega-1}n)$.
- The permutation σ that gives x' can be computed in time $O(n)$ by Lemma 11.
- For $i \leq \tau$, since Y'_i is in $\mathbb{K}^{\mu_i \times \alpha}$ with $\mu_i \geq \alpha$, one can compute E_i in time $O(\alpha^{\omega-1}\mu_i)$ by Lemma 9. Since $\sum_{i=1}^{\tau} \mu_i \leq n$, the total cost is in $O(\alpha^{\omega-1}n)$.
- Since each matrix E_i is given in the form (r_i, J_i, G_i, P_i) of Lemma 9, and thus has $O(\alpha\mu_i)$ non-zero entries, v'' can be deduced from v' in $O(\alpha n)$ operations.
- To put x'' into the form of Eq. (1), we first put it into the repetition-free form of Eq. (6). Since the repeated entries are already grouped together (and thus the multiplicities are known), all we have to do is to sort the multiplicities in decreasing order; this can be done in time $O(n)$ using bucket sorting. Then, Lemma 11 puts the result into the form of Eq. (1) in time $O(n)$.
- All other operations amount to apply permutations to the entries of matrices and vectors and have negligible cost.

The probability analysis comes by remarking that the values $y_1, \dots, y_r$ satisfy our requirements if they do not cancel the polynomial

$$\delta = \prod_{i < j \leq r} (Y_i - Y_j) \times \prod_{i \leq r, j \leq t} (Y_i - \xi_j) \times \prod_{i \leq r} (Y_i^n \psi - 1), \quad (20)$$

which, since r and t are bounded by n , has degree at most $3n^2$. The Zippel–Schwartz lemma gives the required probability estimate. $\square$

5.4. Application: Bivariate interpolation

We conclude by proving Corollary 2 on the complexity of bivariate interpolation. Let $P_1, \dots, P_s$ be the lists of sample points

$$\begin{aligned} P_1 &= [p_{1,1} = (x_1, y_{1,1}), \dots, p_{1,v_1} = (x_1, y_{1,v_1})], \\ &\vdots \\ P_s &= [p_{s,1} = (x_s, y_{s,1}), \dots, p_{s,v_s} = (x_s, y_{s,v_s})], \end{aligned}$$

with $v_1 \geq \dots \geq v_s > 0$. Given a vector of values $v = [v_{i,j}] \in \mathbb{K}^n$, with $1 \leq i \leq s$ and $1 \leq j \leq v_i$, there exists a unique polynomial F in $\mathbb{K}[x, y]$ of the form

$$F = \sum_{0 \leq i < s, 0 \leq j < v_{i+1}} f_{i,j} x^i y^j$$

such that $F(p_{i,j}) = v_{i,j}$ for all i, j . To recover F , we will use the Vandermonde-like structure of the corresponding linear system.

The support of F is thus the set of all monomials $x^i y^j$, with $0 \leq i < s$ and $0 \leq j < v_{i+1}$. To arrange these monomials in a suitable order, we let $\mu = \mu_1 \geq \dots \geq \mu_t$ be the conjugate partition of v , with $s = \mu_1$ and $t = v_1$; for $1 \leq j \leq t$, we define

$$B_j = [x^{i-1} y^{j-1} \mid 1 \leq i \leq \mu_j]$$

and let B be the concatenation of the lists B_j . Thus, for a fixed j , the entries of B_j have the same degree in y and increasing degrees in x .

It will be convenient to rearrange the lists of sample points, by “transposing” the input lists $P_1, \dots, P_s$ to obtain

$$\begin{aligned} Q_1 &= [p_{1,1}, p_{2,1}, \dots, p_{\mu_1,1}], \\ Q_2 &= [p_{1,2}, p_{2,2}, \dots, p_{\mu_2,2}], \\ &\vdots \\ Q_t &= [p_{1,t}, p_{2,t}, \dots, p_{\mu_t,t}]; \end{aligned}$$

we then let Q be the concatenation of $Q_1, \dots, Q_t$. Taking the x -coordinates of the elements in Q , we obtain a vector x in $\mathbb{K}^n$; by construction, x is in the form of Eq. (1). Let further $w \in \mathbb{K}^n$ be obtained by rearranging v in the same way. By Lemma 11, one can deduce Q and w from P and v in time $O(n)$.

Let $\text{Span}(B) \subset \mathbb{K}[x, y]$ be the vector space generated by B ; we are thus concerned by the evaluation map $F \in \text{Span}(B) \mapsto [F(p)]_{p \in Q}$ and its inverse. Let

$$A = [b(p)]_{p \in Q, b \in B} \in \mathbb{K}^{n \times n}$$

be the matrix of this map, with rows indexed by Q and columns by B . Hence,

$$A = [A_1 \quad \dots \quad A_t], \quad \text{with } A_j = [b(p)]_{p \in Q, b \in B_j} \in \mathbb{K}^{n \times \mu_j}.$$

Then $\Delta[\mathbb{D}(x), \mathbb{Z}_{n,0}^t](A)$ can be written GH^t , with G and H in $\mathbb{K}^{n \times t}$; the j th column of H is zero, except for a 1 at row $1 + \mu_{j-1}^*$; the j th column of G is

$$\begin{bmatrix} 1 & \cdots & 1 \end{bmatrix}^t \text{ (for } j = 1) \text{ or } [y(p)^{j-1} - y(p)^{j-2} x(p)^{\mu_{j-1}}]_{p \in Q}^t \text{ (for } j > 1),$$

with $x(p)$ the x -coordinate of $p \in Q$ and $y(p)$ its y -coordinate. Given Q , the matrices G and H can be computed in time $O(t n \log(n))$; [Theorem 2](#) then shows that the system $Af = w$ can be solved in time $O(t^{\omega-1} M(n) \log^2(n))$, where f is the coefficient vector of the polynomial to interpolate and w is the rearranged value vector. Remembering that $t = v_1$ concludes the proof.

6. The Cauchy case

In this section, $x, y \in \mathbb{K}^n$ are as in [Eqs. \(1\), \(4\) and \(5\)](#). The operator associated with the Cauchy structure that we consider here is

$$\Delta[\mathbb{D}(x), \mathbb{D}(y)](A) = A - \mathbb{D}(x) A \mathbb{D}(y).$$

It follows from [Eq. \(5\)](#) and [\[45, Theorem 4.3.2\]](#) that such an operator is invertible. Moreover, given G and H in $\mathbb{K}^{n \times \alpha}$, [Example 4.4.7](#) in [\[45\]](#) shows that the unique matrix $A \in \mathbb{K}^{n \times n}$ such that $\Delta[\mathbb{D}(x), \mathbb{D}(y)](A) = GH^t$ is

$$A = \sum_{i=1}^{\alpha} \mathbb{D}(g_i) \mathbb{C}(x, y) \mathbb{D}(h_i). \quad (21)$$

We will prove [Theorem 3](#) on the complexity of solving Cauchy-like systems of the form $Au = v$ by turning them into Vandermonde-like ones with low multiplicity, and then using the result of [Proposition 7](#). Our reduction follows the one in [\[24, Section 3\]](#) but adds the possibility of handling repetitions among the entries of the vectors x and y above.

The organization of this section is very similar to that of the previous one; the technical arguments are in the same vein as well. We shall start in [Section 6.1](#) with solving another multiplication problem, in the spirit of the one used for the Vandermonde case. Together with the partitions of x and y given in [Eqs. \(1\) and \(4\)](#), this allows to handle in [Section 6.2](#) the cases where multiplicities are small. Large multiplicities are treated in [Section 6.3](#) with the same tools as for the Vandermonde case.

6.1. A multiplication problem

Let G and H be in $\mathbb{K}^{n \times \alpha}$, and let A be the unique $n \times n$ matrix such that $\Delta[\mathbb{D}(x), \mathbb{D}(y)](A) = GH^t$. Splitting A along its rows according to the partition of x and along its columns according to the partition of y , we thus write

$$x = \begin{bmatrix} x_1 \\ \vdots \\ x_s \end{bmatrix}, \quad \text{with } x_j \in \mathbb{K}^{v_j}, \quad y = \begin{bmatrix} y_1 \\ \vdots \\ y_t \end{bmatrix}, \quad \text{with } y_k \in \mathbb{K}^{\delta_k},$$

and

$$A = \begin{bmatrix} A_{1,1} & \cdots & A_{1,t} \\ \vdots & & \vdots \\ A_{s,1} & \cdots & A_{s,t} \end{bmatrix}, \quad \text{with } A_{j,k} \text{ in } \mathbb{K}^{v_j \times \delta_k}. \quad (22)$$

Given vectors $w_1, \dots, w_t$, with $w_k \in \mathbb{K}^{\delta_k}$, we study in this subsection the cost of computing all the products $A_{j,k} w_k \in \mathbb{K}^{v_j}$ when both s and t are bounded by α . The result below, which relies on [Proposition 2](#) of [Section 3](#), will be the key for reducing Cauchy-like systems to Vandermonde-like ones.

Proposition 9. *Given x, y, G, H and $w_1, \dots, w_t$ as above, and assuming that $\max\{s, t\} \leq \alpha$, one can compute all the products $A_{j,k} w_k$ using $O(\alpha^{\omega-1} M(n) \log(n))$ operations in $\mathbb{K}$.*

Proof. Let g_i and h_i be the columns of G and H . We adapt the partition of x to g_i and the one of y to h_i , writing

$$g_i = \begin{bmatrix} g_{i,1} \\ \vdots \\ g_{i,s} \end{bmatrix}, \quad \text{with } g_{i,j} \text{ in } \mathbb{K}^{v_j}, \quad h_i = \begin{bmatrix} h_{i,1} \\ \vdots \\ h_{i,t} \end{bmatrix}, \quad \text{with } h_{i,k} \text{ in } \mathbb{K}^{\delta_k}.$$

For $k \leq t$, let $F_k = \prod_{\gamma \in y_k} (1 - x\gamma)$ and let $v_k \in \mathbb{K}^{\delta_k}$ be such that $\text{Pol}(v_k) = F_k \bmod x^{\delta_k}$. Using [Eq. \(21\)](#) and [Lemma 2](#), it follows that the submatrices $A_{j,k}$ of A are given by $A_{j,k} = \mathbb{D}(F_k(x_j))^{-1} B_{j,k}$, with

$$B_{j,k} = \sum_{i=1}^{\alpha} \mathbb{D}(g_{i,j}) \mathbb{V}(x_j, \delta_k) \mathbb{L}(v_k) \mathbb{V}(y_k, \delta_k)^t \mathbb{D}(h_{i,k}). \quad (23)$$

Remark first that the matrices $\mathbb{D}(F_k(x_j))$ can easily be obtained, using two polynomial operations of [Section 2](#):

- Each polynomial F_k can be constructed from its roots in time $O(M(\delta_k) \log(\delta_k))$. Since $\sum_{k \leq t} \delta_k = n$, the cost of getting all of them is in $O(M(n) \log(n))$.
- By fast evaluation, each $F_k(x) \in \mathbb{K}^n$ can be obtained in $O(M(n) \log(n))$ operations, and thus all of them in time $O(t M(n) \log(n)) \subset O(\alpha M(n) \log(n))$.

We now turn to the computation of the vectors $B_{j,k} w_k$. Since x_j and y_k are repetition-free, we let $Q_{i,j} = \text{Interp}(x_j, g_{i,j})$ and $S_{i,k} = \text{Interp}(y_k, h_{i,k})$. Applying [Lemma 1](#) twice to the right-hand side of Eq. (23) then gives

$$B_{j,k} = \sum_{i=1}^{\alpha} \mathbb{V}(x_j, v_j + \delta_k - 1) \mathbb{M}(Q_{i,j}, \delta_k) \mathbb{L}(v_k) \mathbb{M}(S_{i,k}, \delta_k)^t \mathbb{V}(y_k, 2\delta_k - 1)^t.$$

We can now factor out the leftmost and rightmost (transposed) Vandermonde matrices, which do not depend on the summation index i . Defining

$$f_k = \mathbb{V}(y_k, 2\delta_k - 1)^t w_k \quad (24)$$

and

$$w'_{j,k} = \sum_{i=1}^{\alpha} \mathbb{M}(Q_{i,j}, \delta_k) \mathbb{L}(v_k) \mathbb{M}(S_{i,k}, \delta_k)^t f_k, \quad (25)$$

we deduce from the previous equation that

$$B_{j,k} w_k = \mathbb{V}(x_j, v_j + \delta_k - 1) w'_{j,k}. \quad (26)$$

It remains to estimate the cost of computing all vectors $B_{j,k} w_k$ by means of Eqs. (24)–(26), and deducing the desired vectors $A_{j,k} w_k$. The costs of the first and last steps follow directly from the reminders of Section 2:

- By fast transposed evaluation, each vector f_k is obtained in $O(M(\delta_k) \log(\delta_k))$ operations. Since $\sum_{k \leq t} \delta_k = n$, this gives a total cost of $O(M(n) \log(n))$.
- Assuming that we know $w'_{j,k}$, we deduce $B_{j,k} w_k$ by fast evaluation, in

$$O(M(v_j + \delta_k) \log(v_j + \delta_k))$$

operations. A total cost of $O(\alpha M(n) \log(n))$ then follows from the facts that

$$M(v_j + \delta_k) \leq (v_j + \delta_k) M(2n)/2n \quad \text{and} \quad \sum_{j \leq s} \sum_{k \leq t} (v_j + \delta_k) = (s + t) n.$$

- One recovers the vectors $A_{j,k} w_k$ with $t n \in O(\alpha n)$ divisions.

It only remains to bound the cost of deducing the vectors $w'_{j,k}$ in Eq. (25) from the vectors f_k . Recall that $\text{Pol}(v_k) = F_k \bmod x^{\delta_k}$ and let $S'_{i,k} = \text{Rev}_{\delta_k}(S_{i,k})$ and $F'_k = \text{Pol}(f_k)$. Then, using [Lemmas 4](#) and [5](#), we obtain

$$\text{Pol}(w'_{j,k}) = \sum_{i=1}^{\alpha} Q_{i,j} R_{i,k}, \quad (27)$$

where $Q_{i,j} \in \mathbb{K}[x]_{v_j}$ is as above and where $R_{i,k} \in \mathbb{K}[x]_{\delta_k}$ is given by

$$R_{i,k} = F_k (S'_{i,k} F'_k \text{div } x^{\delta_k - 1}) \bmod x^{\delta_k}.$$

The cost then follows from Section 2 as well as [Proposition 2](#) in Section 3:

- By fast interpolation, each $Q_{i,j}$ can be computed in $O(M(v_j) \log(v_j))$ operations. Since $\sum_{j \leq s} v_j = n$, this gives a total of $O(\alpha M(n) \log(n))$.
- By fast interpolation, each $S_{i,k}$ can be computed in $O(M(\delta_k) \log(\delta_k))$ operations. Since $\sum_{k \leq t} \delta_k = n$, the total cost for all $S_{i,k}$ and $S'_{i,k}$ is $O(\alpha M(n) \log(n))$.
- Since $S'_{i,k} \in \mathbb{K}[x]_{\delta_k}$ and $F'_k \in \mathbb{K}[x]_{2\delta_k - 1}$, a full product $S'_{i,k} F'_k$ has cost $O(M(\delta_k))$, and thus all $R_{i,k}$ follow from $F_k, S'_{i,k}, F'_k$ in time $O(\alpha M(n))$.
- Applying [Proposition 2](#) to Eq. (27) shows that all the polynomials $\text{Pol}(w'_{j,k})$ can be obtained in time $O(\alpha^{\omega-1} M(n) \log(n))$.

□

6.2. The case of low multiplicities

Here we reduce the Cauchy case to the Vandermonde case, assuming that the multiplicities of x and y are bounded by α . As in Section 5.2, we adapt the reduction of [\[24\]](#), now allowing for repetitions in both x and y .

Proposition 10. Let $\mathbf{x}, \mathbf{y} \in \mathbb{K}^n$ be as in Eqs. (1), (4) and (5). If the multiplicity s of $\mathbf{x}$ and the multiplicity t of $\mathbf{y}$ satisfy $\max\{s, t\} \leq \alpha$, then the problem $\text{LinearSystem}(\mathbb{D}(\mathbf{x}), \mathbb{D}(\mathbf{y}), \alpha)$ can be solved in time $O(\alpha^{\omega-1} M(n) \log^2(n))$. The algorithm is probabilistic of type $P(3n-2, n^2+n)$.

Proof. Given $\mathbf{G}, \mathbf{H} \in \mathbb{K}^{n \times \alpha}$ and $\mathbf{v} \in \mathbb{K}^n$, we are looking for solutions $\mathbf{u}$ to the system $\mathbf{A}\mathbf{u} = \mathbf{v}$, where $\mathbf{A}$ is the $n \times n$ matrix such that $\Delta[\mathbb{D}(\mathbf{x}), \mathbb{D}(\mathbf{y})](\mathbf{A}) = \mathbf{G}\mathbf{H}^t$.

For $k \leq t$, let $M_k = \prod_{\gamma \in \mathbf{y}_k} (x - \gamma)$. Since $\mathbf{y}_j$ is repetition-free, Lemma 3 gives

$$\mathbb{D}(\mathbf{y}_k) = \mathbb{V}(\mathbf{y}_k, \delta_k) \mathbb{X}(M_k) \mathbb{V}(\mathbf{y}_k, \delta_k)^{-1}.$$

Hence

$$\mathbb{D}(\mathbf{y}) = \mathbb{V} \mathbb{X} \mathbb{V}^{-1}, \quad (28)$$

where $\mathbb{V}$ and $\mathbb{X}$ are block-diagonal with respective blocks $\mathbb{V}(\mathbf{y}_k, \delta_k)$ and $\mathbb{X}(M_k)$. For $k \leq t$, let $\mathbf{m}_k \in \mathbb{K}^{\delta_k}$ be the coefficient vector of $-M_k \bmod x^{\delta_k}$. Then $\mathbb{X}(M_k) = \mathbb{Z}_{\delta_k, 0} + \mathbf{m}_k \mathbf{e}_{\delta_k, \delta_k}^t$ and, recalling that $\delta_k^* = \delta_1 + \dots + \delta_k$,

$$\mathbb{X} = \mathbb{Z}_{n, 0} + \sum_{k=1}^t \mathbf{d}_k \mathbf{e}_{n, \delta_k^*}^t, \quad (29)$$

where for $k \leq t$, $\mathbf{d}_k \in \mathbb{K}^n$ is obtained by padding $\mathbf{m}_k$ with δ_{k-1}^* zeros on the top and, if $k \neq t$, with -1 followed by $n - \delta_k^* - 1$ zeros on the bottom.

Defining $\mathbf{B} = \mathbf{A}\mathbb{V}^{-t}$, solving $\mathbf{A}\mathbf{u} = \mathbf{v}$ amounts to solve $\mathbf{B}\mathbf{u}' = \mathbf{v}$ and then, if a solution exists, to compute $\mathbf{u} = \mathbb{V}^{-t}\mathbf{u}'$. To do so in the claimed complexity, we exhibit the Vandermonde-like structure of $\mathbf{B}$ and bound the cost of computing $\mathbf{u}$ and a generator for $\mathbf{B}$. By transposing both sides of Eq. (28), we obtain

$$\mathbf{B} - \mathbb{D}(\mathbf{x}) \mathbf{B} \mathbb{X}^t = \mathbf{G} \mathbf{H}^{tt}, \quad \text{with } \mathbf{H}' = \mathbb{V}^{-1} \mathbf{H}.$$

Now, for $k \leq t$, define the vectors $\mathbf{c}_k = \mathbf{B} \mathbf{e}_{n, \delta_k^*}$ and $\mathbf{c}'_k = \mathbb{D}(\mathbf{x}) \mathbf{c}_k$. Using Eq. (29), the previous identity then becomes

$$\Delta[\mathbb{D}(\mathbf{x}), \mathbb{Z}_{n, 0}^t](\mathbf{B}) = \mathbf{C}' \mathbf{D}^t + \mathbf{G} \mathbf{H}^{tt}, \quad (30)$$

where $\mathbf{C}'$ has columns $\mathbf{c}'_k$ and $\mathbf{D}$ has columns $\mathbf{d}_k$. The matrices $[\mathbf{C}' \ \mathbf{G}]$ and $[\mathbf{D} \ \mathbf{H}']$ thus form a $\mathbb{D}(\mathbf{x}), \mathbb{Z}_{n, 0}^t$ -generator of length $t + \alpha$ for $\mathbf{B}$. Since the operator in Eq. (30) is $\Delta[\mathbb{D}(\mathbf{x}), \mathbb{Z}_{n, \psi}^t]$ with $\psi = 0$, the condition in Eq. (2) is clearly satisfied. Moreover, the multiplicity s of $\mathbf{x}$ is bounded by the length $t + \alpha$ of the above generator. Therefore, once this generator is known, the system $\mathbf{B}\mathbf{u}' = \mathbf{v}$ can be solved by the probabilistic algorithm of type $P(3n-2, n^2+n)$ in Proposition 7; since $t \leq \alpha$, the cost is still in $O(\alpha^{\omega-1} M(n) \log^2(n))$.

It remains to bound the cost computing $\mathbf{u}, \mathbf{C}', \mathbf{D}$ and $\mathbf{H}'$. To do so, we will use the result in Proposition 9 and the reminders of Section 2 on the complexity of polynomial operations:

- By fast transposed interpolation and using the block structure of the matrix $\mathbb{V}$, we can deduce the vector $\mathbf{u}$ from $\mathbf{y}$ and $\mathbf{u}'$ in time $O(M(n) \log(n))$.
- By fast interpolation and using the block structure of the matrix $\mathbb{V}$, we can deduce the matrix $\mathbf{H}'$ from $\mathbf{y}$ and $\mathbf{H}$ in time $O(\alpha M(n) \log(n))$.
- All polynomials M_k (and thus all vectors $\mathbf{d}_k$ and the matrix $\mathbf{D}$) can be constructed from their roots in time $O(M(n) \log(n))$.

In order to bound the cost of computing $\mathbf{C}'$, note that, due to the block structure of $\mathbf{A}$ and $\mathbb{V}^{-t}$, each vector $\mathbf{c}_k$ is in fact given by $\mathbf{c}_k = \mathbf{A}_{j, k} \mathbf{w}_k$, where $\mathbf{w}_k \in \mathbb{K}^{\delta_k}$ is the last column of $\mathbb{V}(\mathbf{y}_k, \delta_k)^{-t}$. Thus, we can compute the vectors $\mathbf{w}_1, \dots, \mathbf{w}_t$ and then solve the multiplication problem of Section 6.1:

- Computing the last column of $\mathbb{V}(\mathbf{y}_k, \delta_k)^{-t}$ takes time $O(M(\delta_k) \log(\delta_k))$, which induces a cost of $O(M(n) \log(n))$ for finding all vectors $\mathbf{w}_k$.
- Knowing all $\mathbf{w}_k$, Proposition 9 shows that all vectors $\mathbf{c}_k$ can be computed in time $O(\alpha^{\omega-1} M(n) \log(n))$. Deducing the vectors $\mathbf{c}'_k$ by t multiplications with $\mathbb{D}(\mathbf{x})$ takes time $O(\alpha n)$, which concludes the proof. $\square$

6.3. The case of high multiplicities

We eventually consider the case of high multiplicities ($\max\{s, t\} > \alpha$), reducing it to the case of low multiplicities ($\max\{s, t\} \leq \alpha$) seen in Section 6.2. Our reduction has cost $O(\alpha^{\omega-1} M(n))$ and thus fits in the requested bound; however, like in the Vandermonde case, randomization is used. Combining the complexity and probability results in the next proposition with those of Proposition 10 concludes the proof of Theorem 3.

Proposition 11. Let $\mathbf{x}$ and $\mathbf{y}$ in $\mathbb{K}^n$ be as in Eqs. (1), (4) and (5). If $\mathbf{x}$ or $\mathbf{y}$ has multiplicity larger than α , one can reduce the problem $\text{LinearSystem}(\mathbb{D}(\mathbf{x}), \mathbb{D}(\mathbf{y}), \alpha)$ to the problem $\text{LinearSystem}(\mathbb{D}(\mathbf{a}), \mathbb{D}(\mathbf{b}), \alpha)$, where $\mathbf{a}, \mathbf{b} \in \mathbb{K}^n$ both have multiplicity at most α and satisfy the constraints of Eqs. (1), (4) and (5). The reduction can be done in time $O(\alpha^{\omega-1} n)$ by a probabilistic algorithm of type $P(2n, 6n^2)$.

Proof. If x has multiplicity $s > \alpha$, we start by reducing the problem $\text{LinearSystem}(\mathbb{D}(x), \mathbb{D}(y), \alpha)$ to the problem $\text{LinearSystem}(\mathbb{D}(a), \mathbb{D}(y), \alpha)$, where $a \in \mathbb{K}^n$ has multiplicity at most α and satisfies the constraints of Eqs. (1) and (5).

We can do this first reduction in time $O(\alpha^{\omega-1}n)$ by a probabilistic algorithm of type $P(n, 3n^2)$, using the same technique as in the proof of Proposition 8. The only modifications to do there are to replace $\mathbb{Z}_{n,\psi}^t$ by $\mathbb{D}(y)$, and to replace the polynomial δ in Eq. (20) by one of the form

$$\prod_{i < j \leq r} (Y_i - Y_j) \times \prod_{i \leq r, j \leq r'} (Y_i - \xi_j) \times \prod_{i \leq r, j \leq n} (Y_i y_j - 1), \quad \text{with } r, r' \leq n, \quad (31)$$

which still has degree bounded by $3n^2$. We do not give more details on this part of the algorithm.

If y has multiplicity $t > \alpha$, it remains to reduce $\text{LinearSystem}(\mathbb{D}(a), \mathbb{D}(y), \alpha)$ to $\text{LinearSystem}(\mathbb{D}(a), \mathbb{D}(b), \alpha)$, where b has multiplicity at most α and satisfies the constraints of Eqs. (4) and (5).

To do so, let G and H in $\mathbb{K}^{n \times \alpha}$ and v in $\mathbb{K}^n$ be given, and recall that we are looking for solutions to the system $Au = v$, where A is such that $\Delta[\mathbb{D}(a), \mathbb{D}(y)](A) = GH^t$. We assume that G and H have full rank, for otherwise one can replace (G, H) by a minimal-length generator, whose matrices then have full rank.

We start by reordering the entries of y to obtain a repetition-free decomposition as in Eq. (6) of Section 2:

$$y' = \begin{bmatrix} y_{\sigma(1)} \\ \vdots \\ y_{\sigma(n)} \end{bmatrix} = \begin{bmatrix} y'_1 \\ \vdots \\ y'_{t'} \end{bmatrix},$$

where y'_i is a vector consisting of μ_i repetitions of the same element ξ_i , so that $n = \mu_1 + \dots + \mu_{t'}$, with $\xi_i \neq \xi_j$ for $i \neq j$ and $\mu_1 \geq \dots \geq \mu_{t'} > 0$.

Let A' and H' be obtained by applying the same reordering to, respectively, the columns of A and the rows of H ; hence, (G, H') is a $\mathbb{D}(a), \mathbb{D}(y')$ -generator for A' . Writing Π for the permutation matrix such that $A' = A\Pi$, we see that every solution to the system $Au = v$ corresponds to a solution $\Pi^t u$ to the system $A'u' = v$, and conversely. Therefore we focus on solving $A'u' = v$.

By construction, the matrices A' and H' admit the following decompositions:

$$A' = [A'_1 \cdots A'_{t'}] \quad \text{and} \quad H' = \begin{bmatrix} H'_1 \\ \vdots \\ H'_{t'} \end{bmatrix},$$

with A'_i in $\mathbb{K}^{n \times \mu_i}$ and H'_i in $\mathbb{K}^{\mu_i \times \alpha}$. Remark then that $A'_i - \mathbb{D}(a) A'_i \mathbb{D}(y'_i) = G H_i'^t$. Now, $\mathbb{D}(y'_i) = \xi_i \mathbb{I}_{\mu_i}$ and, since $\xi_i a_j \neq 1$ for all $i \leq t'$ and all $j \leq n$, all matrices $\mathbb{I}_n - \xi_i \mathbb{D}(a)$ are invertible. We thus obtain the equalities

$$A'_i = (\mathbb{I}_n - \xi_i \mathbb{D}(a))^{-1} G H_i'^t \quad \text{for } 1 \leq i \leq t'. \quad (32)$$

We will use dense matrix methods to reduce the number of non-zero entries in A' and H' , while maintaining a Cauchy-like structure. Let thus τ be such that $\mu_\tau > \alpha \geq \mu_{\tau+1}$.

- For $1 \leq i \leq \tau$, we have $\mu_i > \alpha$. We let $J_i \subset \{1, \dots, \mu_i\}$ and $E_i \in \mathbb{K}^{\mu_i \times \mu_i}$ be the index set and the matrix obtained by applying Lemma 9 to H'_i . Let also $r_i = |J_i| = \text{rank}(H'_i)$, so that $r_i \leq \alpha$. Eq. (32) then gives

$$E_i H'_i = \begin{bmatrix} H''_i \\ 0 \end{bmatrix} \quad \text{and} \quad A'_i E_i^t = \begin{bmatrix} A''_i & 0 \end{bmatrix},$$

where H''_i and A''_i have respective sizes $(r_i \times \alpha)$ and $(n \times r_i)$, and H''_i consists of the rows of H'_i indexed by J_i .

- For $\tau < i \leq t'$, we have $\mu_i \leq \alpha$. We let $E_i = \mathbb{I}_{\mu_i}$, $H''_i = H'_i$ and $A''_i = A'_i$.

Let E be the block-diagonal matrix having $E_1, \dots, E_{t'}$ on the diagonal. Hence, $H'' = EH'$ consists of the matrices H''_i , interleaved by blocks of zeros when $i \leq \tau$; the same holds for $A'' = A' E^t$, considering columns instead of rows. Besides, since each $\mathbb{D}(y'_i)$ is a homothety matrix, it commutes with E_i^t ; we deduce that (G, H'') is a $\mathbb{D}(a), \mathbb{D}(y')$ -generator for A'' . Finally, since E is invertible, every solution to the system $A'u' = v$ corresponds to a solution $E^{-t} u'$ to the system $A''u'' = v$, and conversely. Therefore, solving the problem $A''u'' = v$ is enough.

We do so by exhibiting the Cauchy-like structure of A'' for a modified displacement operator whose second diagonal component has lower multiplicity. Define $y'' \in \mathbb{K}^n$ by replacing, for $i \leq \tau$, the last $\mu_i - r_i$ entries of y'_i by new values taken from $\mathbb{K}$. Due to the presence of corresponding blocks of zeros in H'' and A'' , the matrices G and H'' now form a $\mathbb{D}(a), \mathbb{D}(y'')$ -generator for A'' .

With $r = \sum_{i \leq \tau} (\mu_i - r_i)$, suppose that the new values $z_1, \dots, z_r$ inserted in y'' are pairwise distinct and that none of them belongs to y' or satisfies $a_i z_j = 1$. Then, y'' satisfies the constraint of Eq. (5) and has multiplicity at most α (since all r_i are at most α).

It remains to reorder the entries of y'' to obtain a vector b that also satisfies Eq. (1), and let Σ be the permutation matrix such that $b = \Sigma y''$. Defining $F = \Sigma H''$, it follows that (G, F) is a $\mathbb{D}(a), \mathbb{D}(b)$ -generator for the matrix $B = A'' \Sigma^t$. To conclude, every solution to $Bw = v$ corresponds to a solution $u'' = \Sigma^t w$ to $A'' u'' = v$, and conversely.

We have thus reduced solving the system $Au = v$ to solving $Bw = v$ (and then, if a solution exists, to computing $u = \Pi E^t \Sigma^t w$), while providing a generator for B with respect to the Cauchy-like structure $\mathbb{D}(a), \mathbb{D}(b)$.

The complexity analysis of this second reduction is the same as in the proof of Proposition 8 and yields a cost in $O(\alpha^{\omega-1}n)$ as required.

For the probability analysis, remark that the values $z_1, \dots, z_r$ satisfy our requirements if they do not cancel the polynomial

$$\gamma = \prod_{i < j \leq r} (Z_i - Z_j) \times \prod_{i \leq r, j \leq t'} (Z_i - \xi_j) \times \prod_{i \leq n, j \leq r} (a_i Z_j - 1),$$

which has degree at most $3n^2$, since r and t' are bounded by n . Hence, the algorithm for the second reduction has type $P(n, 3n^2)$. Recalling that the first reduction was also by an algorithm of type $P(n, 3n^2)$ concludes the proof. $\square$

Acknowledgments and notes

We thank E. Kaltofen, G. Labahn, M. Morf, V. Y. Pan, B. Salvy, A. Storjohann and G. Villard for useful discussions and comments. An earlier version of this article (lacking in particular the discussion of the Cauchy case) was published in [8].

The first author was supported by the French National Agency for Research (ANR Project “Gecko”) and the Microsoft Research-INRIA Joint Centre. The third author was supported by NSERC (Canada) and by the Canada Research Chair Program.

Appendix

All along this paper, for matrices of Vandermonde or Cauchy type, we assumed that the diagonal matrices used in the operators are as in Eq. (1). One can always reduce to this situation by permuting the entries of these diagonal matrices. We discuss here the cost of this operation.

Given $a = [a_1, \dots, a_n]^t$ in $\mathbb{K}^n$, we actually study the cost of finding a repetition-free decomposition of a , that is, a permutation that puts a into the form of Eq. (6) of Section 2. As seen in Lemma 11, the further cost of converting to the representation of Eq. (1) is then a mere $O(n)$.

If the base field $\mathbb{K}$ is endowed with an order $<$, a sorting algorithm solves the problem in $O(n \log(n))$ comparisons. However, if we do not allow ordering the elements of $\mathbb{K}$, it is far less clear how to achieve a similar complexity. Our solution involves polynomial arithmetic, and bears similarities with techniques used in both elliptic curve factorization algorithms [36] or factorization into coprimes [5].

Proposition 12. *Let $a = [a_1, \dots, a_n]^t$ be in $\mathbb{K}^n$. One can compute a repetition-free decomposition of a in $O(M(n) \log^3(n))$ operations.*

Before giving the proof of this proposition, we start by two lemmas on sequence manipulations. As it turns out, one of the main difficulties is the following (see Lemma 19): given two repetition-free sequences $[c_1, \dots, c_r]$ and $[d_1, \dots, d_r]$, such that $d_i = c_{\sigma(i)}$ for some permutation σ of $\{1, \dots, r\}$, how to recover σ , if ordering elements is not allowed?

Lemma 18. *Let $c = [c_1, \dots, c_s]$ and $d = [d_1, \dots, d_t]$ be repetition-free vectors with entries in $\mathbb{K}$. One can compute sequences of integers $\ell_1 < \dots < \ell_u$ and $m_1 < \dots < m_u$ such that*

$$\{c_1, \dots, c_s\} \cap \{d_1, \dots, d_t\} = \{c_{\ell_1}, \dots, c_{\ell_u}\} = \{d_{m_1}, \dots, d_{m_u}\}$$

in $O(M(q) \log(q))$ operations, with $q = \max\{s, t\}$.

Proof. Let $P = \prod_{i \leq s} (x - c_i)$. Remark that $\alpha \in \mathbb{K}$ belongs to $\{c_1, \dots, c_s\}$ if and only if $P(\alpha) = 0$; hence, evaluating P at d , we obtain the indices m_i as those for which $P(d_{m_i}) = 0$. Similarly, one obtains the indices ℓ_i by evaluating $Q = \prod_{i \leq t} (x - d_i)$ at c . All operations fit into the $O(M(q) \log(q))$ bound. $\square$

Lemma 19. *Let $c = [c_1, \dots, c_r]$ and $d = [d_1, \dots, d_r]$ be repetition-free vectors with entries in $\mathbb{K}$, such that $\{c_1, \dots, c_r\} = \{d_1, \dots, d_r\}$ (as sets). One can compute the unique permutation σ of $\{1, \dots, r\}$ such that $d_i = c_{\sigma(i)}$ in $O(M(n) \log^2(n))$ operations.*

Proof. Defining $s = \lceil r/2 \rceil$ and $t = r - s$, we split d into the subsequences $d' = [d_1, \dots, d_s]$ and $d'' = [d_{s+1}, \dots, d_r]$ of respective lengths s and t . We then compute the polynomial $P = \prod_{i \leq s} (x - d_i)$ and evaluate it at c . Let $\ell_1 < \dots < \ell_s$ be the indices of the entries of c where P vanishes and $m_1 < \dots < m_t$ be those where P is non-zero. We deduce that

$$\{c_{\ell_1}, \dots, c_{\ell_s}\} = \{d_1, \dots, d_s\} \quad \text{and} \quad \{c_{m_1}, \dots, c_{m_t}\} = \{d_{s+1}, \dots, d_r\}.$$

We can then proceed recursively, on $[c_{\ell_1}, \dots, c_{\ell_s}]$ and d' on the one hand, and $[c_{m_1}, \dots, c_{m_t}]$ and d'' on the other hand. This gives us permutations ρ of $\{1, \dots, s\}$ and τ of $\{1, \dots, t\}$, such that $d_i = c_{\ell_{\rho(i)}}$ and $d_{j+s} = c_{m_{\tau(j)}}$ hold for $i \leq s$ and $j \leq t$. We deduce σ , as $\sigma(i) = \ell_{\rho(i)}$ for $i \leq s$ and $\sigma(i) = m_{\tau(i-s)}$ for $i > s$.

The cost of computing P and re-evaluating it is in $O(M(r) \log(r))$ operations in $\mathbb{K}$; the extra cost is $O(r)$ bookkeeping operations. Our claim follows from the super-additivity of the function M . $\square$

We can now prove our claim on the cost of finding a repetition-free decomposition. We give a recursive algorithm that, given the vector $\mathbf{a} = [a_1, \dots, a_n]$, computes a permutation σ of $\{1, \dots, n\}$ such that

$$\begin{bmatrix} a_{\sigma(1)} \\ \vdots \\ a_{\sigma(n)} \end{bmatrix} = \begin{bmatrix} \mathbf{a}_1 \\ \vdots \\ \mathbf{a}_u \end{bmatrix},$$

where the vector $\mathbf{a}_i$ consists of λ_i repetitions of an element α_i , the α_i being pairwise distinct. Remark that this specification is a slight relaxation of the definition given in Section 2, as we do not require that the multiplicities be sorted. The extra cost for sorting them is $O(n)$ integer operations, using bucket sorting.

Given $\mathbf{a} = [a_1, \dots, a_n]$, we define $\ell = \lceil n/2 \rceil$ and $m = n - \ell$ and recursively call the algorithm on $[a_1, \dots, a_\ell]$ and $[a_{\ell+1}, \dots, a_n]$: we obtain permutations ρ' of $\{1, \dots, \ell\}$ and ρ'' of $\{1, \dots, m\}$ such that

$$\begin{bmatrix} a_{\rho'(1)} \\ \vdots \\ a_{\rho'(\ell)} \end{bmatrix} = \begin{bmatrix} \mathbf{a}'_1 \\ \vdots \\ \mathbf{a}'_s \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} a_{\ell+\rho''(1)} \\ \vdots \\ a_{\ell+\rho''(m)} \end{bmatrix} = \begin{bmatrix} \mathbf{a}''_1 \\ \vdots \\ \mathbf{a}''_t \end{bmatrix}, \quad (33)$$

where the vector $\mathbf{a}'_i$ (resp. $\mathbf{a}''_j$) consists of v_i (resp. μ_j) repetitions of an element α'_i (resp. α''_j), the α'_i and α''_j being pairwise distinct.

We now have to determine the common elements between the lists $[\alpha'_1, \dots, \alpha'_s]$ and $[\alpha''_1, \dots, \alpha''_t]$. Using Lemma 18, we compute in $O(M(n) \log(n))$ operations the sequences of indices $\ell_1 < \dots < \ell_r$ and $m_1 < \dots < m_r$ such that

$$\{\alpha'_1, \dots, \alpha'_s\} \cap \{\alpha''_1, \dots, \alpha''_t\} = \{\alpha'_{\ell_1}, \dots, \alpha'_{\ell_r}\} = \{\alpha''_{m_1}, \dots, \alpha''_{m_r}\}.$$

Using $O(n)$ extra operations, we then obtain the complementary sequences $\ell'_1 < \dots < \ell'_{s-r}$ and $m'_1 < \dots < m'_{t-r}$, such that

$$\{\ell'_1, \dots, \ell'_{s-r}\} = \{1, \dots, s\} - \{\ell_1, \dots, \ell_r\} \quad \text{and} \quad \{m'_1, \dots, m'_{t-r}\} = \{1, \dots, t\} - \{m_1, \dots, m_r\}.$$

Since the sequences $[\alpha'_{\ell_1}, \dots, \alpha'_{\ell_r}]$ and $[\alpha''_{m_1}, \dots, \alpha''_{m_r}]$ are repetition-free, there exists a permutation σ of $\{1, \dots, r\}$ such that $\alpha'_{\ell_i} = \alpha''_{m_{\sigma(i)}}$ for $i \leq r$; by Lemma 19, we can compute σ in $O(M(n) \log^2(n))$ operations.

Knowing these subsequences, it remains to reorder the vectors in Eq. (33). We first interleave the entries appearing in both vectors

$$\alpha'_{\ell_1} (v_{\ell_1} \text{ times}), \alpha''_{m_{\sigma(1)}} (\mu_{m_{\sigma(1)}} \text{ times}), \dots, \alpha'_{\ell_r} (v_{\ell_r} \text{ times}), \alpha''_{m_{\sigma(r)}} (\mu_{m_{\sigma(r)}} \text{ times}),$$

followed by the entries appearing only in the first vector,

$$\alpha'_{\ell'_1} (v_{\ell'_1} \text{ times}), \dots, \alpha'_{\ell'_{s-r}} (v_{\ell'_{s-r}} \text{ times}),$$

and by those appearing only in the first vector,

$$\alpha''_{m'_1} (\mu_{m'_1} \text{ times}) \dots, \alpha''_{m'_{t-r}} (\mu_{m'_{t-r}} \text{ times}).$$

The permutation σ that actually puts $\mathbf{a}$ into the above order is readily deduced for $O(n)$ operations. The overall cost is thus $O(M(n) \log^2(n))$, plus that of the two recursive calls; the estimate given in Proposition 12 follows.

References

- [1] B. Beckermann, A reliable method for computing M-Padé approximants on arbitrary staircases, *J. Comput. Appl. Math.* 40 (1) (1992) 19–42.
- [2] B. Beckermann, G. Labahn, A uniform approach for the fast computation of matrix-type Padé approximants, *SIAM J. Matrix Anal. Appl.* 15 (3) (1994) 804–823.
- [3] B. Beckermann, G. Labahn, Fraction-free computation of matrix rational interpolants and matrix GCDs, *SIAM J. Matrix Anal. Appl.* 22 (1) (2000) 114–144.
- [4] M. Ben-Or, P. Tiwari, A deterministic algorithm for sparse multivariate polynomial interpolation, in: 20th Annual ACM Symp. Theory Comp., ACM Press, 1988, pp. 301–309d.
- [5] D.J. Bernstein, Factoring into coprimes in essentially linear time, *J. Algorithms* 54 (1) (2005) 1–30.
- [6] D. Bini, V.Y. Pan, Polynomial and Matrix Computations, in: *Fundamental Algorithms*, vol. 1, Birkhäuser, 1994.
- [7] R.R. Bitmead, B.D.O. Anderson, Asymptotically fast solution of Toeplitz and related systems of linear equations, *Linear Algebra Appl.* 34 (1980) 103–116.
- [8] A. Bostan, C.-P. Jeannerod, É. Schost, Solving Toeplitz- and Vandermonde-like linear systems with large displacement rank, in: *ISSAC'07*, ACM Press, 2007, pp. 33–40.
- [9] A. Bostan, G. Lecerf, É. Schost, Tellegen's principle into practice, in: *ISSAC'03*, ACM Press, 2003, pp. 37–44.
- [10] J. Canny, E. Kaltofen, Y. Lakshman, Solving systems of non-linear polynomial equations faster, in: *ISSAC'89*, ACM Press, 1989, pp. 121–128.
- [11] D.G. Cantor, E. Kaltofen, On fast multiplication of polynomials over arbitrary algebras, *Acta Inform.* 28 (7) (1991) 693–701.
- [12] Z. Chen, V.Y. Pan, An efficient solution for Cauchy-like systems of linear equations, *Comput. Math. Appl.* 48 (2004) 529–537.
- [13] D. Coppersmith, S. Winograd, Matrix multiplication via arithmetic progressions, *J. Symbolic. Comput.* 9 (3) (1990) 251–280.
- [14] R.A. DeMillo, R.J. Lipton, A probabilistic remark on algebraic program testing, *Inform. Process. Lett.* 7 (4) (1978) 193–195.
- [15] J.-G. Dumas, T. Gautier, C. Pernet, Finite field linear algebra subroutines, in: *ISSAC'02*, ACM Press, 2002, pp. 63–74.
- [16] W. Eberly, M. Giesbrecht, P. Giorgi, A. Storjohann, G. Villard, Solving sparse rational linear systems, in: *ISSAC'06*, ACM Press, 2006, pp. 63–70.

- [17] W. Eberly, M. Giesbrecht, P. Giorgi, A. Storjohann, G. Villard, Faster inversion and other black box matrix computations using efficient block projections, in: ISSAC'07, ACM Press, 2007, pp. 143–150.
- [18] T. Finck, G. Heinig, K. Rost, An inversion formula and fast algorithms for Cauchy-Vandermonde matrices, *Linear Algebra Appl.* 183 (1) (1993) 179–191.
- [19] S. Gao, V.M. Rodrigues, J. Stroomer, Gröbner basis structure of finite sets of points, Preprint, 2003.
- [20] M. Gasca, T. Sauer, Polynomial interpolation in several variables, *Adv. Comput. Math.* 12 (4) (2000) 377–410.
- [21] J. von zur Gathen, J. Gerhard, *Modern Computer Algebra*, second ed., Cambridge University Press, 2003.
- [22] J. von zur Gathen, V. Shoup, Computing Frobenius maps and factoring polynomials, *Comput. Complexity* 2 (3) (1992) 187–224.
- [23] P. Giorgi, C.-P. Jeannerod, G. Villard, On the complexity of polynomial matrix computations, in: ISSAC'03, ACM Press, 2003, pp. 135–142.
- [24] I.C. Gohberg, V. Olshevsky, Complexity of multiplication with vectors for structured matrices, *Linear Algebra Appl.* 202 (1994) 163–192.
- [25] O.H. Ibarra, S. Moran, R. Hui, A generalization of the fast LUP matrix decomposition algorithm and applications, *J. Algorithms* 3 (1) (1982) 45–56.
- [26] T. Kailath, S.Y. Kung, M. Morf, Displacement ranks of matrices and linear equations, *J. Math. Anal. Appl.* 68 (2) (1979) 395–407.
- [27] E. Kaltofen, Greatest common divisors of polynomials given by straight-line programs, *J. ACM* 35 (1) (1988) 231–264.
- [28] E. Kaltofen, Asymptotically fast solution of Toeplitz-like singular linear systems, in: ISSAC'94, ACM Press, 1994, pp. 297–304.
- [29] E. Kaltofen, Analysis of Coppersmith's block Wiedemann algorithm for the parallel solution of sparse linear systems, *Math. Comp.* 64 (210) (1995) 777–806.
- [30] E. Kaltofen, Y. Lakshman, Improved sparse multivariate polynomial interpolation algorithms, in: ISSAC'88, in: LNCS, vol. 358, Springer Verlag, 1988, pp. 467–474.
- [31] E. Kaltofen, D. Saunders, On Wiedemann's method of solving sparse linear systems, in: AAECC-9, in: LNCS, vol. 539, Springer Verlag, 1991, pp. 29–38.
- [32] I. Kaporin, The aggregation and cancellation techniques as a practical tool for faster matrix multiplication, *Theoret. Comput. Sci.* 315 (2–3) (2004) 469–510.
- [33] G. Labahn, D.K. Choi, S. Cabay, The inverses of block Hankel and block Toeplitz matrices, *SIAM J. Comput.* 19 (1) (1990) 98–123.
- [34] J. Laderman, V.Y. Pan, X.-H. Sha, On practical algorithms for accelerated matrix multiplication, *Linear Algebra Appl.* 162–164 (1992) 557–588.
- [35] D. Lazard, Ideal bases and primary decomposition: The case of two variables, *J. Symbolic. Comput.* 1 (1985) 261–270.
- [36] P.L. Montgomery, Speeding the Pollard and elliptic curve methods of factorization, *Math. Comp.* 48 (177) (1987) 243–264.
- [37] M. Morf, Fast algorithms for multivariable systems, Ph.D. Thesis, Stanford University, 1974.
- [38] M. Morf, Doubling algorithms for Toeplitz and related equations, in: IEEE Conference on Acoustics, Speech, and Signal Processing, 1980, pp. 954–959.
- [39] T. Mulders, On short multiplications and divisions, *Appl. Algebra Engrg. Comm. Comput.* 11 (1) (2000) 69–88.
- [40] M. Nüsken, M. Ziegler, Fast multipoint evaluation of bivariate polynomials, in: ESA 2004, in: LNCS, vol. 3222, Springer, 2004, pp. 544–555.
- [41] V.Y. Pan, On some computations with dense structured matrices, in: ISSAC'89, ACM Press, 1989, pp. 34–42.
- [42] V.Y. Pan, On computations with dense structured matrices, *Math. Comp.* 55 (191) (1990) 179–190.
- [43] V.Y. Pan, Parametrization of Newton's iteration for computations with structured matrices and applications, *Comput. Math. Appl.* 24 (3) (1992) 61–75.
- [44] V.Y. Pan, Nearly optimal computations with structured matrices, in: SODA'00, ACM Press, 2000, pp. 953–962.
- [45] V.Y. Pan, *Structured Matrices and Polynomials*, Birkhäuser Boston Inc., 2001.
- [46] V.Y. Pan, X. Wang, Inversion of displacement operators, *SIAM J. Matrix Anal. Appl.* 24 (3) (2003) 660–677.
- [47] V.Y. Pan, A. Zheng, Superfast algorithms for Cauchy-like matrix computations and extensions, *Linear Algebra Appl.* 310 (2000) 83–108.
- [48] V.Y. Pan, A. Zheng, M. Abu-Tabaneh, Z. Chen, S. Providence, Superfast computations with singular structured matrices over abstract fields, in: *Computer algebra in scientific computing: CASC'99*, Springer, 1999, pp. 323–338.
- [49] A. Schönhage, V. Strassen, Schnelle Multiplikation großer Zahlen, *Computing* 7 (1971) 281–292.
- [50] J.T. Schwartz, Fast probabilistic algorithms for verification of polynomial identities, *J. ACM* 27 (4) (1980) 701–717.
- [51] A. Storjohann, Algorithms for matrix canonical forms, Ph.D. Thesis, ETH, Zürich, 2000.
- [52] A. Storjohann, Notes on computing minimal approximant bases, Technical report, Symbolic Computation Group, University of Waterloo, 2006.
- [53] V. Strassen, Gaussian elimination is not optimal, *Numer. Math.* 13 (1969) 354–356.
- [54] M. Van Barel, A. Buihel, A general module theoretic framework for vector M-Padé and matrix rational interpolation, *Numer. Algorithms* 3 (1992) 451–461.
- [55] R. Zippel, Probabilistic algorithms for sparse polynomials, in: EUROSAM' 79, in: LNCS, vol. 72, Springer Verlag, 1979.
- [56] R. Zippel, Interpolating polynomials from their values, *J. Symbolic. Comput.* 9 (3) (1990) 375–403.